

Homotopy Hubbard trees for post-singularly finite exponential maps

DAVID PFRANG[†], MICHAEL ROTHGANG[‡] and DIERK SCHLEICHER[§]

[†] *prognostica GmbH, Berliner Platz 6, 97080 Würzburg, Germany*
(e-mail: david.pfrang@gmx.de)

[‡] *Institut für Mathematik, Humboldt-Universität zu Berlin,*
Unter den Linden 6, 10099 Berlin, Germany
(e-mail: rothgami@math.hu-berlin.de)

[§] *Aix-Marseille Université and CNRS, UMR 7373,*
Institut de Mathématiques de Marseille, 163 Avenue de Luminy,
13009 Marseille, France
(e-mail: dierk.schleicher@univ-amu.fr)

(Received 10 April 2019 and accepted in revised form 10 August 2021)

Abstract. We extend the concept of a Hubbard tree, well established and useful in the theory of polynomial dynamics, to the dynamics of transcendental entire functions. We show that Hubbard trees in the strict traditional sense, as invariant compact trees embedded in \mathbb{C} , do not exist even for post-singularly finite exponential maps; the difficulty lies in the existence of asymptotic values. We therefore introduce the concept of a *homotopy Hubbard tree* that takes care of these difficulties. Specifically for the family of exponential maps, we show that every post-singularly finite map has a homotopy Hubbard tree that is unique up to homotopy, and that post-singularly finite exponential maps can be classified in terms of homotopy Hubbard trees, using a transcendental analogue of Thurston’s topological characterization theorem of rational maps.

Key words: Hubbard trees, exponential map, classification, kneading sequence, Thurston theory

2020 Mathematics subject classification: 37F20, 37F10, 37F46 (Primary); 37B10, 37E25 (Secondary)

1. Introduction

The purpose of this paper is to establish the concept of Hubbard trees for post-singularly finite exponential maps, that is, maps of the form $E_\lambda : z \mapsto \lambda \exp(z)$ for which the orbit of the singular value 0 is finite. While Hubbard trees are known to be an extremely useful

concept for the dynamics of polynomials, so far they have not been introduced for the dynamics of entire functions. Part of the problem is that asymptotic values might prevent the existence of Hubbard trees in the strict traditional sense, as invariant compact trees in the plane, and they do so even for exponential maps. We therefore introduce a modified concept that we call *homotopy Hubbard trees* and demonstrate their use for the study of exponential maps. Exponential maps have received a lot of attention over the years, and an understanding of exponential dynamics has often proved to be useful for the study of much more general classes of transcendental entire functions.

This project goes back to the Bachelor's thesis of the second named author [R]. Previously, Hubbard trees were defined only for post-critically finite polynomials in [DH1]. The Hubbard tree is the unique smallest tree embedded in the filled-in Julia set that contains all critical points of the polynomial and is forward invariant under its dynamics (and in addition normalized on bounded Fatou components). It turns out that this definition does not immediately generalize to exponential maps: because of the existence of an asymptotic value, an exponential map cannot have an exactly forward invariant tree containing the post-singular set. We solve this by only requiring the tree to be invariant up to homotopy relative to the post-singular set. Because of the relaxed invariance condition, we call the resulting tree a *homotopy Hubbard tree*. In analogy to polynomials, the underlying abstract graph of a homotopy Hubbard tree, together with the dynamics of its self-map on the vertices and a finite amount of extra information, is a useful combinatorial object. Specifically for exponential maps, we call this an *abstract exponential Hubbard tree*. The main results of our paper can be phrased as follows.

MAIN THEOREM 1.1. *Every post-singularly finite exponential map has a homotopy Hubbard tree, and this tree is unique up to homotopy relative to the post-singular set. Moreover, for every abstract exponential Hubbard tree, there is a unique post-singularly finite exponential map realizing it.*

Hubbard trees are a convenient tool to read off the dynamical properties of the map under consideration (an example is given below). Additionally, abstract Hubbard trees provide a way to define the notion of core entropy (introduced by Thurston [T] for post-critically finite polynomials) for post-singularly finite parameters and to study properties of the core entropy function on the parameter space. This has been done for quadratic polynomials in [DS] and independently in [Ti]. While entropy of transcendental mappings is always infinite [BFP], core entropy of post-singularly finite exponential mappings is always bounded by $\log 2$ [Ha]. This provides relevant information to the dynamics, in a similar sense that entropy of degree d polynomials always equals $\log d$, while core entropy allows to distinguish different polynomial dynamical systems.

We want to point out that the construction of homotopy Hubbard trees in this paper is quite explicit. Given the external address of a dynamic ray landing at the singular value, we show how to construct the abstract exponential Hubbard tree of the corresponding map algorithmically.

Post-singularly finite exponential maps have already been classified in [LSV] in terms of the external addresses of the dynamic rays landing at the singular value. Compared to this

previous classification, our result has the advantage that it establishes a bijection between the class of maps and the instances of the combinatorial model: while several addresses might correspond to the same exponential map, there is a natural bijection between abstract exponential Hubbard trees and post-singularly finite exponential maps. Also, some dynamical properties of the map are obtained more easily from the abstract Hubbard tree than from an external address. For example, the abstract Hubbard tree contains information about all periodic branch points of the Julia set, while it is computationally intensive to determine them from the external address.

Although we restrict to exponential maps in this article, our results have wider implications. Indeed, the insights gained here, together with recent findings on the structure of the escaping sets of entire functions with bounded post-singular set established in [BR], have led to a general theory of homotopy Hubbard trees for all post-singularly finite transcendental entire functions. This is the content of the first named author's PhD thesis [Pf]. The first part of this thesis is available at [PPS].

1.1. *Background and relevance.* The dynamics of a holomorphic map is controlled to a large extent by the orbits of its *singular values* (see [M] for a general introduction to holomorphic dynamical systems and [S2] specifically for transcendental entire functions). Singular values are points in the range of the function that do not have a neighborhood on which all branches of the inverse function are well defined and biholomorphic. For a polynomial $f: \mathbb{C} \rightarrow \mathbb{C}$, every singular value is a *critical value*, that is, a point $w \in \mathbb{C}$ such that $f'(z) = 0$ for some $z \in f^{-1}(w)$. For a transcendental entire function $f: \mathbb{C} \rightarrow \mathbb{C}$, a singular value can also be an asymptotic value or a limit point of asymptotic and critical values. A point $w \in \mathbb{C}$ is called an *asymptotic value* if there exists a curve $\gamma: [0, +\infty) \rightarrow \mathbb{C}$ satisfying $\lim_{t \rightarrow +\infty} \gamma(t) = \infty$ and $\lim_{t \rightarrow +\infty} f(\gamma(t)) = w$. In any parameter space of holomorphic functions, the easiest maps to understand are those for which all singular values have finite orbits. These functions are called *post-singularly finite* (or *post-critically finite* in the case of polynomials). Not only are they the dynamically simplest maps, but also, in many cases, the most important ones for the structure of the parameter space. Figuratively, one reason for studying them is because 'it is easier to search for your lost keys under a lamp post, but there are lamp posts at every important street intersection so they are very helpful to find your way around'. We explain this by way of analogy to the simplest and best-studied space of polynomial maps, the space of quadratic polynomials $p_c: z \mapsto z^2 + c$ for $c \in \mathbb{C}$. Its connectedness locus is the *Mandelbrot set*

$$\mathcal{M} := \{c \in \mathbb{C} : \text{the Julia set } \mathcal{J}(p_c) \text{ is connected}\}.$$

All branch points of the Mandelbrot set (in a sense made precise in [S1, Theorem 3.1]) are post-critically finite parameters, and under the assumption that the Mandelbrot set is locally connected (the famous MLC conjecture), these branch points completely describe its topology. In the space of exponential maps, post-singularly finite parameters play a comparable role.

Naturally, much work has gone into investigating the dynamics of post-singularly finite holomorphic functions. For rational maps, a deep characterization theorem by

Thurston [DH2] about branched self-covers of the sphere that arise from rational maps has made strong classification results possible. Post-critically finite polynomials [P] and, quite recently, post-critically finite Newton maps [LMS1, LMS2] have been completely classified in terms of custom-tailored combinatorial models using Thurston's theorem.

The fundamental idea of classifying post-critically finite polynomials in terms of their Hubbard trees originated from [DH1]. Their program has been carried through in greater generality in [BFH] and in full generality in [P]. The following finite amount of combinatorial information is sufficient to completely describe the dynamics of a post-critically finite polynomial:

- the graph structure of its Hubbard tree (without an embedding into the complex plane);
- the dynamics of the polynomial on the finite set of vertices of the tree;
- the degrees of the critical points of the polynomial; and
- for each vertex, certain (combinatorial) angles between the edges incident to this vertex.

This combinatorial data distinguish all post-critically finite polynomials. Conversely, if we start with an *expansive* (in a sense, defined in Definition 3.8) dynamical tree with consistent degree and angle information (this finite combinatorial object is known as an *abstract Hubbard tree*), there exists a post-critically finite polynomial realizing this tree, and this map is unique up to affine conjugation. Therefore, there exists a natural bijection between post-critically finite polynomials (up to affine equivalence) and abstract Hubbard trees. In the spirit of this classification, we establish a natural bijection between post-singularly finite exponential maps and abstract exponential Hubbard trees in §6.

1.2. *Structure of the article.* In §2, we give a combinatorial description of the escaping set of exponential maps and its dynamics in terms of external addresses. Path-connected components of the escaping set are called *dynamic rays*; topologically, they are arcs terminating at ∞ (in \mathbb{C}). The pre-images of a dynamic ray landing at the singular value form the boundaries of the *dynamic partition*. Itineraries with respect to this partition distinguish (pre-)periodic points and determine which dynamic rays have a common landing point.

In §3, we motivate the concept of homotopy Hubbard trees and give a precise definition of homotopy Hubbard trees for exponential maps. Some technical results on homotopies of embedded trees needed in this paper are discussed.

In §4, we show that a homotopy Hubbard tree can be chosen to not intersect dynamic rays landing at post-singular points. Furthermore, the itinerary of the singular value with respect to a dynamical partition determines the graph structure and the dynamics of the tree. Together, these two facts imply uniqueness of homotopy Hubbard trees.

Section 5 deals with the construction of homotopy Hubbard trees. The combinatorial description of dynamic rays from §2 is used to show that every triple of post-singular points is separated by dynamic rays landing at a common (pre-)periodic point. Connecting the post-singular set without intersecting these separating rays yields a homotopy Hubbard tree.

Finally, in §6, we give a combinatorial classification of post-singularly finite exponential maps in terms of abstract Hubbard trees using the transcendental analogue of Thurston's

characterization theorem for rational maps established in [HSS]. In §7, we discuss possible extensions and generalizations.

1.3. *Notation and terminology.* The complex plane is denoted by \mathbb{C} and the Riemann sphere by $\overline{\mathbb{C}}$. We denote the trace of a curve g by $\text{Tr}(g)$. An *arc* is a simple (injective) curve and a *Jordan curve* is a simple closed curve. A (*pre-*)*periodic point* is a point that is either periodic or pre-periodic under iteration of the map under consideration. We use the term ‘pre-periodic’ in the sense of ‘strictly pre-periodic’, that is, excluding the periodic case.

2. Background on the dynamics of exponential maps

In the following, let $f: \mathbb{C} \rightarrow \mathbb{C}$ denote an entire function. The *singular set* $S(f)$ of f is defined to be the closure of the set of critical and asymptotic values of f . It is the smallest subset of \mathbb{C} such that the restriction of f to $\mathbb{C} \setminus f^{-1}(S(f))$ is a covering map (see e.g. [GK, Lemma 1.1]). The *post-singular set* $P(f)$ of f is given by

$$P(f) := \overline{\bigcup_{n \geq 0} f^{on}(S(f))}.$$

We call the function f *post-singularly finite* if $|P(f)| < \infty$.

Entire functions have an important forward and backward invariant set, the *set of escaping points*; it is often more important than Fatou and Julia sets because it is never empty and never all of \mathbb{C} (see [E1]), so it provides a non-trivial dynamically invariant decomposition of \mathbb{C} . Here is the formal definition.

Definition 2.1. (Escaping set) The *escaping set* of f is given by

$$I(f) := \{z \in \mathbb{C}: f^{on}(z) \rightarrow \infty\}.$$

In [RRRS], it was shown that for a large class of transcendental entire functions (finite compositions of functions of bounded type and finite order), the escaping set is organized in the form of disjoint dynamic rays, which are certain arcs consisting of escaping points that terminate at ∞ .

Definition 2.2. (Ray tails, dynamic rays, and landing points) A *ray tail* of f is an injective curve $g: [\tau, \infty) \rightarrow I(f)$ ($\tau \in \mathbb{R}$) such that for each $n \in \mathbb{N}$, the restriction $f^{on}|_g$ is injective, $\lim_{t \rightarrow \infty} f^{on}(g(t)) = \infty$, and $f^{on}(g(t)) \rightarrow \infty$ as $n \rightarrow \infty$ uniformly in t .

A *dynamic ray* of f is a maximal (in the sense of inclusions of sets) injective curve $g: (\tau_0, \infty) \rightarrow I(f)$ such that $g|_{[\tau, \infty)}$ is a ray tail for every $\tau > \tau_0$. We say that the dynamic ray g *lands* at the point $z \in \mathbb{C}$ if $\lim_{t \rightarrow \tau_0} g(t) = z$; in this case, z is called the *landing point* of g .

We call the dynamic ray g *periodic* if $f^{on}(\text{Tr}(g)) \subseteq \text{Tr}(g)$ for some $n \in \mathbb{N}$; we call it *pre-periodic* if some forward iterate $f^{ok}(g)$ of the ray g (which, by definition, is again a dynamic ray) is periodic, but not g itself.

In this paper, we focus on the case of *post-singularly finite (psf) exponential maps*: these are (up to affine conjugation) maps of the form $E_\lambda: \mathbb{C} \rightarrow \mathbb{C}$, $E_\lambda(z) = \lambda \exp(z)$, where the parameter $\lambda \in \mathbb{C} \setminus \{0\}$ is chosen such that the orbit of the only singular value 0 is finite

and hence strictly pre-periodic (the orbit of 0 cannot be periodic because 0 is an omitted value). For these functions, the set of escaping points has been described and classified in [SZ1], using the combinatorial concept of *external addresses* that distinguish dynamic rays.

Definition 2.3. (External addresses and the shift map) An *external address* \underline{s} is a sequence $\underline{s} = s_1s_2s_3 \dots$ over the integers. We denote by $\mathcal{S} = \mathbb{Z}^{\mathbb{N}}$ the space of all external addresses and by $\sigma : \mathcal{S} \rightarrow \mathcal{S}$, $\sigma(s_1s_2s_3 \dots) = s_2s_3 \dots$ the *left shift map*.

The shift space \mathcal{S} can be totally ordered using lexicographic order (for $\underline{s}, \underline{t} \in \mathcal{S}$, we write $\underline{s} < \underline{t}$ if and only if $\underline{s} = s_1s_2s_3 \dots$ and $\underline{t} = t_1t_2t_3 \dots$ have a k so that $s_1 = t_1, \dots, s_k = t_k$ and $s_{k+1} < t_{k+1}$). The lexicographic order defines the order topology on \mathcal{S} . Any total order induces a *cyclic order* on the same set: for distinct $\underline{s}, \underline{t}, \underline{u} \in \mathcal{S}$, we write

$$\underline{s} < \underline{t} < \underline{u} : \Leftrightarrow (\underline{s} < \underline{t} < \underline{u}) \vee (\underline{t} < \underline{u} < \underline{s}) \vee (\underline{u} < \underline{s} < \underline{t}).$$

We write $(\underline{s}, \underline{t})$ for the open interval between \underline{s} and \underline{t} with respect to the cyclic order on \mathcal{S} , that is, we have $\underline{u} \in (\underline{s}, \underline{t})$ if and only if $\underline{s} < \underline{u} < \underline{t}$.

The following theorem is a weak version of the classification result proved in [SZ1], but it is all we need for this work.

PROPOSITION 2.4. (The escaping set of a post-singularly finite exponential map) *Let E_λ be a pcf exponential map. Then every escaping point either lies on a unique dynamic ray or is the landing point of a unique ray. In particular, distinct dynamic rays are disjoint.*

We can assign to each (pre-)periodic external address \underline{s} a dynamic ray $g_\underline{s} : (0, \infty) \rightarrow \mathbb{C}$ in such a way that

$$E_\lambda(g_\underline{s}(t)) = g_{\sigma(\underline{s})}(F(t)) \quad \text{for all } t > 0,$$

where $F(t) = e^t - 1$, and so that

$$\lim_{t \rightarrow \infty} \operatorname{Re} g_\underline{s}(t) = +\infty \quad \text{and} \quad \lim_{t \rightarrow \infty} \operatorname{Im} g_\underline{s}(t) = -\operatorname{Im} \log \lambda + 2\pi s_1. \tag{1}$$

Here, we chose the branch of the logarithm for which $\operatorname{Im} \log \lambda \in (-\pi, \pi]$.

Remark 2.5. Note that the landing point of a dynamic ray can be an escaping point; in this case, there cannot be another dynamic ray with the same landing point. In contrast, a non-escaping point can be the landing point of several dynamic rays. Indeed, (pre-)periodic rays landing together play an important role in this paper.

We partition the complex plane into horizontal strips of the form

$$S_k := \{z \in \mathbb{C} : -\operatorname{Im} \log \lambda - \pi + 2\pi k < \operatorname{Im} z < -\operatorname{Im} \log \lambda + \pi + 2\pi k\}$$

with $k \in \mathbb{Z}$. This partition is called the *static partition* for E_λ .

Equation (1) implies that every dynamic ray is, for all sufficiently large potentials t , contained in a single sector of the static partition, that is, there exists a $k \in \mathbb{Z}$ such that $g(t) \in S_k$ for t large enough.

Definition 2.6. (External address of a dynamic ray) Let g be a dynamic ray for a psf exponential map E_λ . The *external address* $\text{Ad}(g) = s_1s_2 \dots \in \mathcal{S}$ of the dynamic ray g is the unique external address such that for every $k \in \mathbb{N}$, we have

$$E_\lambda^{\circ(k-1)}(g(t)) \in S_{s_k} \quad \text{for } t \text{ large enough (depending on } k\text{)}.$$

Remark 2.7. We see that $\text{Ad}(g_t) = \underline{t}$ for the dynamic rays g_t defined in Proposition 2.4. Different dynamic rays have different external addresses (as follows from the full version of the classification result in [SZ1]), but not every external address is the address of a dynamic ray. A sequence $\underline{s} \in \mathcal{S}$ occurs as the external address of a dynamic ray if and only if it is *exponentially bounded* (see [SZ1, Definition 4.1 and Theorem 4.2]).

Because dynamic rays are disjoint and converge to ∞ in a controlled way, as described in (1), they have a well-defined vertical order, defined as follows.

Definition 2.8. (Vertical order of dynamic rays) Let g and g' be two dynamic rays of E_λ . Then for sufficiently large $\xi > 0$, the ray g disconnects the right half-plane $\{z \in \mathbb{C} : \text{Re}(z) > \xi\}$ into exactly two unbounded parts (plus possibly some bounded ones), and the curve $g'(t)$ must be contained in a single one of them for all sufficiently large t . We say that g *lies above* g' if g' (for large potentials) is contained in the lower of these two unbounded components, and write $g > g'$; otherwise we say that g *lies below* g' and write $g < g'$.

It follows easily from the mapping properties of the exponential map that the vertical order of dynamic rays coincides with the lexicographical order of their external addresses (see [FS, Lemma 3.9]).

LEMMA 2.9. (Order of rays and external addresses) *For any two dynamic rays g_t and g'_t of E_λ , the ray g_t lies above g'_t if and only if $\underline{t} > \underline{t}'$ in lexicographic ordering.*

As described above for the space of external addresses, the linear order induces a cyclic order on the set of dynamic rays; we write $g < g' < g''$ if g' lies between g and g'' in this cyclic order. The map assigning to each ray its external address preserves the cyclic order (as it preserves the linear order inducing it).

For the purposes of this paper, we are most interested in (pre-)periodic dynamic rays and their landing behavior (compare Definition 2.2). It was shown in [SZ2, Theorem 3.2] that every (pre-)periodic dynamic ray of a psf exponential map that is not eventually mapped onto a ray landing at the singular value 0 lands at a (pre-)periodic point. Conversely, every (pre-)periodic point is the landing point of at least one (pre-)periodic dynamic ray by [SZ2, Theorem 5.3]. This result is in analogy to the Douady–Hubbard landing theorem for polynomials.

THEOREM 2.10. (A landing theorem [SZ2]) *For every post-singularly finite exponential map, every periodic or pre-periodic point is the landing point of at least one and at most finitely many periodic, respectively pre-periodic, dynamic rays. Rays landing at the same point have the same pre-period and period.*

We construct a partition of the plane that allows us to describe in combinatorial terms which dynamic rays are landing together: Choose a pre-periodic dynamic ray $g_{\underline{s}}$ landing at 0. The pre-image $E_{\lambda}^{-1}(\text{Tr}(g_{\underline{s}}))$ consists of countably many disjoint simple curves which are translations of each other by integer multiples of $2\pi i$. Furthermore, if g is any lift of $g_{\underline{s}}$ by E_{λ} , then we have

$$\lim_{t \rightarrow 0} \text{Re}(g(t)) = -\infty \quad \text{and} \quad \lim_{t \rightarrow \infty} \text{Re}(g(t)) = +\infty$$

by the mapping properties of exponential maps.

We define a *sector* to be a connected component of $\mathbb{C} \setminus E_{\lambda}^{-1}(\text{Tr}(g_{\underline{s}}))$. These sectors partition \mathbb{C} so that the sector boundaries are exactly the pre-images of the ray $g_{\underline{s}}$.

Because distinct dynamic rays are disjoint, there is a unique sector D_0 containing the ray $g_{\underline{s}}$ and its landing point 0. For $k \in \mathbb{Z}$, the sector obtained by translating D_0 by $2k\pi i$ is called D_k . Observe that for all $k \in \mathbb{Z}$, the restriction $E_{\lambda} : D_k \rightarrow \mathbb{C} \setminus (\text{Tr}(g_{\underline{s}}) \cup \{0\})$ is biholomorphic. We call $\mathcal{D} := \bigcup_{k \in \mathbb{Z}} \{D_k\}$ the *dynamical partition* for E_{λ} with respect to $g_{\underline{s}}$ and we call $\partial\mathcal{D} := E_{\lambda}^{-1}(\text{Tr}(g_{\underline{s}}))$ the *boundary* of the partition \mathcal{D} . Note that \mathcal{D} is a collection of subsets of \mathbb{C} , and $\partial\mathcal{D} \subset \mathbb{C}$.

As dynamic rays are parameterized by external addresses, the dynamical partition can also be constructed on the level of external addresses. Consider again the shift space \mathcal{S} endowed with the lexicographic order. In the following, terms like $t_0\underline{t} \in \mathcal{S}$ with $t_0 \in \mathbb{Z}$ and $\underline{t} \in \mathcal{S}$ will denote concatenation of an integer and an external address.

We start with the external address $\underline{s} = s_1s_2 \dots$ of the dynamic ray $g_{\underline{s}}$. This address is not constant, so we either have $\underline{s} \in ((s_1 - 1)\underline{s}, s_1\underline{s})$ or $\underline{s} \in (s_1\underline{s}, (s_1 + 1)\underline{s})$. Denote the interval containing \underline{s} by I_0 . For all $k \in \mathbb{Z}$, we define the interval

$$I_k := \{t_1t_2t_3 \dots \in \mathcal{S} : (t_1 - k)t_2t_3 \dots \in I_0\}.$$

Observe that $\mathcal{I} := \bigcup_{k \in \mathbb{Z}} \{I_k\}$ is a partition of the shift space \mathcal{S} , which we call the *dynamical partition* of \mathcal{S} with respect to \underline{s} . We denote by $\partial\mathcal{I} := \sigma^{-1}(\underline{s}) \subset \mathcal{S}$ the *boundary* of the partition \mathcal{I} .

It follows from Lemma 2.9 that a dynamic ray g is contained in the sector D_k of the dynamical partition \mathcal{D} of the complex plane if and only if its address $\text{Ad}(g) \in \mathcal{S}$ is contained in the sector I_k of the dynamical partition of the shift space. Recording the sectors in \mathcal{D} to which a point in the plane is mapped under iteration of E_{λ} yields so-called itineraries. The same applies to the partition \mathcal{I} of the shift space.

Definition 2.11. (Itineraries of external addresses and the kneading sequence) Let \star be a formal symbol not contained in \mathbb{Z} and let $\underline{t} \in \mathcal{S}$ be an external address. Then the *itinerary* $\text{It}(\underline{t} \mid \underline{s})$ of \underline{t} with respect to \underline{s} is the unique sequence $\text{It}(\underline{t} \mid \underline{s}) = \underline{\tau} = \tau_1\tau_2 \dots \in (\mathbb{Z} \cup \{\star\})^{\mathbb{N}}$ such that

$$\tau_k := \begin{cases} i & \text{if } \sigma^{\circ(k-1)}(\underline{t}) \in I_i, \\ \star & \text{if } \sigma^{\circ(k-1)}(\underline{t}) \in \partial\mathcal{I}. \end{cases}$$

We call the itinerary $\nu := \text{It}(\underline{s} \mid \underline{s})$ the *kneading sequence* of the external address \underline{s} .

Definition 2.12. (Itineraries of points and dynamic rays) Let \star be a formal symbol not contained in \mathbb{Z} . For a point $z \in \mathbb{C}$, we define the *itinerary* of z with respect to \mathcal{D} to be the sequence $\text{It}(z \mid \mathcal{D}) = \underline{t} = t_1 t_2 \dots \in (\mathbb{Z} \cup \{\star\})^{\mathbb{N}}$ such that

$$t_k := \begin{cases} i & \text{if } E_\lambda^{\circ(k-1)}(z) \in D_i, \\ \star & \text{if } E_\lambda^{\circ(k-1)}(z) \in \partial\mathcal{D}. \end{cases}$$

Let g be a dynamic ray of E_λ . As distinct dynamic rays do not intersect and the partition boundary $\partial\mathcal{D}$ consists of rays, every iterate of g is either fully contained in a single sector of the dynamical partition or part of the partition boundary. Therefore, all points on g have equal itineraries and we set $\text{It}(g \mid \mathcal{D}) := \text{It}(w \mid \mathcal{D})$ for an arbitrary point $w \in \text{Tr}(g)$.

Itineraries are written in a different font to distinguish them from external addresses. Observe that the itinerary of a (pre-)periodic point is itself (pre-)periodic. Furthermore, it does not contain the symbol \star because $\partial\mathcal{D}$ consists of escaping points.

The following proposition describes the landing behavior of (pre-)periodic dynamic rays in terms of itineraries. It is crucial for the construction of homotopy Hubbard trees for exponential maps. A proof can be found in [SZ2, Lemma 5.2 and Theorem 5.3].

PROPOSITION 2.13. (Landing behavior of (pre-)periodic rays [SZ2]) *Let E_λ be a psf, let \underline{s} be the external address of a dynamic ray landing at the singular value, and let \mathcal{D} be the corresponding dynamical partition; denote the kneading sequence of \underline{s} by v .*

The map $z \mapsto \text{It}(z \mid \mathcal{D})$ is a bijection between (pre-)periodic points and (pre-)periodic itineraries $\underline{t} \in \mathbb{Z}^{\mathbb{N}}$ satisfying $\sigma^{\text{on}}(\underline{t}) \neq v$ for all $n \in \mathbb{N}$. The (pre-)periodic ray $g_{\underline{t}}$ lands at the (pre-)periodic point z if and only if $\text{It}(z \mid \mathcal{D}) = \text{It}(\underline{t} \mid \underline{s})$.

One of the main goals of this paper is to classify post-singularly finite exponential maps in terms of their abstract Hubbard trees. A different classification of psf exponential maps has already been obtained in [LSV, Theorems 2.6 and 2.7].

THEOREM 2.14. (Classification of post-singularly finite exponential maps [LSV]) *For every pre-periodic external address \underline{s} starting with the entry 0, there is a unique post-singularly finite exponential map such that the dynamic ray at external address \underline{s} lands at the singular value.*

Every post-singularly finite exponential map is associated in this way to a positive finite number of pre-periodic external addresses starting with 0. Two such external addresses \underline{s} and \underline{s}' are associated to the same exponential map if and only if $\text{It}(\underline{s}' \mid \underline{s}) = \text{It}(\underline{s} \mid \underline{s})$.

We do not use this result, neither for the construction of homotopy Hubbard trees for exponential maps nor in the proof of our own classification result. Still, it is a nice insight to have in mind during the upcoming constructions.

Convention 2.15. Let E_λ be a post-singularly finite exponential map. For the rest of the paper, \underline{s} will always denote the external address of a dynamic ray landing at the singular value and v will denote its kneading sequence.

3. Homotopy Hubbard trees

Hubbard trees were defined for polynomials over thirty years ago in [DH1]. We give here a more conceptual but equivalent definition of Hubbard trees, which has a better chance to be generalized to the case of exponential maps than the original one given in [DH1] (which uses the concept of filled-in Julia sets). To avoid difficulties arising from unnecessary generality, we do not give the most general definition, but a definition valid for *unicritical polynomials* (polynomials with only one critical point in the complex plane) with pre-periodic critical value. This makes sense from a conceptual viewpoint as exponential maps are a dynamical limit of unicritical polynomials [DGH] and for post-singularly finite exponential maps, the unique singular value is pre-periodic.

Let us make the concept of an embedded tree precise. Our definition differs from the standard definition (based on a topological quotient) in the case of infinite trees: infinite trees occur naturally as part of our construction and the topology of these trees at vertices of infinite degree deviates from the usual quotient topology.

Definition 3.1. (Embedded graphs and trees) A topological space G is called a *topological graph* if it is homeomorphic to a space

$$X = \left(\bigsqcup_{i \in I} [0^{(i)}, 1^{(i)}] \right) / \sim$$

of disjoint copies of the unit interval $[0, 1]$, where I is some index set (for our purposes, we may restrict to countable index sets) and \sim is an equivalence relation identifying some of the endpoints of the intervals $[0^{(i)}, 1^{(i)}]$.

Away from equivalence classes of infinite cardinality, we define the topology on X to be the usual quotient topology. In particular, if I is finite, we equip X with the usual quotient topology.

If $x \in X$ is a vertex of infinite degree, that is, an equivalence class of infinite cardinality, there are index sets $I_0, I_1 \subset I$, at least one of which is infinite, such that $0^{(i)} \in x$ if and only if $i \in I_0$ and $1^{(i)} \in x$ if and only if $i \in I_1$. For all $0 < \epsilon < 1$, we set

$$U_\epsilon := \left(\left(\bigcup_{i \in I} [0^{(i)}, \epsilon^{(i)}] \right) \cup \left(\bigcup_{i \in I} [1 - \epsilon^{(i)}, 1^{(i)}] \right) \right) / \sim,$$

and define $\{U_\epsilon\}_{0 < \epsilon < 1}$ to be a neighborhood basis of x in X .

We call G *finite* if the index set I may be chosen to be finite. The topological graph G is called a *topological tree* if it is connected and has trivial fundamental group.

Definition 3.2. (Branch and endpoints in embedded trees) A *branch* of the topological tree H at the point $p \in H$ is the closure of a connected component of $H \setminus \{p\}$. We denote the number of branches of H at p by $\text{deg}_H(p)$ and call p a *branch point* of H if $\text{deg}_H(p) \geq 3$. If $\text{deg}_H(p) = 1$, we call p an *endpoint* of H .

Definition 3.3. (Hubbard trees for pre-periodic unicritical polynomials) Let p be a post-critically finite unicritical polynomial with pre-periodic critical value. The *Hubbard tree* $H \subset \mathbb{C}$ of p is the unique finite embedded tree such that:

- $P(p) \subset H$, that is, the forward orbit of every critical point of p is contained in the tree;
- all endpoints of H are contained in $P(p)$; and
- $p(H) \subset H$, that is, H is forward invariant under the dynamics of p .

The naive way to generalize this definition to the case of exponential maps fails because of the existence of an asymptotic value.

THEOREM 3.4. (Hubbard trees of exponential maps must contain escaping points) *Let E_λ be a post-singularly finite exponential map. There does not exist a finite embedded tree $H \subset \mathbb{C}$ that is forward invariant under the dynamics of E_λ and contains $P(E_\lambda)$.*

Proof. Let \underline{s} be the external address of a dynamic ray landing at the singular value 0 and let \mathcal{D} be the dynamical partition for E_λ with respect to $g_{\underline{s}}$. Furthermore, let $p, q \in H$, $p \neq q$, be arbitrary points of well-defined and different itineraries with respect to \mathcal{D} . For example, one can easily verify that there are two points on the forward orbit of 0 which differ by a multiple of $2\pi i$ and hence lie in different sectors of the dynamical partition \mathcal{D} , so one could take p and q to be those two points. Because H is a tree, there is a unique arc $\gamma : [0, 1] \rightarrow H$ connecting p and q . Because H is forward-invariant, for every $n \geq 0$, the points $E_\lambda^{on}(p)$ and $E_\lambda^{on}(q)$ are connected within H by $E_\lambda^{on}(\gamma)$. By hypothesis, for some $k \in \mathbb{N}$, the points $E_\lambda^{ok}(p)$ and $E_\lambda^{ok}(q)$ lie in different sectors of the dynamical partition, thus, by continuity, $E_\lambda^{ok}(\gamma)$ crosses the partition boundary and contains an escaping point, which contradicts the forward invariance of H . □

The idea that leads us to a meaningful definition of Hubbard trees for exponential maps is to relax the invariance condition: we do not require a Hubbard tree to be exactly forward invariant, but only invariant up to homotopy relative to the post-singular set. Note that this relaxation is not only necessary (because there is no exactly invariant tree), but also natural from the point of view of Thurston’s theory: an exponential Hubbard tree should determine a psf exponential map up to Thurston equivalence, and the Thurston equivalence class of a map is invariant under homotopies relative to the post-singular set in the domain and co-domain of the map (see Definition 6.4).

Because homotopies rel $P(E_\lambda)$ cannot be pushed forward by E_λ (because of the existence of non-post-singular pre-images of post-singular points), the right way to formulate the forward invariance condition is through backwards iteration. We want to say that the pre-image $E_\lambda^{-1}(H)$ of a Hubbard tree H contains H up to homotopy rel $P(E_\lambda)$. This statement bears a problem. The pre-image $E_\lambda^{-1}(H)$ is disconnected because H contains the singular value 0. Different connected components of $E_\lambda^{-1}(H)$ contain post-singular points which are, by definition, contained in H and are not allowed to move during the homotopy. Hence, by connectedness, H cannot be homotoped into its pre-image $E_\lambda^{-1}(H)$. By adding a pre-image $-\infty$ of 0 to the complex plane, the pre- image of H becomes connected and in fact becomes an (infinite) embedded tree, so it makes sense to require that H can be homotoped into $E_\lambda^{-1}(H) \cup \{-\infty\}$ rel $P(E_\lambda)$ in the extended plane.

Let us make this idea precise. Let \mathbb{C}_T (as a set) be defined as the disjoint union $\mathbb{C}_T := \mathbb{C} \dot{\cup} \{-\infty\}$, where, for now, the point $-\infty$ is just an abstract point not contained in the

complex plane. We turn \mathbb{C}_T into a topological space by choosing a neighborhood basis (V_n) of 0 and declaring the sets $U_n := E_\lambda^{-1}(V_n) \cup \{-\infty\}$ to be a neighborhood basis of $\{-\infty\}$. The extension $\widehat{E}_\lambda: \mathbb{C}_T \rightarrow \mathbb{C}$ of E_λ defined by $\widehat{E}_\lambda(-\infty) := 0$ is continuous by definition. The completion \mathbb{C}_T of \mathbb{C} is a special case of a far more general construction of defining a completion of the domain of a holomorphic function by adding all transcendental singularities of its inverse function. See [E2] for further information. The extended map \widehat{E}_λ is not a covering map any more, but we can still lift homeomorphisms and homotopies of the complex plane that fix 0 under \widehat{E}_λ .

LEMMA 3.5. (Lifting homeomorphisms) *Let $A \subset \mathbb{C}$ be a set containing 0 and $\varphi: \mathbb{C} \rightarrow \mathbb{C}$ be a homeomorphism which is homotopic to the identity relative to A . There exists a unique homeomorphism $\Phi: \mathbb{C}_T \rightarrow \mathbb{C}_T$ which is homotopic to the identity relative to $\widehat{E}_\lambda^{-1}(A)$ such that the diagram*

$$\begin{array}{ccc} \mathbb{C}_T & \xrightarrow{\Phi} & \mathbb{C}_T \\ \widehat{E}_\lambda \downarrow & & \downarrow \widehat{E}_\lambda \\ \mathbb{C} & \xrightarrow{\varphi} & \mathbb{C} \end{array}$$

commutes. We call Φ the preferred lift of φ . Every homotopy between φ and id rel A lifts to a homotopy between Φ and id rel $\widehat{E}_\lambda^{-1}(A)$.

Proof. It follows from the homotopy lifting principle that every homotopy between id and φ rel A lifts to a homotopy between id and a homeomorphism $\Phi: \mathbb{C} \rightarrow \mathbb{C}$ relative to $E_\lambda^{-1}(A)$. Both the map Φ and the homotopy between Φ and the identity extend continuously to \mathbb{C}_T and fix $-\infty$ because φ as well as the homotopy between φ and the identity fix 0. □

To state our definition of homotopy Hubbard trees, we need a bit more vocabulary.

Definition 3.6. (Spanned subtrees) For an embedded tree $H \subset X$ and a finite subset $W \subset H$, we write $[W]_H$ (or just $[W]$) for the smallest subtree of H containing W . We say that H is spanned by W if $H = [W]$. Usually, we write $[p_1, p_2, \dots, p_n]_H$ for the smallest subtree containing the points $p_i \in H$ (omitting the curly brackets).

Assume that $H \subset \mathbb{C}$ is a finite embedded tree spanned by $P(E_\lambda)$. By the mapping properties of exponential maps, $\widehat{E}_\lambda^{-1}(H) \subset \mathbb{C}_T$ is an infinite embedded tree where the only point of infinite degree is $-\infty$ and every branch of $\widehat{E}_\lambda^{-1}(H)$ at $-\infty$ is homeomorphic to H . (This step uses our tailored definition of an infinite tree: the topology near the vertex $-\infty$ is compatible with the extension \mathbb{C}_T .)

As $P(E_\lambda)$ is forward invariant, we have $P(E_\lambda) \subset \widehat{E}_\lambda^{-1}(H)$, and therefore it makes sense to talk about the subtree $H' := [P(E_\lambda)]_{\widehat{E}_\lambda^{-1}(H)}$ of the pre-image tree spanned by $P(E_\lambda)$. We say that H is invariant up to homotopy rel $P(E_\lambda)$ if H' is homotopic to H in the extended plane \mathbb{C}_T relative to $P(E_\lambda)$. Let us make the last statement precise.

Definition 3.7. (Relative homotopies of embedded trees) Given a subset $A \subset X$ of the space X and two embedded trees $H, H' \subset X$, we say that H is homotopic to H' relative

to (rel) A if there exists a continuous map $I: H \times [0, 1] \rightarrow X$ with the following properties.

- $I^0 := I(\cdot, 0) = \text{id}$ and $I^1 := I(\cdot, 1): H \rightarrow H'$ is a homeomorphism.
- For all $p \in H \setminus A$, we have $I(p, t) \in X \setminus A$ for all $t \in [0, 1]$.
- For all $p \in H \cap A$, the homotopy I is constant on $\{p\} \times [0, 1]$.

As one might expect, this defines an equivalence relation: if two embedded trees H and H' are homotopic rel A through $I: H \times [0, 1] \rightarrow X$, then also H' and H are homotopic rel A through the *reversed homotopy* $\bar{I}: H' \times [0, 1] \rightarrow X$, $\bar{I}(p, t) := I((I^1)^{-1}(p), 1 - t)$. Given a second homotopy $I': H' \times [0, 1] \rightarrow X$ between H' and an embedded tree $H'' \subset X$ rel A , we obtain a homotopy between H and H'' rel A by forming the *concatenation* $I \cdot I': H \times [0, 1] \rightarrow X$ via

$$I \cdot I'(p, t) := \begin{cases} I(p, 2t) & \text{for } t \leq \frac{1}{2}, \\ I'(I^1(p), 2(t - \frac{1}{2})) & \text{for } t \geq \frac{1}{2}. \end{cases}$$

The concatenation satisfies $(I \cdot I')^1 = (I')^1 \circ I^1$.

Assume from now on that H is invariant up to homotopy rel $P(E_\lambda)$. While H is not forward invariant as a set, we still obtain a self-map of H which is well defined up to a certain equivalence relation: the homotopy $I: H \times [0, 1] \rightarrow \mathbb{C}_T$ from Definition 3.7 yields an identification of H and H' via the homeomorphism $\psi := I_1: H \rightarrow H'$ and the composition $f := \widehat{E}_\lambda \circ \psi: H \rightarrow H$ is a self-map of the tree H . We call f the *induced self-map* of H . There is a distinguished point $v_T := \psi^{-1}(-\infty) \in H$, the *singular point* of (H, f) , with the property that f is not injective at v_T while it is a local homeomorphism elsewhere and its image $f(v_T) = 0$ is the singular value. We call $V_f := \{f^{on}(v_T) \mid n \in \mathbb{N}_0\} \cup \{p \in H \mid \deg_H(p) \geq 3\}$ the set of *marked points* of (H, f) . It is the union of the forward orbit of v_T under f and the set of branch points of H , and, as such, contains $P(E_\lambda)$.

Different homotopies between H and H' yield different self-maps of H and we want to investigate this ambiguity. Let $\varphi: H \rightarrow H'$ be another identification of H with H' obtained by a homotopy between H and H' rel $P(E_\lambda)$, and let $f' := \widehat{E}_\lambda \circ \varphi$ be the corresponding self-map of H and set $v'_T := \varphi^{-1}(-\infty)$. As both maps ψ and φ are induced by homotopies relative to $P(E_\lambda)$, we have $\psi(p) = \varphi(p) = p$ for all $p \in P(E_\lambda)$. Therefore the ‘change of identification’ $\theta := \varphi^{-1} \circ \psi: H \rightarrow H$ restricts to the identity on $P(E_\lambda)$ and, in particular, on the set of endpoints of H . A (graph-theoretic) isomorphism between finite trees is uniquely determined by its values on endpoints; this implies $\theta(b) = b$ for all branch points $b \in H$. It might happen that $v_T \neq v'_T$, but we still have $\theta(v_T) = v'_T$ by definition of the singular point. We claim that $\theta(f(v)) = f'(v)$ for all $v \in V_f$. Indeed, if $v \neq v_T$ is a branch point of H , then f is a local homeomorphism at v , so $f(v)$ is also a branch point of H , and we have $\theta(f(v)) = f(v)$. If $v \in P(E_\lambda)$, then $f(v) \in P(E_\lambda)$ because $P(E_\lambda)$ is forward invariant under the dynamics of f , so again we have $\theta(f(v)) = f(v)$. Finally, if $v = v_T$, then $f(v_T) = 0 \in P(E_\lambda)$, so we have $\theta(f(v)) = f(v)$. It follows that θ restricts to a conjugation between f and f' on V_f because $\theta(f(v)) = f(v) = f'(\theta(v))$ for all $v \in V_f$. It turns out that if we just replace the forward invariance condition of Hubbard trees by the weaker condition of being invariant up to homotopy rel $P(E_\lambda)$, the graph structure of the resulting tree is not uniquely determined. For example, if there exists a periodic branch




FIGURE 1. Replacing one degree four branch point b by two degree three branch points b_1 and b_2 .

point $b \in H$ of degree four, one can split it into two degree three branch points b_1 and b_2 , changing the tree only in an arbitrary small neighborhood of b as indicated in Figure 1. If we perform this change consistently along the (forward and backward) orbit of b under the self-map f , we again obtain an invariant tree. Note, however, that in the new tree, we have $f^{on}([b_1, b_2]) = [b_1, b_2]$, where n is the period of b (and hence also of b_1 and b_2) under f , that is, the in-tree connection of b_1 and b_2 is forward invariant.

To obtain uniqueness of the graph structure, we want f to be expansive in the following sense.

Definition 3.8. (Expansive self-map) The self-map $f : H \rightarrow H$ is called *expansive* if, for every pair of marked points $p, q \in V_f$ with $p \neq q$, there exists an $n \geq 0$ such that $v_T \in f^{on}([p, q])$.

Because different self-maps of H are conjugate on their sets of marked points, the definition of expansivity does not depend on the choice of homotopy between H and H' . Again, the motivation for this definition comes from the polynomial case: the map obtained by restricting a unicritical polynomial with pre-periodic critical value to its Hubbard tree is always expansive in the sense that some iterated image of the in-tree connection of two marked points contains the unique critical point.

Summing up, we obtain the following definition of homotopy Hubbard trees for exponential maps.

Definition 3.9. (Homotopy Hubbard trees for exponential maps) Let E_λ be a post-singularly finite exponential map. A *homotopy Hubbard tree* $H \subset \mathbb{C}$ for E_λ is a finite embedded tree H such that the following conditions apply.

- H is spanned by $P(E_\lambda)$.
- H is invariant up to homotopy:

$$H' := [P(E_\lambda)]_{\widehat{E_\lambda}^{-1}(H)} \subset \mathbb{C}_T \text{ is homotopic to } H \text{ relative to } P(E_\lambda).$$

- The induced self-map of H is expansive.

Homotopy Hubbard trees for exponential maps are only required to be invariant up to homotopy. To prove a meaningful uniqueness statement, we have to define a suitable equivalence relation on homotopy Hubbard trees that deals with the increased flexibility compared to the polynomial case. Naturally, being a homotopy Hubbard tree is a property of homotopy classes of embedded trees relative to the post-singular set.

LEMMA 3.10. (Equivalent homotopy Hubbard trees) *Let $H, \tilde{H} \subset \mathbb{C}$ be finite embedded trees spanned by $P(E_\lambda)$ and assume that H and \tilde{H} are homotopic rel $P(E_\lambda)$. Then H is a homotopy Hubbard tree for E_λ if and only if \tilde{H} is a homotopy Hubbard tree for E_λ .*

Proof. By symmetry of Definition 3.7, both directions are equivalent and it suffices to show one of them. Assume that H is a homotopy Hubbard tree for E_λ , and let I_0 be a homotopy between H and \tilde{H} rel $P(E_\lambda)$. In analogy to Lemma 3.5, one shows that I_0 lifts to a homotopy $\hat{I}_1: \widehat{E}_\lambda^{-1}(H) \times [0, 1] \rightarrow \mathbb{C}_T$ between the pre-images of H and \tilde{H} rel $\widehat{E}_\lambda^{-1}(P(E_\lambda))$. As $P(E_\lambda) \subset \widehat{E}_\lambda^{-1}(P(E_\lambda))$, the homotopy \hat{I}_1 fixes $P(E_\lambda)$ pointwise, so the restriction $I_1 := \hat{I}_1|_{H' \times [0,1]}$ is a homotopy between H' and \tilde{H}' rel $P(E_\lambda)$. By the invariance of H , there exists a homotopy $I: H \times [0, 1] \rightarrow \mathbb{C}_T$ between H and H' . We see that \tilde{H} is homotopic to \tilde{H}' rel $P(E_\lambda)$ via the concatenation $\bar{I}_0 \cdot I \cdot I_1$. The self-map $\tilde{f} := \widehat{E}_\lambda \circ I_1^1 \circ I^1 \circ \bar{I}_0^{-1}$ of \tilde{H} obtained via this homotopy is conjugate to the self-map $f := \widehat{E}_\lambda \circ I^1$ of H . This can be seen using the fact that $\widehat{E}_\lambda \circ I_1^1 = I_0^1 \circ \widehat{E}_\lambda$ because I_1^1 is a lift of I_0^1 . Indeed,

$$\tilde{f} = \widehat{E}_\lambda \circ I_1^1 \circ I^1 \circ \bar{I}_0^{-1} = I_0^1 \circ \widehat{E}_\lambda \circ I^1 \circ \bar{I}_0^{-1} = I_0^1 \circ f \circ (I_0^1)^{-1}.$$

Therefore, the expansivity of \tilde{f} follows from the expansivity of f . □

Let us look more closely at the homotopy involved in the invariance condition of homotopy Hubbard trees. We want to see that after a small modification, this homotopy can be replaced by a stronger kind of homotopy called an ambient isotopy.

Definition 3.11. (Ambient isotopies) Let X be a topological space and $A \subset X$ a subspace. Two embedded trees $H, H' \subset X$ are called *ambient isotopic* relative to A if there exists a continuous map $I: X \times [0, 1] \rightarrow X$ such that the following conditions are satisfied.

- $I(\cdot, 0) = \text{id}$ and $I(H \times \{1\}) = H'$.
- For each $a \in A$, the homotopy I is constant on $\{a\} \times [0, 1]$.
- For each $t \in [0, 1]$, the time- t map $I(\cdot, t): X \rightarrow X$ is a homeomorphism.

Two homeomorphisms $\varphi_0, \varphi_1: X \rightarrow X$ are called *isotopic* relative to A if there exists a homotopy between φ_0 and φ_1 which is constant on A and restricts to a homeomorphism for each fixed $t \in [0, 1]$.

The distinction between relative homotopy and isotopy is known to be subtle; in general, ϕ being homotopic to the identity id rel A does not imply ϕ being isotopic to id rel A . In our setting, however, these notions are equivalent; we will use this in §6.

PROPOSITION 3.12. (E.g. [FM, Theorem 1.12]) *Let S be a closed oriented surface (that is, compact without boundary) and $A \subset S$ a finite set of marked points. An*

orientation-preserving homeomorphism $\phi: S \rightarrow S$ is homotopic to id rel A if and only if it is isotopic to id rel A .

For later use (in the proof of Proposition 4.9), we want to prove the following result regarding the homotopy involved in Definition 3.9.

LEMMA 3.13. (On the invariance condition) *Let H be a homotopy Hubbard Tree for E_λ . For every neighborhood $U \subset \mathbb{C}_T$ of $-\infty$, there exists an embedded tree $H'' \subset \mathbb{C}$ such that H'' is homotopic to H' rel $\mathbb{C}_T \setminus U$, and H'' is ambient isotopic to H in \mathbb{C} rel $P(E_\lambda)$.*

One might hope to obtain an ambient isotopy between H and H' rel $P(E_\lambda)$. Unfortunately, such an isotopy does not exist. The tree H' contains $-\infty$, while H does not, so the ambient isotopy would have to send $-\infty$ to some point in the complex plane. However, the space \mathbb{C}_T is not locally compact at $-\infty$, while it is locally compact at every point in the complex plane, so there does not exist a homeomorphism of \mathbb{C}_T sending $-\infty$ to a point in the complex plane.

Therefore, we have to take an intermediate step, homotoping *only the tree H'* to a tree $H'' \subset \mathbb{C}$. Then, we use a classical result on homotopies between graphs on surfaces to find an ambient isotopy between the modified tree H'' and H .

LEMMA 3.14. (Homotoping $-\infty$ into the plane) *Let $H \subset \mathbb{C}$ be a finite embedded tree spanned by $P(E_\lambda)$ and assume that 0 is an endpoint of H . Let $H' = [P(E_\lambda)]_{\widehat{E_\lambda}^{-1}(H)} \subset \mathbb{C}_T$ be the subset of the pre-image tree spanned by $P(E_\lambda)$. For every neighborhood U of $-\infty$ in \mathbb{C}_T , there is a homotopy between H' and an embedded tree $H'' \subset \mathbb{C}$ relative to $\mathbb{C}_T \setminus U$.*

Proof. The image $\tilde{V} := \widehat{E_\lambda}(U)$ is a neighborhood of 0 . As 0 is an endpoint of H , we can find a Jordan domain $V \subset \tilde{V}$ such that $H \cap V = \text{Tr}(\gamma)$ for some arc $\gamma: [0, 1] \rightarrow \overline{V}$ connecting 0 to some point $v \in \partial V$ and satisfying $\gamma([0, 1]) \subset V$. The pre-image $\tilde{U} := \widehat{E_\lambda}^{-1}(V) \subset U$ is a neighborhood of $-\infty$ and $\partial \tilde{U} \subset \mathbb{C}$ is a $2\pi i$ -periodic arc. The pre-image curves of γ have a natural vertical order and only finitely many of them are contained in H' . Let γ_u be the uppermost and γ_l be the lowermost pre-image curves of γ contained in H' . The arcs γ_u and γ_l , together with the part of $\partial \tilde{U}$ between their endpoints on $\partial \tilde{U}$, bound a simply connected domain $D \subset \tilde{U}$ which is contained in \mathbb{C} . The map $\Phi: \mathbb{C}_T \rightarrow \overline{\mathbb{C}}$, defined by $\Phi(z) = z$ for $z \in \mathbb{C}$ and $\Phi(-\infty) = \infty$, is continuous and injective, but its inverse is not continuous. The domain $\Phi(D)$ is a Jordan domain in $\overline{\mathbb{C}}$ and it is easy to see that $\tilde{H}' := \Phi(H')$ is homotopic to a tree $\tilde{H}'' \subset \overline{\mathbb{C}}$ rel $\overline{\mathbb{C}} \setminus \text{cl}_{\overline{\mathbb{C}}}(D)$. The closure $\text{cl}_{\mathbb{C}_T}(D)$ in the extended plane \mathbb{C}_T is homeomorphic to its closure $\text{cl}_{\overline{\mathbb{C}}}(D)$ on the Riemann sphere because every sequence $(z_n)_n \subset D$ with $\lim_{n \rightarrow \infty} |z_n| = \infty$ satisfies $\lim_{n \rightarrow \infty} \text{Re}(z_n) = -\infty$. (Note that this does not hold if we just require $(z_n)_n \in U$ because the sequence could escape in the vertical direction.) Therefore, the homotopy between \tilde{H}' and \tilde{H}'' in $\overline{\mathbb{C}}$ is pushed forward to a homotopy between H' and some $H'' \subset \mathbb{C}$ relative to $\mathbb{C}_T \setminus \text{cl}_{\mathbb{C}_T}(D)$. □

It remains to show that H'' is ambient isotopic to H rel $P(E_\lambda)$ in \mathbb{C} . The question under which conditions homotopic embedded graphs on a surface are ambient isotopic has

already been studied extensively. One rather general result can be found in [FM, Lemma 2.9]. We state a weaker version of this result that is sufficient for our purposes.

LEMMA 3.15. (Isotopies of curve systems on marked spheres) *Let $A \subset \overline{\mathbb{C}}$ be a finite set. Let $\gamma_1, \dots, \gamma_n$ be a collection of pairwise non-homotopic proper arcs in $\overline{\mathbb{C}}$ that do not intersect each other except possibly at their endpoints. If $\gamma'_1, \dots, \gamma'_n$ is another such collection so that γ'_i is homotopic to γ_i relative to A for each i , then there exists an ambient isotopy $I: \overline{\mathbb{C}} \times [0, 1] \rightarrow \overline{\mathbb{C}}$ rel A satisfying $I^1(\text{Tr}(\gamma_i)) = \text{Tr}(\gamma'_i)$ for all i simultaneously.*

Here, a *proper arc* is an arc $\gamma: [0, 1] \rightarrow \overline{\mathbb{C}}$ satisfying $\gamma^{-1}(A) = \{0, 1\}$, and two arcs are called *homotopic rel A* if they are homotopic rel A in the sense of Definition 3.7. We cannot apply Lemma 3.15 directly to our setting because no post-singular branch points of H are allowed to move during the homotopy, so there is no decomposition of H into proper arcs. The following lemma allows us to deal with this problem.

LEMMA 3.16. *Let $D \subset \overline{\mathbb{C}}$ be a Jordan domain and let $H, \tilde{H} \subset \overline{D}$ be finite embedded trees such that $H \cap \partial D = \tilde{H} \cap \partial D = \{p_1, \dots, p_n\}$, where the p_i are indexed according to their cyclic order on ∂D . Assume that there exists a homeomorphism $\varphi: H \rightarrow \tilde{H}$, such that $\varphi(p_i) = p_i$ for all $i \in \{1, \dots, n\}$, and that the set of endpoints of H (and therefore also of \tilde{H}) is a subset of $\{p_1, \dots, p_n\}$. Then H and \tilde{H} are ambient isotopic rel ∂D .*

Proof. We prove a more general statement with H and \tilde{H} replaced by finite unions of pairwise disjoint embedded trees. The points p_i and the branch points of H are the *vertices* of H , and the subarcs of H joining adjacent vertices are the *edges* of H . The proof works through induction on the number of edges.

There exists a component T of H and an index i such that $p_i, p_{i+1} \in T$. Let D' be the subdomain of D bounded by $[p_i, p_{i+1}]_H$ and $[p_{i+1}, p_i]_{\partial D}$. By hypothesis, we have $H \subset \overline{D'}$. Let \tilde{T} be the component of \tilde{H} with the same endpoints. There exists an isotopy of \overline{D} rel ∂D mapping $[p_i, p_{i+1}]_{\tilde{T}}$ to $[p_i, p_{i+1}]_T$. Both $H' = \overline{H \cap D'}$ and $\tilde{H}' = \overline{\tilde{H} \cap D'}$ consist of finitely many embedded trees with endpoints $p'_1, \dots, p'_m \in \partial D$, but the total number of edges of these trees has been reduced. □

LEMMA 3.17. *Let H be a homotopy Hubbard tree, and let $H'' \subset \mathbb{C}$ be homotopic to $H' = [P(E_\lambda)]_{\widehat{E_\lambda}^{-1}(H)}$ rel $P(E_\lambda)$ in \mathbb{C}_T . Then H'' is ambient isotopic to H rel $P(E_\lambda)$ in \mathbb{C} .*

Proof. It is easy to see that H'' is homotopic to H rel $P(E_\lambda)$ in \mathbb{C} . Hence, it remains to show that H'' is ambient isotopic to H relative to $P(E_\lambda)$ in \mathbb{C} . Choose a positively oriented simple closed curve $\Gamma: [0, 1] \rightarrow \mathbb{C}$ with $\Gamma(0) = \Gamma(1) = 0$ containing $P(E_\lambda)$ and satisfying the following properties.

- We can index the post-singular set as $P(E_\lambda) = \{0 = p_0, p_1, \dots, p_n = p_0\}$ and subdivide Γ into arcs $\gamma_i: [t_i, t_{i+1}] \rightarrow \mathbb{C}$ ($i \in \{0, \dots, n-1\}$) with $\gamma_i(t_i) = p_i, \gamma_i(t_{i+1}) = p_{i+1}$.
- We have $\gamma_i((t_i, t_{i+1})) \cap H = \emptyset$ for all $i \in \{0, \dots, n-1\}$.
- γ_i is homotopic to the arc $[p_i, p_{i+1}]_H$ relative to $P(E_\lambda)$.

Hence, Γ traverses the boundary of H touching it only once at every post-singular point. In the same way, we choose a simple closed curve Γ'' that traverses the boundary of H'' .

As H and H'' are homotopic relative to $P(E_\lambda)$, γ_i and γ_i'' are homotopic rel $P(E_\lambda)$ by the third property on the list. By Lemma 3.15, there is an ambient isotopy $I : \overline{\mathbb{C}} \times [0, 1] \rightarrow \overline{\mathbb{C}}$ between $\cup \gamma_i$ and $\cup \gamma_i''$ relative to $P(E_\lambda) \cup \{\infty\}$. The curve Γ'' traverses the boundaries of $H^1 := I^1(H)$ and H'' simultaneously. By Lemma 3.16, the trees H^1 and H'' are ambient isotopic rel $\text{Tr}(\Gamma'') \supset P(E_\lambda)$. □

Proof of Lemma 3.13. By Lemma 4.3, the singular value 0 is an endpoint of H . Hence, Lemma 3.13 is a direct consequence of Lemma 3.14 and Lemma 3.17. □

4. The triod algorithm: determining the graph structure

In this section, we show that for a fixed post-singularly finite exponential map, the structure of a homotopy Hubbard tree as a dynamical tree (see Definition 4.2) is uniquely determined by the kneading sequence (see Definition 2.11) of the external address of a ray landing at the singular value. Together with the fact that a homotopy Hubbard tree does not intersect dynamic rays landing at post-singular points (up to homotopy), this implies uniqueness of homotopy Hubbard trees.

If several dynamic rays land at the singular value, the kneading sequences of their external addresses agree by the following lemma, so we speak of the *kneading sequence* of a post-singularly finite exponential map. A proof of the result can be found in [LSV, Proof of Theorem 2.7, (1) \Rightarrow (5)].

LEMMA 4.1. (Different dynamical partitions) *Let E_λ be a psf exponential map for which the dynamic rays $g_{\underline{s}}$ and $g_{\underline{s}'}$ both land at 0. Then the kneading sequences $\text{It}(\underline{s}' \mid \underline{s}')$ and $\text{It}(\underline{s} \mid \underline{s})$ agree. Therefore, a post-singular point $p \in P(E_\lambda)$ is contained in the i th sector D_i of \mathcal{D} if and only if it is contained in the i th sector D'_i of \mathcal{D}' .*

Let H be a homotopy Hubbard tree and f be an induced self-map of H . We will see that two post-singular points are contained in the same branch of H at the singular point v_T if and only if their itineraries start with the same integer k_i , and we index the different branches B_{k_i} by the corresponding integers. If we accept this fact for the moment, the triple (H, f, B_{k_i}) satisfies the following definition.

Definition 4.2. (Exponential dynamical tree) An *exponential dynamical tree* is a triple (H, f, B_{k_i}) , where H is a finite topological tree and $f : H \rightarrow H$ is a self-map of H satisfying the following conditions.

- There exists a distinguished point $v_T \in H$, the *singular point*, such that f is not injective at v_T , whereas f is a local homeomorphism at each point $p \in H \setminus \{v_T\}$.
- All endpoints of H are on the singular orbit.
- The *singular value* $f(v_T)$ is strictly pre-periodic.
- If $p, q \in H$ with $p \neq q$ are branch points or points on the singular orbit, then there is an $n \geq 0$ such that $f^{on}([p, q])$ contains the singular point v_T (expansivity).

In addition, the connected components $B_{k_0}, B_{k_1}, \dots, B_{k_n}$ of $H \setminus \{v_T\}$ are indexed by distinct integers $k_i \in \mathbb{Z}$ such that $k_0 = 0$ and $f(v_T) \in B_0$.

The set V_f of *marked points* consists of the forward orbit of v_T under f (including v_T) and the set of branch points of H .

We call two exponential dynamical trees (H, f, B_{k_i}) and $(\tilde{H}, \tilde{f}, \tilde{B}_{k_i})$ *equivalent* if there exists a homeomorphism $\varphi: H \rightarrow \tilde{H}$ that restricts to a conjugation between f and \tilde{f} on the set of marked points V_f and maps each branch B_{k_i} to the branch \tilde{B}_{k_i} of the same index.

Let us state some simple properties of exponential dynamical trees first. Their proof is not very hard and works in complete analogy to the proof of [BKS, Lemma 2.3], so we omit it here.

LEMMA 4.3. (Basic properties of exponential dynamical trees) *Let (H, f, B_{k_i}) be an exponential dynamical tree. The singular value $f(v_T)$ is an endpoint of H . Each branch point is periodic or pre-periodic. The restriction of f to any branch of H at v_T is injective.*

We have seen in the previous section that different self-maps of a homotopy Hubbard tree H yield equivalent exponential dynamical trees (see the paragraphs after Definition 3.7). Furthermore, if \tilde{H} and H are equivalent homotopy Hubbard trees, we can choose their induced self-maps to be conjugate to each other (see the proof of Lemma 3.10). Therefore, we obtain a well-defined map from the set of equivalence classes of homotopy Hubbard trees to the set of equivalence classes of exponential dynamical trees. The main goal of this section is to give a constructive proof of the following statement.

THEOREM 4.4. *Let H and \tilde{H} be two homotopy Hubbard trees for a post-singularly finite exponential map E_λ . Then H and \tilde{H} yield equivalent exponential dynamical trees.*

Note that we do not require H and \tilde{H} to be equivalent. The proof of this theorem works in two steps. To explain the outline of the proof, we also have to define kneading sequences for exponential dynamical trees.

Definition 4.5. (Itineraries and kneading sequence) Let (H, f, B_{k_i}) be an exponential dynamical tree. The *itinerary* of a point $p \in H$ is the infinite sequence $\text{It}(p) = \underline{t} = t_1 t_2 \dots$ where

$$t_i := \begin{cases} i & \text{if } f^{\circ(k-1)}(p) \in B_{k_i}, \\ \star & \text{if } f^{\circ(k-1)}(p) = v_T. \end{cases}$$

The itinerary $\text{It}(f(v_T)) =: \nu = \nu_1 \nu_2 \dots$ of the singular value is called the *kneading sequence* of (H, f, B_{k_i}) .

The proof of Theorem 4.4 works in two steps. First we show that different homotopy Hubbard trees for the same psf exponential map yield exponential dynamical trees with equal kneading sequences. The second step is to show that an exponential dynamical tree is already determined (up to equivalence) by its kneading sequence.

To prove the first statement, we will show that for every homotopy Hubbard tree, there exists a dynamic ray landing at 0 such that, up to homotopy, the tree does not intersect the ray. Therefore, the pre-image of the tree does not intersect the boundary of the corresponding dynamical partition (up to homotopy). This implies the statement above because itineraries are independent of the choice of partition by Lemma 4.1 and every partition sector contains at most one branch of $H' = [P(E_\lambda)]_{\widehat{E}_\lambda^{-1}(H)}$ at $-\infty$.

For the proof, we need a topological characterization of dynamic rays landing at periodic post-singular points. Let $p \in P(E_\lambda)$ be periodic. An arc $\gamma : [0, 1] \rightarrow \overline{\mathbb{C}}$ with $\gamma(0) = p$, $\gamma(1) = \infty$, and $\text{Tr}(\gamma) \cap P(E_\lambda) = \{p\}$ is called a *leg at p*. We denote its homotopy class relative to $P(E_\lambda) \cup \{\infty\}$ by $[\gamma]$. By definition, a leg $\tilde{\gamma}$ at p is in the same homotopy class as γ if there is a homotopy between γ and $\tilde{\gamma}$ in $\overline{\mathbb{C}} \setminus (P(E_\lambda) \setminus \{p\})$ fixing the endpoints.

For every leg γ , we set $\tilde{\mathcal{L}}(\gamma) = \Gamma$, where $\Gamma : [0, 1] \rightarrow \overline{\mathbb{C}}$ with $\Gamma(0) = p'$ and $\Gamma(1) = \infty$ is the lift of γ starting at the unique periodic pre-image p' of p . Lifting is always possible as $0 \notin \text{Tr}(\gamma)$ and $E_\lambda : \mathbb{C} \rightarrow \mathbb{C} \setminus \{0\}$ is a covering map. As $P(E_\lambda)$ is forward invariant, the arc Γ satisfies $\text{Tr}(\Gamma) \cap P(E_\lambda) = \{p'\}$, so Γ is a leg at p' .

The map $\tilde{\mathcal{L}}$ defined in this way descends to a well-defined map $\mathcal{L}([\gamma]) = [\Gamma]$ because homotopies lift under covering maps. We call \mathcal{L} the *leg map*, using the terminology from [B, Definition 4.3].

PROPOSITION 4.6. (Topological characterization of dynamic rays) *Let $p \in P(E_\lambda)$ be periodic and let γ be a leg at p . There exists a dynamic ray $g \in [\gamma]$ if and only if $[\gamma]$ is periodic under the leg map.*

If $[\gamma]$ is periodic under the leg map, the ray $g \in [\gamma]$ is unique. Stated differently, distinct dynamic rays are not homotopic rel $P(E_\lambda) \cup \infty$.

Proof. It is clear that $[\gamma]$ is periodic under iteration of \mathcal{L} if there is a dynamic ray $g \in [\gamma]$: every dynamic ray landing at a periodic point is periodic as a set, so $[g] = [\gamma]$ is periodic under the leg map. The proof of the other direction is essentially the same as the proof of [B, Theorem 4.11]. Every iterate $f := E_\lambda^{\circ n}$ has the same dynamic rays as E_λ , so by passing to a suitable iterate f , we can assume that $[\gamma]$ is fixed under iteration of the leg map for f . Applying the leg map to $[\gamma]$ at least twice, we can further assume that γ is eventually contained in a single fundamental domain F_0 for f for some static partition of f . One can then prove just as in [B, Theorem 4.11] that γ is homotopic to the dynamic ray $g_{\underline{t}}$ of f with external address $\underline{t} = F_0 F_0 \dots$. If there was a second dynamic ray $g' \in [\gamma]$, it would also have external address \underline{t} , which would contradict the fact that every external address is the address of at most one ray. See [B] for the definition of the terminology and the details of the proof. □

Remark 4.7. We are slightly imprecise in the formulation of Proposition 4.6: a dynamic ray $g : (0, \infty) \rightarrow \overline{\mathbb{C}}$ landing at p neither contains p nor ∞ by Definition 2.2. When talking about the homotopy class of g , we treat it as an arc in $\overline{\mathbb{C}}$ containing its endpoints p and ∞ , so that g becomes a leg at p .

We prove that for every access to a homotopy Hubbard tree at a post-singular point, there exists a dynamic ray that approaches H through this access and does not intersect H up to homotopy. Before stating the theorem, let us formally define what we mean by an access.

Definition 4.8. (Accesses) Let H be a homotopy Hubbard tree for E_λ . An *access* to H at $p \in P(E_\lambda)$ is a prime end $A = \{[N_j]\}$ for the pair $(\overline{\mathbb{C}} \setminus H, H)$ with impression $\{p\}$ (see [M, §17]).

We say that a leg $\gamma : [0, 1] \rightarrow \overline{\mathbb{C}}$ approaches p through A if $\gamma(0) = p, \gamma((0, 1]) \cap H = \emptyset$, and, for every $j \in \mathbb{N}$, we have $\gamma(t) \in N_j$ for $t > 0$ small enough.

PROPOSITION 4.9. (Accesses contain dynamic rays) *Let H be a homotopy Hubbard tree for E_λ . There exists an equivalent homotopy Hubbard tree $\tilde{H} \cong H$ such that for each post-singular point $p \in P(E_\lambda)$ and each access A to p in \tilde{H} , there exists a dynamic ray g approaching p through A .*

Proof. Let $\{0 = p_0, p_1 = E_\lambda(p_0), \dots, p_n = E_\lambda^{on}(p_0), p_{n+1} = p_{l+1}\}$ be the post-singular orbit, where p_{l+1} is the first periodic point on the forward orbit of 0. Denote the accesses to H at p_m by $A_m^{(1)}, \dots, A_m^{(d_m)}$, where $d_m \in \mathbb{N}$ is the number of branches of H at p_m and the accesses $A_m^{(i)}$ are numbered according to their cyclic order at p_m . For every access $A_m^{(i)}$, choose a leg $\gamma_m^{(i)} : [0, 1] \rightarrow \overline{\mathbb{C}}$ approaching p_m through $A_m^{(i)}$. Note that if $\Gamma_m^{(i)}$ is another leg approaching p_m through $A_m^{(i)}$, then $\gamma_m^{(i)}$ is homotopic to $\Gamma_m^{(i)}$ rel $P(E_\lambda)$ (which is easy to see using the theory of prime ends). We first consider the periodic part of the forward orbit of 0. For $m \in \{l + 1, \dots, n\}$, the number of branches of H at p_m is independent of m (say equal to $d \in \mathbb{N}$) because the induced self-map $f : H \rightarrow H$ is a local homeomorphism at p_m and p_m is periodic. Choose a neighborhood V of 0 such that none of the legs $\gamma_m^{(i)}$ for $m \in \{l + 1, \dots, n\}$ intersect V . We set $U := \widehat{E_\lambda}^{-1}(V)$ and apply Lemma 3.13 to obtain an embedded tree $H'' \subset \mathbb{C}$ homotopic to $H' \text{ rel } \mathbb{C}_T \setminus U$. The lifts $\tilde{\mathcal{L}}(\gamma_m^{(i)})$ of the $\gamma_m^{(i)}$ along the periodic post-singular orbit do not intersect U and therefore are disjoint from H'' (as H' and H'' are equal on $\mathbb{C} \setminus U$) except for their endpoints. Now we see why we need Lemma 3.13: there exists an ambient isotopy $I : \mathbb{C} \times [0, 1] \rightarrow \mathbb{C}$ between H'' and $H \text{ rel } P(E_\lambda)$, and it extends to an ambient isotopy $I : \overline{\mathbb{C}} \times [0, 1] \rightarrow \overline{\mathbb{C}}$ fixing ∞ . Hence, $\Gamma_m^{(i)} := I|_{\overline{\mathbb{C}} \times \{1\}}(\tilde{\mathcal{L}}(\gamma_m^{(i)}))$ approaches p_{m-1} (or p_n for $m = l + 1$) through an access $A_{m-1}^{(\sigma_m(i))}$. By the above, $\Gamma_m^{(i)}$ is homotopic to $\gamma_{m-1}^{(\sigma_m(i))}$, so we have

$$\mathcal{L}([\gamma_m^{(i)}]) = [\gamma_{m-1}^{(\sigma_m(i))}] \quad \text{for all } m \in \{l + 1, \dots, n\} \text{ and all } i \in \{1, \dots, d\}.$$

As E_λ is locally an orientation-preserving homeomorphism and as ambient isotopies preserve the cyclic order of curves landing at a common point, we see that each $\sigma_m : \{1, \dots, d\} \rightarrow \{1, \dots, d\}$ is a power of the d -cycle $(1 \ 2 \ \dots \ d)$. Therefore, there exists an $N \in \mathbb{N}$ such that $\mathcal{L}^{\circ N}([\gamma_m^{(i)}]) = [\gamma_m^{(i)}]$ for all m and i . By Proposition 4.6, each $\gamma_m^{(i)}$ is homotopic to a dynamic ray $g_m^{(i)}$ rel $P(E_\lambda) \cup \{\infty\}$.

Because the $\gamma_m^{(i)}$ as well as the $g_m^{(i)}$ are pairwise non-homotopic rel $P(E_\lambda) \cup \{\infty\}$ and intersect each other at most at their endpoints, the two curve systems fulfill the hypotheses of Lemma 3.15. Thus, there exists an ambient isotopy $I : \overline{\mathbb{C}} \times [0, 1] \rightarrow \overline{\mathbb{C}}$ rel $P(E_\lambda) \cup \{\infty\}$ satisfying $I^1(\gamma_m^{(i)}) = g_m^{(i)}$ for all $m \in \{l + 1, \dots, n\}$ and all $i \in \{1, \dots, d\}$. By Lemma 3.10, the tree $H_0 := I^1(H)$ is again a homotopy Hubbard tree and it is equivalent to H .

The last part of the proof works by induction. Assume that $H_j \cong H$ is equivalent to our initial homotopy Hubbard tree and for each post-singular point $p \in P(E_\lambda)$, for which $E_\lambda^{\circ j}(p)$ is periodic and every access A to p in H_j , there exists a dynamic ray g approaching p through A . Choose a small neighborhood V_j of 0 such that V_j does not intersect the dynamic rays landing at the points p_{l+1-j}, \dots, p_n . The pre-image $U_l := \widehat{E_\lambda}^{-1}(V_l)$ does

not intersect any of the dynamic rays landing at the points p_{l-j}, \dots, p_n . By Lemma 3.13, there exists a tree $H'_j \subset \mathbb{C}$ which is isotopic to $H'_j \text{ rel } \mathbb{C}_T \setminus U_j$ and isotopic to $H_j \text{ rel } P(E_\lambda)$ in \mathbb{C} . By Lemma 3.10, the tree $H_{j+1} := H'_j$ is itself a homotopy Hubbard tree and we have $H_{j+1} \cong H_j \cong H$. By construction, for every $p \in \{p_{l-j}, \dots, p_n\}$ and every access to H_{j+1} at p , there exists a dynamic ray approaching p through this access. Iterating this procedure l times yields a tree $\tilde{H} := H_l$ satisfying the conditions of the theorem. \square

COROLLARY 4.10. (Equal kneading sequences) *Let H be a homotopy Hubbard tree for the post-singularly finite exponential map E_λ , let $f : H \rightarrow H$ be an induced self-map of H , and let v_T be the singular point of f . Two post-singular points $p, q \in P(E_\lambda)$ are contained in the same branch of H at v_T if and only if the first entries of the itineraries of p and q with respect to some, and hence any, dynamical partition (compare Lemma 4.1) are equal.*

As this property is independent of the choice of H , different homotopy Hubbard trees yield exponential dynamical trees with equal kneading sequences.

Proof. If $\tilde{H} \cong H$ and \tilde{f} is an induced self-map of \tilde{H} , then there exists a homeomorphism $\theta : H \rightarrow \tilde{H}$ that restricts to a conjugation between f and \tilde{f} on the set V_f of marked points. This follows from the elaborations on different self-maps of the same tree after Definition 3.7 together with the proof of Lemma 3.10. As θ maps branches of H at v_T to branches of \tilde{H} at \tilde{v}_T , it is enough to prove Corollary 4.10 for an equivalent homotopy Hubbard tree \tilde{H} .

By Proposition 4.9, there exist a dynamic ray g landing at 0 and a homotopy Hubbard tree $\tilde{H} \cong H$ such that $\text{Tr}(g) \cap \tilde{H} = \emptyset$. Let \mathcal{D} be the dynamical partition with respect to g . Then we have $\tilde{H}' \cap \partial\mathcal{D} = \emptyset$, where \tilde{H}' is the subtree of the pre-image tree of \tilde{H} spanned by $P(E_\lambda)$. It follows that post-singular points from different sectors cannot be contained in the same branch of \tilde{H}' at $-\infty$. Furthermore, each partition sector contains at most one branch of \tilde{H}' at $-\infty$, as the singular value is an endpoint of \tilde{H} by Lemma 4.3. Hence, p and q are contained in the same branch of \tilde{H}' at $-\infty$ if and only if they are contained in the same sector of \mathcal{D} , that is, if their itineraries start with the same integer. Branches of \tilde{H}' at $-\infty$ get identified with branches of H at \tilde{v}_T , so the corollary follows. \square

We will now show that an exponential dynamical tree is already determined (up to equivalence) by its kneading sequence. The following ideas, leading to the proof of this fact, are inspired by and, in many parts, analogous to results of [BKS]. Let us start with a simple observation.

PROPOSITION 4.11. (Distinct itineraries) *Let (H, f, B_{k_i}) be an exponential dynamical tree and let $p, q \in V_f$ be distinct marked points. Then $\text{It}(p) \neq \text{It}(q)$.*

Proof. We write $\text{It}(p) = p_1 p_2 \dots$ and $\text{It}(q) = q_1 q_2 \dots$. If $p_1 \neq q_1$, we are done. Else, we have $p_1 = q_1 = k_1$, and hence $p, q \in B_{k_1}$ for the branch B_{k_1} of H at v_T . By Lemma 4.3, the restriction $f|_{B_{k_1}}$ is injective, so we have $f([p, q]) = [f(p), f(q)]$. By expansivity of (H, f) , there exists a smallest $n \geq 0$, such that $v_T \in f^{on}([p, q]) = [f^{on}(p), f^{on}(q)]$. Therefore, the itineraries $\text{It}(f^{on}(p))$ and $\text{It}(f^{on}(q))$ have different initial entries. \square

Let (H, f, B_{k_i}) be an exponential dynamical tree and let $p, q, r \in V_f$ be distinct marked points. The subtree $[p, q, r]$ spanned by p, q , and r is called a *triad*. It must look in one of two ways: if $[p, q, r]$ is homeomorphic to the letter Y, we call it *branched*. Else, it is homeomorphic to the letter I and we call it *linear*. It should now be clear how to define the *middle point* $b[p, q, r] \in H$ of the triad $[p, q, r]$. Note that the middle point $b[p, q, r] \in V_f$ is also a marked point.

The shape of the triad $[p, q, r]$ as well as the itinerary of the middle point $b[p, q, r]$ can be determined algorithmically from the itineraries of the marked points p, q , and r . This algorithm is known as the *triad algorithm*. The triad algorithm is purely combinatorial and could be applied to any triple of pre-periodic sequences (not just the itineraries of p, q , and r). We apply it only to itineraries of vertices of an abstract exponential tree (H, f, B_{k_i}) ; let ν be its kneading sequence.

Definition 4.12. (Formal triads and the formal triad map) Let \star be a formal symbol not contained in \mathbb{Z} . The *space of formal (pre-)periodic points* $\mathcal{S}_\nu \subset (\mathbb{Z} \cup \{\star\})^\mathbb{N}$ consists of all (pre-)periodic sequences $\underline{t} \in (\mathbb{Z} \cup \{\star\})^\mathbb{N}$ such that either $\underline{t} \in \mathbb{Z}^\mathbb{N}$ and $\sigma^{on}(\underline{t}) \neq \nu$ for all $n \in \mathbb{N}$ or \underline{t} is contained in the backwards orbit

$$\Omega^-(\star\nu) := \{k_1 \dots k_n \star \nu \mid n \in \mathbb{N}_0, k_i \in \mathbb{Z}\} \subset (\mathbb{Z} \cup \{\star\})^\mathbb{N}$$

of finite integer sequences followed by $\star\nu$. We call a point $\underline{t} \in \mathcal{S}_\nu$ a *pre-singular point* if $\underline{t} \in \Omega^-(\star\nu)$.

We have the *shift map* $\sigma: (\mathbb{Z} \cup \{\star\})^\mathbb{N} \rightarrow (\mathbb{Z} \cup \{\star\})^\mathbb{N}$ acting on the full space of sequences. Note that we are using the same symbol for the shift map on \mathcal{S} . It will always be clear from the context which map we are considering.

Any triple of distinct sequences $\underline{t}, \underline{u}, \underline{v} \in \mathcal{S}_\nu$ is called a *formal triad* $[\underline{t}, \underline{u}, \underline{v}]$. Given a formal triad, we define the *formal triad map* \mathcal{T} as follows:

$$\mathcal{T}[\underline{t}, \underline{u}, \underline{v}] := \begin{cases} [\sigma(\underline{t}), \sigma(\underline{u}), \sigma(\underline{v})] & \text{if } t_1 = u_1 = v_1, \\ [\sigma(\underline{t}), \sigma(\underline{u}), \nu] & \text{if } t_1 = u_1 \neq v_1, \\ [\sigma(\underline{t}), \nu, \sigma(\underline{v})] & \text{if } t_1 = v_1 \neq u_1, \\ [\nu, \sigma(\underline{u}), \sigma(\underline{v})] & \text{if } t_1 \neq u_1 = v_1, \\ \text{stop} & \text{if } t_1, u_1, v_1 \text{ distinct.} \end{cases}$$

In all cases other than the **stop** case, $\mathcal{T}[\underline{t}, \underline{u}, \underline{v}]$ is again a formal triad: because \mathcal{S}_ν is forward invariant under σ and $\nu \in \mathcal{S}_\nu$, the three image sequences are contained in \mathcal{S}_ν and they are distinct, as we have $\{\underline{t}, \underline{u}, \underline{v}\} \cap \{k\nu \mid k \in \mathbb{Z}\} = \emptyset$. By construction, the only sequence that starts with \star is $\star\nu$, so in all cases other than the **stop** case, at least two of the first entries of the involved sequences are equal integers. If exactly two of the three first entries are equal, we say that the sequence whose first entry differs from the other two gets *chopped off* under iteration of \mathcal{T} .

The formal triad map can be iterated as long as the **stop** case is not reached. We write $\mathcal{T}^{on}[\underline{t}, \underline{u}, \underline{v}]$ for the resulting formal triad after n iterations of \mathcal{T} (if iteration is possible). Note that if the triad can be iterated indefinitely, at least two sequences each get chopped off under iteration of \mathcal{T} infinitely many times: otherwise, there would exist an $n \in \mathbb{N}$ such that

two of the three sequences of the triod $\mathcal{T}^{on}[\underline{t}, \underline{u}, \underline{v}]$ never get chopped off. However, this implies that these sequences are equal, which contradicts the fact that all three sequences stay distinct under iteration of \mathcal{T} .

Definition 4.13. (Majority vote and middle point of a formal triod) Let $[\underline{t}, \underline{u}, \underline{v}]$ be a formal triod. If $\mathcal{T}[\underline{t}, \underline{u}, \underline{v}] \neq \mathbf{stop}$, then, as noted above, at least two of the three sequences start with the same integer. We denote this integer by $m(\underline{t}, \underline{u}, \underline{v})$ and call it the *majority vote* of the triod $[\underline{t}, \underline{u}, \underline{v}]$. Let $i_0 \in \mathbb{N}_0$ be chosen such that $\mathcal{T}^{o i_0}[\underline{t}, \underline{u}, \underline{v}] = \mathbf{stop}$, if the triod eventually reaches the stop case, and let $i_0 = \infty$ otherwise. We define a sequence $\underline{b}(\underline{t}, \underline{u}, \underline{v}) \in \mathbb{Z}^{\mathbb{N}} \cup \Omega^-(\star v)$ by setting

$$(\underline{b}(\underline{t}, \underline{u}, \underline{v}))_i := \begin{cases} m(\mathcal{T}^{o(i-1)}[\underline{t}, \underline{u}, \underline{v}]) & \text{if } i < i_0, \\ \star & \text{if } i = i_0, \\ v_{i-i_0} & \text{if } i > i_0, \end{cases}$$

and we call $\underline{b}(\underline{t}, \underline{u}, \underline{v})$ the *middle point* of the triod. If we have $\underline{b}(\underline{t}, \underline{u}, \underline{v}) \notin \{\underline{t}, \underline{u}, \underline{v}\}$, we call $[\underline{t}, \underline{u}, \underline{v}]$ *branched*, otherwise we call it *linear*. Sometimes, we want to be more precise and call a triod *pre-singularly branched* or *pre-singularly linear* if $i_0 < \infty$, that is, if it eventually reaches the **stop** case under iteration of the triod map.

We now prove that the triod algorithm really determines the itineraries of the branch points of an exponential dynamical tree. The proof is an adaption of [BKS, Proposition 3.5].

LEMMA 4.14. (Correctness of the triod algorithm) *Let (H, f, B_{k_i}) be an exponential dynamical tree and let $[p_1, p_2, p_3] \subset H$ be a triod. Then $[p_1, p_2, p_3]$ is branched if and only if the formal triod $[\text{It}(p_1), \text{It}(p_2), \text{It}(p_3)]$ is branched. Furthermore, we have*

$$\text{It}(b[p_1, p_2, p_3]) = \underline{b}(\text{It}(p_1), \text{It}(p_2), \text{It}(p_3)).$$

Proof. If p_1, p_2 , and p_3 are contained in distinct branches of H at v_T , then $b[p_1, p_2, p_3] = v_T$, and this is the branch point determined (on the level of itineraries) by the triod algorithm. If one of the p_i equals v_T and the other two points are contained in different branches of H , then $b[p_1, p_2, p_3] = p_i$, and again this is the output of the triod algorithm.

If both of these cases do not occur, there exist a branch B_{k_1} of H at v_T and distinct indices $i, j \in \{1, 2, 3\}$, such that $p_i, p_j \in B_{k_1} \setminus \{v_T\}$. Hence, we also have $b[p_1, p_2, p_3] \in B_{k_1} \setminus \{v_T\}$. We see, that the first entry of $\text{It}(b[p_1, p_2, p_3])$ is calculated correctly by the triod algorithm.

If all of the p_l are contained in B_{k_1} , the spanned subtree $[p_1, p_2, p_3]$ is also entirely contained in B_{k_1} . By Lemma 4.3, the restriction $f|_{[p_1, p_2, p_3]}$ is injective, so $[f(p_1), f(p_2), f(p_3)]$ is also a triod and we have $b[f(p_1), f(p_2), f(p_3)] = f(b[p_1, p_2, p_3])$. On the level of formal triods, we have $\mathcal{T}[\text{It}(p_1), \text{It}(p_2), \text{It}(p_3)] = [\text{It}(f(p_1)), \text{It}(f(p_2)), \text{It}(f(p_3))]$. If instead $p_l \notin B_{k_1}$, while $p_i, p_j \in B_{k_1}$, then the chopped-off triod $[v_T, p_i, p_j]$ still gets mapped forward injectively. Hence, we have $b[0, f(p_i), f(p_j)] = f(b[p_1, p_2, p_3])$. On the level of formal triods, we have $\mathcal{T}[\text{It}(p_1), \text{It}(p_2), \text{It}(p_3)] = [\text{It}(0), \text{It}(f(p_i)), \text{It}(f(p_j))]$.

If we are in one of the two preceding cases (we have not reached the **stop** case in the first iteration step), we apply the same reasoning as before to the image triod. Hence, the triod algorithm correctly calculates $\text{It}(b[p_1, p_2, p_3])$. It remains to show that $[p_1, p_2, p_3]$ is branched if and only if the formal triod $[\text{It}(p_1), \text{It}(p_2), \text{It}(p_3)]$ is branched. If $[p_1, p_2, p_3]$ is branched, then

$$\text{It}(b[p_1, p_2, p_3]) = \underline{b}(\text{It}(p_1), \text{It}(p_2), \text{It}(p_3)) \neq \text{It}(p_i) \quad \text{for all } i \in \{1, 2, 3\}$$

by Proposition 4.11, so the formal triod $[\text{It}(p_1), \text{It}(p_2), \text{It}(p_3)]$ is also branched. If $[p_1, p_2, p_3]$ is linear, then

$$\text{It}(b[p_1, p_2, p_3]) = \underline{b}(\text{It}(p_1), \text{It}(p_2), \text{It}(p_3)) = \text{It}(p_i) \quad \text{for some } i \in \{1, 2, 3\},$$

so the formal triod $[\text{It}(p_1), \text{It}(p_2), \text{It}(p_3)]$ is also linear. □

Proof of Theorem 4.4. It remains to show that two exponential dynamical trees (H, f, B_{k_i}) and $(\tilde{H}, \tilde{f}, \tilde{B}_{k_i})$ with the same kneading sequence are equivalent. Because every endpoint of H is a post-singular point, every branch point of H is also the branch point of a post-singular triod. Thus, the kneading sequence fully determines the itineraries of all marked points of H and for every triod $[p, q, r]$ of marked points, it determines their incidence relation by Lemma 4.14. As (H, f, B_{k_i}) and $(\tilde{H}, \tilde{f}, \tilde{B}_{k_i})$ have the same kneading sequence, their marked points have the same itineraries, and we can define a homeomorphism $\varphi: H \rightarrow \tilde{H}$ by sending the marked points of (H, f, B_{k_i}) to the marked points of $(\tilde{H}, \tilde{f}, \tilde{B}_{k_i})$ with the same itinerary and extending this map to the edges between the marked points. The exponential dynamical trees (H, f, B_{k_i}) and $(\tilde{H}, \tilde{f}, \tilde{B}_{k_i})$ are equivalent via φ . □

Theorem 4.4 is an important step in proving uniqueness of homotopy Hubbard trees, but we still have to see that the embedding of the tree into the plane is also unique. Let H be a homotopy Hubbard tree and let $[p, q, r]$ be a triod of post-singular points. If $[p, q, r]$ is linear with middle point p , Proposition 4.9 implies that there are two dynamic rays $g_p, g_{p'}$ landing at p and separating q from r in the sense that q and r are contained in different connected components of $\mathbb{C} \setminus (\text{Tr}(g_p) \cup \text{Tr}(g_{p'}) \cup \{p\})$. Conversely, if such separating rays exist, the triod $[p, q, r]$ is linear with middle point p .

LEMMA 4.15. (Separating rays determine triod type) *Let H be a homotopy Hubbard tree and let $[p, q, r]$ be a triod of post-singular points. If there are two dynamic rays $g_p, g_{p'}$ landing at p , such that q and r are contained in different connected components of $\mathbb{C} \setminus (\text{Tr}(g_p) \cup \text{Tr}(g_{p'}) \cup \{p\})$, the triod $[p, q, r]$ is linear with middle point p .*

Proof. Let $\text{It}(p) = \underline{p} = p_1 p_2 \dots$, $\text{It}(q) = \underline{q} = q_1 q_2 \dots$, and $\text{It}(r) = \underline{r} = r_1 r_2 \dots$ be the itineraries of the points p, q , and r . By correctness of the triod algorithm (see Lemma 4.14), it is enough to show that $\underline{b}(\underline{p}, \underline{q}, \underline{r}) = \underline{p}$.

By Proposition 2.13 and Theorem 2.10, there exist addresses \underline{q} and \underline{r} satisfying $\text{It}(\underline{q} \mid \underline{s}) = \underline{q}$ and $\text{It}(\underline{r} \mid \underline{s}) = \underline{r}$. The triod $[\underline{p}, \underline{q}, \underline{r}]$ is called a triod of external addresses associated to $[\underline{p}, \underline{q}, \underline{r}]$ (this terminology is introduced rigorously in §5). It will be shown in

§5 that $[\underline{p}, \underline{q}, \underline{r}]$ has the same shape as the triod $[\underline{p}, \underline{q}, \underline{r}]$ of associated external addresses. Finally, Lemma 5.8 implies that $\underline{b}(\underline{p}, \underline{q}, \underline{r}) = \underline{p}$. □

We are now in the position to prove uniqueness of homotopy Hubbard trees.

THEOREM 4.16. (Uniqueness of exponential Hubbard trees) *Let H and \tilde{H} be homotopy Hubbard trees for the post-singularly finite exponential map E_λ . Then H and \tilde{H} are homotopic relative to the post-singular set $P(E_\lambda)$.*

Proof. Let $P(E_\lambda) = \{p_0 = 0, p_1 = E_\lambda(p_0), \dots, p_n = E_\lambda^{on}(0), p_{n+1} = p_{l+1}\}$ denote the post-singular orbit of E_λ . At every post-singular point p_i , the trees H and \tilde{H} have the same number of branches by Theorem 4.4; denote this number by d_i . By Proposition 4.9, we can assume without loss of generality that there are dynamic rays $g_1^{(i)}, \dots, g_{d_i}^{(i)}$ landing at p_i indexed according to their cyclic order at p_i and satisfying $H \cap \text{Tr}(g_l^{(i)}) = \emptyset$ such that $[p_i, p_j, p_k]_H$ is linear with middle point p_i if and only if p_j and p_k are contained in different connected components of $\mathbb{C} \setminus ((\cup_l \text{Tr}(g_l^{(i)})) \cup \{p_i\})$. In the same way, we choose rays $\gamma_l^{(i)}$ for \tilde{H} , and we claim that (possibly after re-indexing) $g_l^{(i)}$ is homotopic to $\gamma_l^{(i)}$ rel $P(E_\lambda)$.

Otherwise, there are indices i_0 and l_0 such that $\gamma_{l_0}^{(i_0)}$ is not homotopic to any of the rays $g_l^{(i_0)}$. There is an index l' such that $g_{l'}^{(i_0)} < \gamma_{l_0}^{(i_0)} < g_{l'+1}^{(i_0)}$ holds. Furthermore, there are rays $g_{l_j}^{(j)}$ and $g_{l_k}^{(k)}$ landing at distinct post-singular points p_j and p_k different from p_{i_0} such that $g_{l'}^{(i_0)} < g_{l_j}^{(j)} < \gamma_{l_0}^{(i_0)}$ and $\gamma_{l_0}^{(i_0)} < g_{l_k}^{(k)} < g_{l'+1}^{(i_0)}$ hold; otherwise, $\gamma_{l_0}^{(i_0)}$ would either be homotopic to $g_{l'}^{(i_0)}$ or to $g_{l'+1}^{(i_0)}$. However, then p_j and p_k are separated by $\gamma_{l_0}^{(i_0)}$ and $g_{l'}^{(i_0)}$, so $[p_{i_0}, p_j, p_k]$ is linear with middle point p_{i_0} by Lemma 4.15, which yields a contradiction.

By Proposition 4.6, homotopic dynamic rays landing together at a *periodic* post-singular point are equal, so we actually have $\gamma_l^{(i)} = g_l^{(i)}$ for all $i \in \{l+1, \dots, n\}$. The argument in the preceding paragraph also shows that the number of rays landing at p_i equals d_i . Analyzing the proof of Proposition 4.9, we conclude that we can choose H and \tilde{H} as to not intersect any dynamic ray landing at any post-singular point (periodic or not). In particular, we can find a ray g landing at 0 such that $H \cap E_\lambda^{oi}(\text{Tr}(g)) = \tilde{H} \cap E_\lambda^{oi}(\text{Tr}(g)) = \emptyset$ for all $i \in \{0, \dots, n\}$. Let $\varphi: \mathbb{D} \rightarrow \mathbb{C} \setminus \cup_i (E_\lambda^{oi}(\text{Tr}(g)) \cup \{p_i\})$ be a conformal map. By Carathéodory's theorem, φ extends continuously to $\partial\mathbb{D}$ and by Lemma 3.16, the trees $\varphi^{-1}(H)$ and $\varphi^{-1}(\tilde{H})$ are homotopic rel $\partial\mathbb{D}$. Pushing this homotopy forward via φ , we see that H and \tilde{H} are homotopic rel $P(E_\lambda)$, and hence they are equivalent. □

5. *Separating dynamic rays: embedding the tree into the plane*

In this section, we show the existence of homotopy Hubbard trees through an explicit construction. We use the triod algorithm from the preceding section to determine the middle points of post-singular triods, and thereby the set of marked points of the homotopy Hubbard tree on the level of itineraries. Using Proposition 2.13, we pick (pre-)periodic points in the complex plane realizing the itineraries in the formal set of marked points (for pre-singular itineraries, the construction is a bit more involved because there is no actual (pre-)periodic point in the plane realizing this itinerary).

It remains to find the right way to embed the edges of the tree into the complex plane. Our main idea is to find an embedded tree which does not intersect any dynamic rays landing at a marked point. (This statement is only approximately true; some rays landing at pre-singular marked points need to be intersected, but we do so in a controlled way.)

This is partially motivated by Proposition 4.9: if there exists a homotopy Hubbard tree, then, up to homotopy rel $P(E_\lambda)$, every access to the tree at every post-singular point contains a dynamic ray. In particular, this ray does not intersect the homotopy Hubbard tree. Another reason is the analogy to polynomials. A polynomial Hubbard tree does not intersect any dynamic rays. If a polynomial does not have bounded Fatou components, every access to every point on its Hubbard tree contains a dynamic ray.

The key observation is that the union of the rays landing at marked points partitions the plane in a meaningful way: there exists an embedded tree spanned by $P(E_\lambda)$ that does not intersect these rays (except for some rays landing at pre-singular points, as noted above), and this tree is unique up to homotopy rel $P(E_\lambda)$. As the rays landing at marked points form a forward invariant set, the pre-image tree also does not intersect them, so the embedded tree is invariant up to homotopy. Expansivity of the induced self-map follows because different marked points have different itineraries, so the embedded tree is a homotopy Hubbard tree for E_λ .

Let us begin with the construction of homotopy Hubbard trees. At first, we determine the set of marked points on the level of itineraries by using the triod algorithm. Throughout this section, ν denotes the kneading sequence of the exponential map E_λ (recall Lemma 4.1). Let $\Omega^+(\star\nu) := \{\sigma^{ok}(\star\nu) : k \geq 0\}$ denote the forward orbit of $\star\nu$ under the shift map σ . This is a finite set because ν is pre-periodic. We set

$$V_\nu := \Omega^+(\star\nu) \cup \bigcup_{[\underline{t}, \underline{u}, \underline{v}]} \{\underline{b}(\underline{t}, \underline{u}, \underline{v})\},$$

where the union runs over all formal triods consisting of sequences in $\Omega^+(\star\nu)$: we are adding all branch points of triods formed by post-singular points and $\star\nu$ to the set of post-singular points. We call V_ν the *formal vertex set*. The following properties of V_ν will be proved later in this section.

LEMMA 5.1. (The formal vertex set) *The formal vertex set V_ν has the following properties.*

- (1) *It is forward invariant under the shift map, that is, $\sigma(V_\nu) \subset V_\nu$.*
- (2) *It consists entirely of formal (pre-)periodic points, that is, $V_\nu \subset \mathcal{S}_\nu$.*
- (3) *It is closed under taking triods: for every triod $[\underline{t}, \underline{u}, \underline{v}]$ of formal vertices $\underline{t}, \underline{u}, \underline{v} \in V_\nu$, we have $\underline{b}(\underline{t}, \underline{u}, \underline{v}) \in V_\nu$.*

Next, we embed the formal vertex set into the complex plane. The resulting set will become the set of marked points of the yet to be constructed homotopy Hubbard tree. First, assume that $\underline{t} \in V_\nu$ is not a pre-singular point, that is, the itinerary $\underline{t} \in \mathbb{Z}^\mathbb{N}$ is (pre-)periodic and $\sigma^{on}(\underline{t}) \neq \nu$ for all $n \in \mathbb{N}$. By Proposition 2.13, there is exactly one (pre-)periodic point $v_{\underline{t}} \in \mathbb{C}$ with $\text{It}(v_{\underline{t}} | s) = \underline{t}$. This is the point we assign to our formal vertex \underline{t} .

For a pre-singular itinerary $\underline{t} \in V_\nu$, there is no (pre-)periodic point in \mathbb{C} of itinerary \underline{t} . This issue could be addressed by adding further points at infinity corresponding to iterated

pre-images of $-\infty$ to the plane. For brevity, we take a more hands-on approach: we choose surrogate points in the plane which are sufficiently close to these iterated pre-images. In our terms, this means that the itinerary of the vertex $v_{\underline{t}} \in \mathbb{C}$ shares sufficiently many entries with \underline{t} . To make this precise, let

$$N := \max\{n \in \mathbb{N}_0 \mid \text{there exists } \underline{t} \in V_\nu : \underline{t} = k_1 \dots k_n \star \nu\}$$

and pick a closed neighborhood U of the singular value 0 with the following properties.

- (1) $E_\lambda^{-n}(U) \cap E_\lambda^{-m}(U) = \emptyset$ for all distinct $n, m \in \{0, \dots, N + 1\}$.
- (2) For every $n \in \{1, \dots, N + 1\}$, the pre-image $E_\lambda^{-n}(U)$ does not intersect the union of the set of non-pre-singular vertices and the dynamic rays landing at these points.
- (3) U is bounded by a Jordan curve such that the dynamic ray $g_{\underline{s}}$ landing at the singular value intersects the boundary of U exactly once, that is, there is a unique potential t_0 such that $g_{\underline{s}}(t_0) \in \partial U$, whereas $g_{\underline{s}}(t) \in U$ for $t < t_0$ and $g_{\underline{s}}(t) \in \mathbb{C} \setminus \overline{U}$ for $t > t_0$. This is only to simplify topological considerations involving U .

The first condition is equivalent to $U \cap E_\lambda^{\circ n}(U) = \emptyset$ for all $n \in \{1, \dots, N + 1\}$ and this is obviously fulfilled for U small enough because 0 is not periodic. The second condition is equivalent to U not intersecting the union of the set of non-pre-singular vertices and the dynamic rays landing at those except for the singular value 0. Again, this is true for U small enough. Finally, condition (3) can always be ensured by shrinking a neighborhood that fulfills conditions (1) and (2).

Let $\log_{\underline{s},k} : \mathbb{C} \setminus (\text{Tr}(g_{\underline{s}}) \cup \{0\}) \rightarrow D_k$ be the branch of the inverse of E_λ with the stated domain and co-domain. For every itinerary $\underline{t} \in S_\nu$ of the form $\underline{t} = k_1 \dots k_n \star \nu$ with $n \in \{0, \dots, N\}$, we define

$$U_{\underline{t}} := \log_{\underline{s},k_1} \dots \log_{\underline{s},k_n}(E_\lambda^{-1}(U)),$$

that is, $U_{\underline{t}}$ is the iterated pre-image under E_λ of the domain U constructed above by the branches of the logarithm prescribed by the entries of \underline{t} . Note that by property (1) of U , distinct $U_{\underline{t}}$ and $U_{\underline{u}}$ are disjoint. We assign to $\star \nu \in V_\nu$ an arbitrary point $v_{\star \nu} \in U_{\star \nu}$ which will later become the singular point (see the paragraph after Definition 3.7) of the homotopy Hubbard tree. If $\underline{t} = k_1 \dots k_n \star \nu$ is pre-singular, we associate to \underline{t} the point

$$v_{\underline{t}} := \log_{\underline{s},k_1} \dots \log_{\underline{s},k_n}(v_{\star \nu}) \in U_{\underline{t}}.$$

To every formal vertex $\underline{t} \in V_\nu$, we have thus associated a point $v_{\underline{t}} \in \mathbb{C}$.

Definition 5.2. (Vertex set and triods of vertices) We define the *vertex set* $V \subset \mathbb{C}$ to be the set of all $v_{\underline{t}}$ for $\underline{t} \in V_\nu$. A *triod of vertices* $[v_{\underline{t}}, v_{\underline{u}}, v_{\underline{v}}]$ is a triple of distinct vertices $v_{\underline{t}}, v_{\underline{u}}, v_{\underline{v}} \in V$.

As triods of vertices are in natural bijection to triods of formal vertices, we use the terminology introduced for formal triods for vertex triods too. In particular, we call a triod of vertices (pre-singularly) branched or (pre-singularly) linear if the corresponding formal triod is of this type.

We now turn our attention to constructing the edges of a homotopy Hubbard tree. The partitioning property mentioned at the beginning of this section rests on a result about

dynamic rays separating the vertex set V : if $p, q, r \in V$ are distinct vertices and the corresponding formal triod $[p, q, r]$ is branched, then there exist (pre-)periodic dynamic rays g_1, g_2, g_3 of itinerary $\underline{b}[p, q, r]$, such that p, q , and r are separated by the rays g_i , that is, such that these three points are contained in different connected components of $\mathbb{C} \setminus \bigcup \text{Tr}(g_i)$. If $[p, q, r]$ is linear, there are two (pre-)periodic rays landing at the middle point and separating the other two points from each other. To prove the existence of such separating dynamic rays, we use the language of external addresses. This requires a variant of the triod algorithm operating on the level of external addresses.

Definition 5.3. (Formal triods and the formal triod map) Recall our notation $\mathcal{I} = \{I_k\}_{k \in \mathbb{Z}}$ for the dynamical partition of the shift space \mathcal{S} with respect to \underline{s} (see §2). In particular, \underline{s} is the external address of a dynamic ray $g_{\underline{s}}$ landing at the singular value 0.

A formal triod of external addresses is a triple $[\underline{t}, \underline{u}, \underline{v}]$ of external addresses $\underline{t}, \underline{u}, \underline{v} \in \mathcal{S}$ such that $\underline{t} < \underline{u} < \underline{v}$, and $\text{It}(\underline{t} | \underline{s}), \text{It}(\underline{u} | \underline{s}), \text{It}(\underline{v} | \underline{s}) \in \mathcal{S}_v$ are distinct itineraries. We define the formal triod map $\mathcal{T}_{\mathcal{S}}$ on the level of external addresses as follows:

$$\mathcal{T}_{\mathcal{S}}[\underline{t}, \underline{u}, \underline{v}] := \begin{cases} [\sigma(\underline{t}), \sigma(\underline{u}), \sigma(\underline{v})] & \text{if } \underline{t}, \underline{u}, \underline{v} \in I_k \text{ for some } k \in \mathbb{Z}, \\ [\sigma(\underline{t}), \sigma(\underline{u}), \underline{s}] & \text{if } \underline{t}, \underline{u} \in I_k \text{ for some } k \in \mathbb{Z}, \underline{v} \notin I_k, \\ [\sigma(\underline{t}), \underline{s}, \sigma(\underline{t})] & \text{if } \underline{t}, \underline{v} \in I_k \text{ for some } k \in \mathbb{Z}, \underline{u} \notin I_k, \\ [\underline{s}, \sigma(\underline{u}), \sigma(\underline{v})] & \text{if } \underline{u}, \underline{v} \in I_k \text{ for some } k \in \mathbb{Z}, \underline{t} \notin I_k, \\ \text{stop} & \text{otherwise.} \end{cases}$$

We claim that if the **stop** case is not reached, the image is again a formal triod. The following result will help us to show this fact.

LEMMA 5.4. (Order-preserving restrictions) *Let $I = (k\underline{s}, (k + 1)\underline{s})$ be a partition sector and let $I^- := (k\underline{s}, (k + 1)\underline{s}]$. The restriction*

$$\sigma|_{I^-} : I^- \rightarrow \mathcal{S}$$

is a bijection and preserves the cyclic order.

Proof. The shift map σ is strictly monotonically increasing with respect to the linear order on \mathcal{S} on both of the sets $I_l := \{\underline{t} \in I : \underline{t} = k \dots\}$ and $I_u := \{\underline{t} \in I : \underline{t} = (k + 1) \dots\} \cup \{(k + 1)\underline{s}\}$. It maps I_l bijectively onto $\{\underline{t} \in \mathcal{S} : \underline{t} > \underline{s}\}$ and I_u bijectively onto $\{\underline{t} \in \mathcal{S} : \underline{t} \leq \underline{s}\}$. In particular, it swaps the two sets globally, that is, $\sigma(\underline{t}) > \sigma(\underline{u})$ for all $\underline{t} \in I_l$ and all $\underline{u} \in I_u$. See Figure 2 for a sketch of the mapping behavior of σ . It follows that $\sigma|_{I^-}$ is a bijection. One can see that $\sigma|_{I^-}$ preserves the cyclic order by checking all possible configurations of the three external addresses with respect to I_l and I_u . See Figure 2 for a configuration under which the cyclic order, but not the linear order, of the addresses $\underline{t}, \underline{u}$, and \underline{v} is preserved. □

Let us write $[\tilde{\underline{t}}, \tilde{\underline{u}}, \tilde{\underline{v}}] := \mathcal{T}_{\mathcal{S}}[\underline{t}, \underline{u}, \underline{v}]$. Lemma 5.4 shows that $\tilde{\underline{t}} < \tilde{\underline{u}} < \tilde{\underline{v}}$. It remains to show that $\tilde{\underline{x}} := \text{It}(\tilde{\underline{x}} | \underline{s})$ are distinct for $x \in \{t, u, v\}$. The map $[\underline{t}, \underline{u}, \underline{v}] \mapsto [\underline{t}, \underline{u}, \underline{v}]$ (where $\underline{x} := \text{It}(\underline{x} | \underline{s})$ as before) is a semi-conjugation between $\mathcal{T}_{\mathcal{S}}$ and \mathcal{T} . The sequences $\tilde{\underline{x}}$ are distinct because $[\tilde{\underline{t}}, \tilde{\underline{u}}, \tilde{\underline{v}}] = \mathcal{T}[\underline{t}, \underline{u}, \underline{v}]$, and triods of formal (pre-)periodic points are




FIGURE 2. Sketch illustrating the mapping behavior of the shift map. The left sketch shows how the intervals I_l and I_u are exchanged globally, while the restriction of the shift to each of them is strictly monotonically increasing. The right figure shows that the linear order of the image addresses might change, but the cyclic order stays the same. It also illustrates that an interval splits if it contains the address s .

mapped to triods of formal (pre-)periodic points under iteration of \mathcal{T} by the considerations after Definition 4.12. We call the triods $[t, u, v]$ and $[\underline{t}, \underline{u}, \underline{v}]$ associated to each other. We define the *middle point* $\underline{b}[t, u, v]$ to be the middle point of $[\underline{t}, \underline{u}, \underline{v}]$, and the majority vote $\underline{m}[t, u, v]$ to be the majority vote of $[\underline{t}, \underline{u}, \underline{v}]$. We call $[v, t, u]$ branched if $[\underline{t}, \underline{u}, \underline{v}]$ is branched, and linear otherwise. If $\mathcal{T}^{oi_0}[\underline{t}, \underline{u}, \underline{v}] = \mathbf{stop}$, then $[t, u, v]$ reaches the **stop** case at the same iteration step and *vice versa*. We have seen in §4 that if $[\underline{t}, \underline{u}, \underline{v}]$ can be iterated indefinitely, all of the three sequences will eventually be contained in $\mathcal{S}_v \setminus \Omega^-(\star v)$, and $[\underline{t}, \underline{u}, \underline{v}]$ is (pre-)periodic under iteration of \mathcal{T} . By Theorem 2.10 and Proposition 2.13, there are only finitely many external addresses associated to a formal (pre-)periodic point $\underline{t} \in \mathcal{S}_v \setminus \Omega^-(\star v)$. Therefore, $[t, u, v]$ is eventually periodic under iteration of \mathcal{T}_S .

For the proof of the existence of separating rays, which works on the combinatorial level, the following results will be useful.

LEMMA 5.5. (Pullbacks of intervals) *Let $[t^1, t^2, t^3]$ be a formal triod of external addresses which can be iterated at least once before reaching the **stop** case, let $[u^1, u^2, u^3] := \mathcal{T}_S[t^1, t^2, t^3]$ be its image triod, and let $k := m(t^1, t^2, t^3)$ be its majority vote. Write $J_n := (t^n, t^{n+1})$ and $J'_n := (u^n, u^{n+1})$ for $n \in \{1, 2, 3\}$ (with indices labeled mod 3) for the intervals of the partition of the shift space by the initial triod and the image triod, respectively. Then we have*

$$\sigma^{-1}(J'_n) \cap I_k \subseteq J_n.$$

Proof. Because not all t^i are contained in distinct partition sectors by hypothesis, we have $t^j \notin I_k$ for at most one j . We define a triod $[\tilde{t}^1, \tilde{t}^2, \tilde{t}^3]$ by replacing such a t^j (if any) by the unique pre-image of s in I_k^- . Then $\mathcal{T}_S[t^1, t^2, t^3] = \mathcal{T}_S[\tilde{t}^1, \tilde{t}^2, \tilde{t}^3] = [\sigma(\tilde{t}^1), \sigma(\tilde{t}^2), \sigma(\tilde{t}^3)]$.

Furthermore, we have $\tilde{J}_n \cap I_k = J_n \cap I_k$, where $\tilde{J}_n := (\tilde{t}^n, \tilde{t}^{n+1})$. Let $\underline{v} \in \sigma^{-1}(J'_n) \cap I_k$ be an external address. We have $\sigma(\underline{v}) \in J'_n$, that is, $\sigma(\tilde{t}^n) < \sigma(\underline{v}) < \sigma(\tilde{t}^{n+1})$. It follows from Lemma 5.4 that $\tilde{t}^n < \underline{v} < \tilde{t}^{n+1}$ and, therefore, $\underline{v} \in \tilde{J}_n \cap I_k \subset J_n$. \square

LEMMA 5.6. (Splitting of intervals) *Let $I, I' \in \mathcal{I}$ be partition sectors and let $J \subset I$ be an interval. If $\underline{s} \notin J$, then $J' := \sigma^{-1}(J) \cap I'$ is an interval and it is of the form $J' = \{k\underline{t} : \underline{t} \in J\}$ for some $k \in \mathbb{Z}$.*

Proof. We have $I' = (k\underline{s}, (k+1)\underline{s})$ for some $k \in \mathbb{Z}$. As J is entirely contained in some partition sector and $\underline{s} \notin J$, we either have $\underline{t} > \underline{s}$ for all $\underline{t} \in J$ or $\underline{t} < \underline{s}$ for all $\underline{t} \in J$. Assume that the first case is true (the second case works analogously). Then we have $k\underline{t} \in I'$ for all $\underline{t} \in J$. As $\sigma|_{I'}$ is injective, we have $J' = \{k\underline{t} : \underline{t} \in J\}$. The right part of Figure 2 illustrates why an interval containing \underline{s} splits. \square

LEMMA 5.7. (Unlinked addresses) *Let $\underline{t}, \underline{u} \in \mathbb{Z}^{\mathbb{N}} \cup \Omega^-(\star v)$ be distinct (pre-)periodic itineraries, let $T := (\underline{t}^i)_{i \in \mathbb{I}}$ be the set of external addresses with $\text{It}(\underline{t}^i | \underline{s}) = \underline{t}$, and let $U := (\underline{u}^j)_{j \in J}$ be the set of external addresses with $\text{It}(\underline{u}^j | \underline{s}) = \underline{u}$. Then T and U are unlinked in the sense that there are no addresses $\underline{t}, \underline{t}' \in T$ and $\underline{u}, \underline{u}' \in U$ such that $\underline{t} < \underline{u} < \underline{t}' < \underline{u}'$.*

Proof. Assume to the contrary that there exist four external addresses as above. First consider the case that $\underline{t}_1 = \star$. Then \underline{t} and \underline{t}' are contained in the boundary $\partial\mathcal{I}$ of the dynamical partition, so \underline{u} and \underline{u}' are contained in different partition sectors, which contradicts the fact that they have the same itinerary. Hence, we can assume $\underline{t}_1, \underline{u}_1 \neq \star$. However then, $\underline{t} < \underline{u} < \underline{t}' < \underline{u}'$ implies $\underline{t}_1 = \underline{u}_1$. As $\sigma|_{I_{\underline{t}_1}}$ preserves the cyclic order, we have $\sigma(\underline{t}) < \sigma(\underline{u}) < \sigma(\underline{t}') < \sigma(\underline{u}')$. Repeating this argument inductively, we obtain $\underline{u} = \underline{t}$, which contradicts our assumptions. \square

Let us now prove the main combinatorial lemma of this section. Major ideas for the proof are taken from [SZ2, Lemma 5.2].

LEMMA 5.8. (Combinatorial version of separating dynamic rays) *Let $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ be a triod of external addresses and let $\underline{b} \in \mathbb{Z}^{\mathbb{N}} \cup \Omega^-(\star v)$ be (pre-)periodic.*

The triod $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ is branched with $\underline{b} = \underline{b}[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ if and only if there are three distinct (pre-)periodic external addresses $\underline{s}^1, \underline{s}^2, \underline{s}^3 \in \mathcal{S}$ such that $\text{It}(\underline{s}^i | \underline{s}) = \underline{b}$ and $\underline{s}^i \in (\underline{t}^i, \underline{t}^{i+1})$ (where again indices are labeled modulo 3).

The triod $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ is linear with $\underline{b}[\underline{t}^1, \underline{t}^2, \underline{t}^3] = \text{It}(\underline{t}^j | \underline{s}) =: \underline{b}$ if and only if there are two distinct (pre-)periodic external addresses $\underline{s}^j, \underline{s}^{j+1} \in \mathcal{S}$ with $\underline{s}^j = \underline{t}^j$ and $\underline{s}^{j+1} \in (\underline{t}^{j+1}, \underline{t}^{j+2})$, such that $\text{It}(\underline{s}^i | \underline{s}) = \underline{b}$ holds.

Proof. We start by proving the ‘only if’ direction.

CLAIM 1. *If $\mathcal{T}_{\mathcal{S}}[\underline{t}^1, \underline{t}^2, \underline{t}^3] = [\underline{u}^1, \underline{u}^2, \underline{u}^3]$ and the result is true for $[\underline{u}^1, \underline{u}^2, \underline{u}^3]$, then it also holds for $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$.*

Proof. Assume first that $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ and, hence also, $[\underline{u}^1, \underline{u}^2, \underline{u}^3]$ is branched, and let \underline{s}^j be separating addresses for the image triod. Then, setting $k := m(\underline{t}^1, \underline{t}^2, \underline{t}^3)$ and $\underline{s}^j :=$

$\sigma|_{I_k}^{-1}(\underline{s}^j)$, we have $\underline{s}^j \in (\underline{t}^j, \underline{t}^{j+1})$ by Lemma 5.5. Furthermore, we have $\text{It}(\underline{s}^j | \underline{s}) = \underline{b}[\underline{t}^1, \underline{t}^2, \underline{t}^3]$, so the addresses \underline{s}^j are separating addresses for the triod $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$. The linear case works analogously and is left to the reader. \square

CLAIM 2. *The result is true if $\underline{b}[\underline{t}^1, \underline{t}^2, \underline{t}^3] \in \Omega^-(\star\nu)$.*

Proof. There exists an $n \geq 0$ such that $\mathcal{T}_S[\underline{t}^1, \underline{t}^2, \underline{t}^3] = \text{stop}$. Assume first that $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ is branched. The image triod $[\underline{u}^1, \underline{u}^2, \underline{u}^3] := \mathcal{T}_S^{\circ(n-1)}[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ consists of external addresses lying in distinct sectors of the dynamical partition \mathcal{I} . Therefore, there exists addresses $\underline{s}^i := k_i \underline{s}$ for suitable $k_i \in \mathbb{Z}$ such that $\underline{s}^i \in (\underline{u}^i, \underline{u}^{i+1})$, and we have $\text{It}(\underline{s}^i | \underline{s}) = \star\nu = \underline{b}[\underline{u}^1, \underline{u}^2, \underline{u}^3]$ for every $i \in \{1, 2, 3\}$. Hence, we have proved the lemma for the image triod $[\underline{u}^1, \underline{u}^2, \underline{u}^3]$. By Claim 1, it also holds for $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$. The linear case works analogously and is left to the reader. \square

Now assume that $\underline{b}[\underline{t}^1, \underline{t}^2, \underline{t}^3] \notin \Omega^-(\star\nu)$, so the triod $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ can be iterated indefinitely under \mathcal{T}_S . Every such triod is eventually periodic, so by Claim 1, we can assume without loss of generality that $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ is periodic, say of period N . Hence, $\underline{b}[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ is also periodic (under iteration of \mathcal{T}), possibly of smaller period. Let us write $\underline{b} := \underline{b}[\underline{t}^1, \underline{t}^2, \underline{t}^3] =: \underline{b}_0 \underline{b}_1 \dots \underline{b}_{n-1}$, where the indices of \underline{b} are labeled modulo n .

CLAIM 3. *For every $i \in \{1, 2, 3\}$, there exists an address $\underline{s}^i \in \text{cl}(\underline{t}^i, \underline{t}^{i+1})$ satisfying $\text{It}(\underline{s}^i | \underline{s}) = \underline{b}$.*

Proof. Our strategy for the proof is to pull back a suitable interval along the inverse branches of σ prescribed by \underline{b} to obtain a nested sequence of intervals whose intersection is an external address with the desired properties. The forward orbit $T := \Omega^+(\underline{s})$ is finite, so $I_{\underline{b}_0} \setminus T$ consists of a finite number of disjoint open intervals. Write $J_{1,i}^{(0)}, \dots, J_{n_i,i}^{(0)}$ for the subintervals of $(\underline{t}^i, \underline{t}^{i+1})$ contained in $I_{\underline{b}_0} \setminus T$. Inductively, we define

$$J_{j,i}^{(m)} := \sigma^{-1}(J_{j,i}^{(m-1)}) \cap I_{\underline{b}_{-m}}.$$

As $\mathcal{S} \setminus T$ is backward invariant, we have $J_{j,i}^{(m)} \cap T = \emptyset$, so every set $J_{j,i}^{(m)}$ is an interval by Lemma 5.6. Writing $[\underline{t}^{1,(j)}, \underline{t}^{2,(j)}, \underline{t}^{3,(j)}] := \mathcal{T}_S^{\circ j}[\underline{t}^1, \underline{t}^2, \underline{t}^3]$, we have $J_{j,i}^{(m)} \subseteq (\underline{t}^{i,(-m)}, \underline{t}^{i+1,(-m)})$ by Lemma 5.5.

After N pullbacks, we arrive at subintervals of the initial partition sector $I_{\underline{b}_0}$ satisfying $J_{j,i}^{(N)} \subseteq (\underline{t}^i, \underline{t}^{i+1})$. Hence, each $J_{j,i}^{(N)}$ must be a subinterval of one of our initial intervals $J_{j,i}^{(0)}$. More precisely, we get three self-maps $\rho_i : \{1, \dots, n_i\} \rightarrow \{1, \dots, n_i\}$ such that

$$J_{j,i}^{(N)} \subseteq J_{\rho_i(j),i}^{(0)} \quad \text{for each } i \in \{1, 2, 3\} \text{ and } j \in \{1, \dots, n_i\}.$$

For each $i \in \{1, 2, 3\}$, we can find a $j_i \in \{1, \dots, n_i\}$ and an $l_i \in \mathbb{N}$ such that $\rho_i^{ol_i}(j_i) = j_i$, and setting $N_i := l_i \cdot N$, we have $J_{j_i,i}^{(N_i)} \subset J_{j_i,i}^{(0)}$. Furthermore, the first N_i entries are the same for all $\underline{t} \in J_{j_i,i}^{(N_i)}$, say equal to $s_0^{(i)} s_1^{(i)} \dots s_{N_i-1}^{(i)}$. Setting $J_i^{(m)} := J_{j_i,i}^{(mN_i)}$, the $J_i^{(m)}$ form a nested sequence of (open) intervals and every address $\underline{t} \in J_i^{(m)}$ begins with m times the sequence $s_0^{(i)} s_1^{(i)} \dots s_{N_i-1}^{(i)}$. We define $\underline{s}^i := s_0^{(i)} s_1^{(i)} \dots s_{N_i-1}^{(i)}$. Recall that for an

address $\underline{t} \in \mathcal{S}$, the cylinder sets $\{\underline{u} \in \mathcal{S} : u_0 \dots u_i = t_0 \dots t_i\}$ form an open neighborhood basis of \underline{t} . As every address $\underline{t} \in J_i^{(m)}$ begins with m times the sequence $s_0^{(i)} s_1^{(i)} \dots s_{N_i-1}^{(i)}$, we see that $U \cap J_i^{(m)} \neq \emptyset$ for every open neighborhood U of \underline{s}^i and every $m \geq 0$. Hence, we have $\underline{s}^i \in \bigcap_{m \in \mathbb{N}} \text{cl}(J_i^{(m)})$. This implies $\sigma^{om}(\underline{s}^i) \in \overline{I_{\mathfrak{b}_m}}$ for all $m \in \mathbb{N}_0$. However, we have $\sigma^{om}(\underline{s}^i) \notin \partial I_{\mathfrak{b}_m}$ because otherwise $\sigma^{om}(\underline{s}^i)$ would be strictly pre-periodic, which contradicts the fact that \underline{s}^i is periodic. Hence, we have $\text{It}(\underline{s}^i | \mathfrak{s}) = \underline{\mathfrak{b}}$. \square

If $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ is branched, we actually have $\underline{s}^i \in (\underline{t}^i, \underline{t}^{i+1})$ because $\text{It}(\underline{t}^i | \mathfrak{s}) \neq \underline{\mathfrak{b}}$ for all i . If instead $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ is linear with \underline{t}^j in the middle, we have $\underline{s}^{j+1} \in (\underline{t}^{j+1}, \underline{t}^{j+2})$ because $\text{It}(\underline{s}^{j+1} | \mathfrak{s}) = \underline{\mathfrak{b}}$ is distinct from both $\text{It}(\underline{t}^{j+1} | \mathfrak{s})$ and $\text{It}(\underline{t}^{j+2} | \mathfrak{s})$. Hence, \underline{t}^j and \underline{s}^{j+1} form separating external addresses in this case. This finishes the proof of the ‘only if’ direction.

Let us now prove the other direction. Assume first that there exists a (pre-)periodic $\underline{\mathfrak{b}} \in \mathbb{Z}^{\mathbb{N}} \cup \Omega^-(\star\nu)$ and three distinct (pre-)periodic external addresses $\underline{s}^1, \underline{s}^2, \underline{s}^3 \in \mathcal{S}$ such that $\text{It}(\underline{s}^i | \mathfrak{s}) = \underline{\mathfrak{b}}$ and $\underline{s}^i \in (\underline{t}^i, \underline{t}^{i+1})$. We want to see that $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ is branched with $\underline{\mathfrak{b}}[\underline{t}^1, \underline{t}^2, \underline{t}^3] = \underline{\mathfrak{b}}$. Assume to the contrary that $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ is branched with $\underline{\mathfrak{b}}[\underline{t}^1, \underline{t}^2, \underline{t}^3] \neq \underline{\mathfrak{b}}$. By the above, there are three distinct (pre-)periodic external addresses $\underline{\tilde{s}}^1, \underline{\tilde{s}}^2, \underline{\tilde{s}}^3 \in \mathcal{S}$ such that $\text{It}(\underline{\tilde{s}}^i | \mathfrak{s}) = \underline{\mathfrak{b}}[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ and $\underline{\tilde{s}}^i \in (\underline{t}^i, \underline{t}^{i+1})$. This contradicts Lemma 5.7 because the sets $\{\underline{s}^i\}$ and $\{\underline{\tilde{s}}^j\}$ would not be unlinked. The same lemma leads to a contradiction if we assume that $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ is linear. The linear case works analogously and is left to the reader. \square

Let us translate Lemma 5.8 into the language of dynamic rays and their landing points. This requires Lemma 5.1 which is why we prove it now.

Proof of Lemma 5.1. The formal vertex set V_ν is forward invariant under σ because $\Omega^+(\star\nu)$ is forward invariant and $\sigma(\underline{\mathfrak{b}}(\underline{\mathfrak{t}}, \underline{\mathfrak{u}}, \underline{\mathfrak{v}})) = \underline{\mathfrak{b}}(\mathcal{T}[\underline{\mathfrak{t}}, \underline{\mathfrak{u}}, \underline{\mathfrak{v}}])$ as long as $\mathcal{T}[\underline{\mathfrak{t}}, \underline{\mathfrak{u}}, \underline{\mathfrak{v}}] \neq \mathbf{stop}$. Also, if $\mathcal{T}[\underline{\mathfrak{t}}, \underline{\mathfrak{u}}, \underline{\mathfrak{v}}] = \mathbf{stop}$, then $\underline{\mathfrak{b}}(\underline{\mathfrak{t}}, \underline{\mathfrak{u}}, \underline{\mathfrak{v}}) = \star\nu \in V_\nu$.

For the proof of (2), assume to the contrary that there exists a triod $[\underline{\tilde{\mathfrak{t}}}, \underline{\tilde{\mathfrak{u}}}, \underline{\tilde{\mathfrak{v}}}]$ with $\underline{\tilde{\mathfrak{t}}}, \underline{\tilde{\mathfrak{u}}}, \underline{\tilde{\mathfrak{v}}} \in \Omega^+(\star\nu)$ and $\underline{\mathfrak{b}}(\underline{\tilde{\mathfrak{t}}}, \underline{\tilde{\mathfrak{u}}}, \underline{\tilde{\mathfrak{v}}}) \notin \mathcal{S}_\nu$. Because $\underline{\tilde{\mathfrak{t}}}, \underline{\tilde{\mathfrak{u}}}, \underline{\tilde{\mathfrak{v}}}$ are all (pre-)periodic, so is $\underline{\mathfrak{b}}(\underline{\tilde{\mathfrak{t}}}, \underline{\tilde{\mathfrak{u}}}, \underline{\tilde{\mathfrak{v}}})$, and we must have $\underline{\mathfrak{b}}(\underline{\tilde{\mathfrak{t}}}, \underline{\tilde{\mathfrak{u}}}, \underline{\tilde{\mathfrak{v}}}) = k_1 \dots k_n \nu$ (because all other (pre-)periodic sequences are contained in \mathcal{S}_ν). Iterating forward n times, we obtain another triod $[\underline{\mathfrak{t}}, \underline{\mathfrak{u}}, \underline{\mathfrak{v}}] := \mathcal{T}^{\circ n}[\underline{\tilde{\mathfrak{t}}}, \underline{\tilde{\mathfrak{u}}}, \underline{\tilde{\mathfrak{v}}}]$ with $\underline{\mathfrak{t}}, \underline{\mathfrak{u}}, \underline{\mathfrak{v}} \in \Omega^+(\nu)$ and $\underline{\mathfrak{b}}(\underline{\mathfrak{t}}, \underline{\mathfrak{u}}, \underline{\mathfrak{v}}) = \nu$. We can assume without loss of generality that $\underline{\mathfrak{t}} = \sigma^{\circ j}(\nu)$, $\underline{\mathfrak{u}} = \sigma^{\circ k}(\nu)$, and $\underline{\mathfrak{v}} = \sigma^{\circ l}(\nu)$, where $j, k > 0$ and $l \geq 0$. The triod $[\underline{\mathfrak{t}}, \underline{\mathfrak{u}}, \underline{\mathfrak{v}}]$ might be branched or linear. In both cases, by passing to an associated triod of external addresses, Lemma 5.8 guarantees the existence of external addresses $\underline{s}', \underline{t}, \underline{u} \in \mathcal{S}$ with the following properties.

- We have $\text{It}(\underline{s}' | \mathfrak{s}) = \nu$, $\text{It}(\underline{t} | \mathfrak{s}) = \sigma^{\circ j}(\nu) = \underline{\mathfrak{t}}$ and $\text{It}(\underline{u} | \mathfrak{s}) = \sigma^{\circ k}(\nu) = \underline{\mathfrak{u}}$.
- We have $\underline{t} \in (\underline{s}, \underline{s}')$ and $\underline{u} \in (\underline{s}', \underline{s})$.

By Proposition 2.13, the dynamic rays $g_{\underline{s}}$ and $g_{\underline{s}'}$ land at 0, while $g_{\underline{t}}$ lands at $p = E_\lambda^{\circ j}(0)$ and $g_{\underline{u}}$ lands at $q = E_\lambda^{\circ k}(0)$. By Lemma 2.9, the relations $\underline{t} \in (\underline{s}, \underline{s}')$ and $\underline{u} \in (\underline{s}', \underline{s})$ imply that p and q are contained in different connected components of $\mathbb{C} \setminus (\text{Tr}(g_{\underline{s}}) \cup \text{Tr}(g_{\underline{s}'}) \cup \{0\})$. However then, for one of these points, say p , the itineraries of the pre-images of p have different initial entries with respect to the dynamical partitions \mathcal{D} induced by $g_{\underline{s}}$ and \mathcal{D}' induced by $g_{\underline{s}'}$. At least one pre-image of p is itself a post-singular point and it has different itineraries with respect to the two dynamical partitions, which contradicts Lemma 4.1.




FIGURE 3. Sketch of the separating sets A_v associated to the embedded vertices for the exponential map $E_{i\pi}$.

Finally, let us prove (3). We can assume that $[\underline{t}, \underline{u}, \underline{v}]$ is branched because otherwise, the middle point of the triod is contained in V_v anyway. Let $T := (\underline{t}^i)_{i \in I}$, $U := (\underline{u}^j)_{j \in J}$ and $V := (\underline{v}^k)_{k \in K}$ be the sets of external addresses with itinerary \underline{t} , \underline{u} , and \underline{v} , respectively. By Lemma 5.8 and Lemma 5.7, there are three external addresses $\underline{s}^1, \underline{s}^2, \underline{s}^3$ (without loss of generality indexed with respect to their cyclic order) with itinerary $\underline{b}(\underline{t}, \underline{u}, \underline{v})$ such that $T \subset (\underline{s}^1, \underline{s}^2)$, $U \subset (\underline{s}^2, \underline{s}^3)$, and $V \subset (\underline{s}^3, \underline{s}^1)$. If $\underline{t} \in \Omega^+(\star\nu)$, we set $\underline{t}' := \underline{t}$. Otherwise, \underline{t} is the branch point of a triod $[\underline{t}^1, \underline{t}^2, \underline{t}^3]$ of post-singular points $\underline{t}^i \in \Omega^+(\star\nu)$ and by Lemma 5.8, we can find an external address $\underline{t} \in (\underline{s}^1, \underline{s}^2)$ with itinerary \underline{t}^j for a suitable value of $j \in \{1, 2, 3\}$. We set $\underline{t}' := \underline{t}^j$. Proceeding with \underline{u} and \underline{v} in the same way, we obtain a triod $[\underline{t}', \underline{u}', \underline{v}']$ of post-singular points $\underline{t}', \underline{u}', \underline{v}' \in \Omega^+(\star\nu)$, and it follows from Lemma 5.8 (the if-direction) that $\underline{b}(\underline{t}', \underline{u}', \underline{v}') = \underline{b}(\underline{t}, \underline{u}, \underline{v})$. We have $\underline{b}(\underline{t}', \underline{u}', \underline{v}') \in V_v$ by definition of the formal vertex set. \square

Definition 5.9. (Separating sets) If $v = v_{\underline{t}}$ is a pre-singular vertex,

$$A_v := \bigcup_{\text{It}(\underline{t} | \underline{s}) = \underline{t}} \text{Tr}(g_{\underline{t}}) \cup U_{\underline{t}}.$$

If v is not a pre-singular vertex,

$$A_v := \bigcup_{\text{It}(\underline{t} | \underline{s}) = \underline{t}} \text{Tr}(g_{\underline{t}}) \cup \{v\}.$$

It is clear (by the properties (1) and (2) imposed on U) that if $v, w \in V$ are different vertices, then $A_v \cap A_w = \emptyset$. See Figure 3 for a sketch of the separating sets A_v for a particularly simple pspf exponential map.

LEMMA 5.10. (Separating dynamic rays) *Let $[p, q, r]$ be a triod of vertex points. The triod $[p, q, r]$ is linear with p in the middle if and only if q and r lie in different connected components of $\mathbb{C} \setminus A_p$, and it is branched with branch point $b = b[p, q, r] \in V$ if and only if $p, q,$ and r lie in different connected components of $\mathbb{C} \setminus A_b$.*

Proof. This is just a matter of translating Lemma 5.8 and Lemma 5.1 into the language of dynamic rays and their landing points using Lemma 2.9 and Proposition 2.13. \square

We continue our construction of homotopy Hubbard trees. Choose a point $p_0 \in \partial U \setminus \text{Tr}(g_{\underline{s}})$ and write $p_{\star v, j}$ for the unique pre-image of p (under E_λ) in the sector D_j of the dynamical partition. We choose disjoint arcs $\gamma_{p_{\star v, j}} : [0, 1] \rightarrow U_{\star v}$ from $v_{\star v}$ to $p_{\star v, j}$ such that $\gamma_{p_{\star v, j}}((0, 1)) \subset \text{int}(U_{\star v})$. These curves are there to normalize the ends of the edges of the yet to be constructed homotopy Hubbard tree that have $v_{\star v}$ as an endpoint. For a pre-singular vertex $v_{\underline{t}}$ with $\underline{t} = k_1 \dots k_n \star v$, we define the pulled-back path $\gamma_{p_{\underline{t}, j}} := \log_{\underline{s}, k_1} \dots \log_{\underline{s}, k_n}(\gamma_{p_{\star v, j}})$ with endpoints $v_{\underline{t}}$ and $p_{\underline{t}, j}$. The points $p_{\underline{t}, j}$ are just auxiliary points that will lie on the interior of an edge at the end of the construction.

We call an arc $\gamma : [0, 1] \rightarrow \mathbb{C}$ between two different points $v, w \in V$ allowable if the following conditions are satisfied.

- If $r \in V$ is not a pre-singular vertex, then either $\text{Tr}(\gamma) \cap A_r = \emptyset$ or $\text{Tr}(\gamma) \cap A_r = \{r\}$.
- If $r = v_{\underline{t}} \in V$ is a pre-singular vertex, then either $\text{Tr}(\gamma) \cap A_r = \emptyset$ or $\text{Tr}(\gamma) \cap A_r = \text{Tr}(\gamma_{p_{\underline{t}, i}})$ for some $i \in \mathbb{Z}$ if r is one of the endpoints of γ or $\text{Tr}(\gamma) \cap A_r = \text{Tr}(\gamma_{p_{\underline{t}, i}}) \cup \text{Tr}(\gamma_{p_{\underline{t}, j}})$ for indices $i \neq j$ if r lies in the interior of the path.

An allowable arc γ satisfying $\gamma^{-1}(V) = \{0, 1\}$ is called an *edge*. We call two vertices *incident* if they can be connected by an edge. Note that every allowable arc is a concatenation of finitely many edges; conversely, concatenating edges of a finite acyclic path always yields an allowable arc.

LEMMA 5.11. (Existence and uniqueness of allowable arcs) *For any pair of distinct vertices $w, w' \in V$, there exists an allowable arc $\gamma : [0, 1] \rightarrow \mathbb{C}$ with $\gamma(0) = w$ and $\gamma(1) = w'$. If $\tilde{\gamma} : [0, 1] \rightarrow \mathbb{C}$ is another allowable arc from w to w' , then γ and $\tilde{\gamma}$ are homotopic rel $\bigcup_{v \in V} A_v$.*

Proof. Let D be a connected component of $\mathbb{C} \setminus \bigcup_{v \in V} A_v$. We start by investigating the topology of D and ∂D . It is well known that if $A \subset \overline{\mathbb{C}}$ is closed and connected, every connected component of $\overline{\mathbb{C}} \setminus A$ is simply connected. As $\text{cl}_{\overline{\mathbb{C}}}(\bigcup_{v \in V} A_v)$ is closed and connected, D is simply connected. Let B be a connected component of ∂D . We have $B \subset \partial A_v$ for some $v \in V$. If $v = v_{\underline{t}}$ is a pre-singular vertex, then there exists a continuous bijection $\Gamma : (-\infty, \infty) \rightarrow B$ such that $\lim_{t \rightarrow -\infty} \Gamma(t) = \lim_{t \rightarrow +\infty} \Gamma(t) = \infty$ (in $\overline{\mathbb{C}}$) and $\Gamma(0) = p_{\underline{t}, j}$ for some $j \in \mathbb{Z}$. We call $p_{\underline{t}, j}$ the *distinguished boundary point* of D on A_v . If v is not pre-singular, and there is only one dynamic ray g landing at v , then we have $B = \partial D \cap A_v = \text{Tr}(g) \cup \{v\}$. If there are at least two dynamic rays landing at v , we have $B = \partial D \cap A_v = \text{Tr}(g) \cup \text{Tr}(g') \cup \{v\}$ for some rays g, g' landing at v . In both cases, we call v the *distinguished boundary point* of D on A_v .

We claim that ∂D has at most two connected components. Assume to the contrary that there are vertices $v_1, v_2, v_3 \in V$ such that $\partial D \cap A_{v_i} \neq \emptyset$ for all i . Then, we can

connect any two distinct v_i and v_j by a curve $\gamma_{(i,j)}$ such that $\text{Tr}(\gamma) \cap A_v = \emptyset$ for all $v \in V \setminus \{v_i, v_j\}$. Consider the triod $[v_1, v_2, v_3]$. If it is linear with middle point v_i , then $A_{v_{i-1}}$ and $A_{v_{i+1}}$ are separated by A_{v_i} by Lemma 5.10, and if $[v_1, v_2, v_3]$ is branched with branch point b , then each pair A_{v_i} and A_{v_j} is separated by A_b by Lemma 5.10. In both cases, we get a contradiction. Therefore, there are only three possibilities for the topology of $(D, \partial D)$: there exists a homeomorphism $\varphi: (D, \partial D) \rightarrow (\Omega_i, \partial\Omega_i)$, that is, a homeomorphism $\varphi: \overline{D} \rightarrow \overline{\Omega_i}$ satisfying $\varphi(\partial D) = \partial\Omega_i$ for exactly one of three uniformizing domains $\Omega_i \subset \mathbb{C}$, where:

- (1) $\Omega_1 := \{z \in \mathbb{C} : \text{Re}(z) > 0\}$ is the right half-plane;
- (2) $\Omega_2 := \{z \in \mathbb{C} : \text{Re}(z) > 0\} \setminus [1, +\infty)$ is the slit right half-plane;
- (3) and $\Omega_3 := \{z \in \mathbb{C} : 0 < \text{Re}(z) < 1\}$ is a vertical strip.

Let us now prove the existence and uniqueness of allowable arcs. First, consider the special case that there exists a connected component D of $\mathbb{C} \setminus \bigcup_{v \in V} A_v$ such that $\partial D \cap A_w \neq \emptyset$ and $\partial D \cap A_{w'} \neq \emptyset$. Let $p_w, p_{w'} \in \partial D$ be the distinguished boundary points. The domain D is of type (2) or (3), so there exists an arc $\delta: [0, 1] \rightarrow \overline{D}$ from p_w to $p_{w'}$ such that $\delta((0, 1)) \in D$, and δ is unique up to homotopy $\text{rel } \bigcup_{v \in V} A_v$. We obtain an edge γ connecting w and w' after concatenating δ with γ_{p_w} or $\gamma_{p_{w'}}$ if $p_w \neq w$ or $p_{w'} \neq w'$.

Finally, consider the general case and set $v_0 := w$. Let $E(v_0)$ be the component of $\mathbb{C} \setminus A_{v_0}$ containing w' , and let D_1 be the unique component of $\mathbb{C} \setminus \bigcup_{v \in V} A_v$ such that $\partial D_1 \cap A_{v_0} \neq \emptyset$ and $D_1 \subset E(v_0)$. The component D_1 is of type (2) or (3) because $E(v_0)$ contains at least the vertex w' , so there exists $v_1 \in V$ such that $\partial D_1 \cap A_{v_1} \neq \emptyset$. If $v_1 = w'$, we are reduced to the special case. Else D_1 is of type (3) and there is a unique component $E(v_1) \subset E(v_0)$ of $\mathbb{C} \setminus A_{v_1}$ containing w' . We continue this argument inductively, and as the number of vertices contained in $E(v_i)$ decreases in each step, we obtain a finite sequence $w = v_0, v_1, \dots, v_n = w'$ of vertices such that $A_{v_{i-1}}$ and A_{v_i} both intersect ∂D_i . The concatenation $\gamma := \gamma_0 \cdot \gamma_1 \cdot \dots \cdot \gamma_{n-1}$, where γ_i is an edge from v_i to v_{i+1} , is an allowable path from w to w' , and it is unique up to homotopy $\text{rel } \bigcup_{v \in V} A_v$ because every part of γ_i is unique. □

THEOREM 5.12. (Existence of homotopy Hubbard trees) *For every post-singularly finite exponential map E_λ , there exists a homotopy Hubbard tree.*

Proof. For every pair $w, w' \in V$ of incident vertices, choose an edge $e_{(w,w')}: [0, 1] \rightarrow \mathbb{C}$ connecting w and w' . We choose the unique edge $e_{(0,w_0)}: [0, 1] \rightarrow \mathbb{C}$ connecting 0 to another vertex w_0 in such a way that it intersects ∂U only at the point $p_0 \in \partial U \setminus \text{Tr}(g_s)$ (see the paragraph before Lemma 5.11 for the definition of p_0). Define

$$H := \bigcup_{w,w' \in V \text{ incident}} \text{Tr}(e_{(w,w')}).$$

As every edge is contained in some component of $\mathbb{C} \setminus \bigcup_{v \in V} A_v$ (except possibly for its normalized ends), and every such component contains at most one edge by Lemma 5.11, different edges are disjoint (except possibly for their endpoints). Clearly, they are not homotopic $\text{rel } \bigcup_{v \in V} A_v$. The embedded graph H is connected by Lemma 5.11, and the existence of a cycle of edges would contradict uniqueness in Lemma 5.11. Hence, H is

an embedded tree. Moreover, every endpoint of H is a post-singular point: every vertex $v \in V \setminus P(E_\lambda)$ is the middle point of a triod of the form $[p_1, v, p_2]$ with $p_1, p_2 \in P(E_\lambda)$ and the in-tree connection of p_1 and p_2 has to contain v by Lemma 5.10 and the definition of allowable arcs.

Next, we want to see that $H' := [P(E_\lambda)]_{\widehat{E}_\lambda^{-1}(H)}$ is homotopic to $H \text{ rel } P(E_\lambda)$. By Lemma 3.14, we can find a homotopy I_0 between H' and an embedded tree $H'' \subset \mathbb{C} \text{ rel } \mathbb{C}_T \setminus \widehat{E}_\lambda^{-1}(U)$, and by our choice of the edge $e_{(0, w_0)}$, we can make sure that $H'' \cap \widehat{E}_\lambda^{-1}(U) = H \cap \widehat{E}_\lambda^{-1}(U) = \bigcup_{i \in \{1, \dots, k\}} \text{Tr}(\gamma_{p_{\star v, j_i}})$ for certain $j_i \in \mathbb{Z}$.

For a vertex $v \neq v_{\star v}$, the map $E_\lambda: A_v \rightarrow A_{E_\lambda(v)}$ is a homeomorphism, and if $v = v_{\underline{v}}$ is pre-singular, we have $E_\lambda(\bigcup_j \text{Tr}(\gamma_{\underline{v}, j})) = \bigcup_j \text{Tr}(\gamma_{\sigma(\underline{v}), j})$. If v is not pre-singular, we have $H \cap A_{E_\lambda(v)} = \{E_\lambda(v)\}$, and therefore $\widehat{E}_\lambda^{-1}(H) \cap A_v = \{v\}$. If instead $v = v_{\underline{v}} \neq v_{\star v}$ is pre-singular, we have $H \cap A_{E_\lambda(v)} \subset \bigcup_j \text{Tr}(\gamma_{\sigma(\underline{v}), j})$, and therefore $\widehat{E}_\lambda^{-1}(H) \cap A_v \subset \bigcup_j \text{Tr}(\gamma_{\underline{v}, j})$. Hence, for every pair of distinct post-singular points $p, q \in P(E_\lambda)$, the in-tree connection $[p, q]_{H''}$ is an allowable arc.

It follows that $V \subset H''$, because every $v \in V \setminus (P(E_\lambda) \cup \{v_{\star v}\})$ is the branch point of a post-singular triod. Moreover, the trees H and H'' have the same graph structure on V . To see this, observe that H is defined to contain an edge for every pair of incident vertices. Hence, vertices that are connected by an edge in H'' are also connected by an edge in H . Conversely, the tree H'' must contain all edges in H because otherwise it would be disconnected. By the uniqueness of allowable arcs, there exists a homotopy I_1 between H'' and $H \text{ rel } \bigcup_{v \in V} A_v$. We obtain a homotopy I between H and $H' \text{ rel } P(E_\lambda)$ by concatenating I_0 and I_1 .

It remains to show that the induced self-map $f := \widehat{E}_\lambda \circ I^1: H \rightarrow H$ is expansive (see the paragraph before Definition 3.9 for the definition of expansivity). By construction, we have $I^1(v_{\star v}) = -\infty$ and we deduce that $V_f = V$ for the set of marked points of the self-map f . Let $p, q \in V_f = V$ be two different marked points of f . By construction, two vertices $v_{\underline{v}}, v_{\underline{u}} \in V \setminus \{v_{\star v}\}$ are contained in different branches of H at $v_{\star v}$ if and only if the first entries of \underline{v} and \underline{u} are different. As different vertices have different itineraries, there exists a small $n \geq 0$ such that the itineraries of $f^{on}(p)$ and $f^{on}(q)$ have different initial entries. Hence, we have $v_{\star v} \in [f^{on}(p), f^{on}(q)] \subset f^{on}[p, q]$. □

6. Classification of post-singularly finite exponential maps

Post-singularly finite exponential maps have been classified in terms of the external addresses of the dynamic rays landing at the singular value in [LSV], and we have stated their main result in Theorem 2.14. As an application of our construction, we give another combinatorial classification of this class of maps in terms of abstract Hubbard trees. Post-singularly finite polynomials have been classified combinatorially by Poirier in [P] in terms of so-called *abstract Hubbard trees*. These are graph-theoretic trees equipped with a self-map and certain extra information, but without an embedding into \mathbb{C} . The exponential dynamical trees from Definition 4.2 come quite close to what we mean by an abstract exponential Hubbard tree. We have to add a cyclic order on the branches of the tree at each marked point to specify a homotopy type of embeddings of the tree into the complex plane.

Definition 6.1. (Abstract exponential Hubbard tree) An *abstract exponential Hubbard tree* $(H, f, B_{k_i}, \angle_v)$ is an exponential dynamical tree (H, f, B_{k_i}) together with a cyclic order \angle_v on the set of branches of H at v for each marked point $v \in V_f \setminus \{v_T\}$ such that f preserves the cyclic order. Two abstract Hubbard trees $(H, f, B_{k_i}, \angle_v)$ and $(\tilde{H}, \tilde{f}, \tilde{B}_{k_i}, \angle_{\tilde{v}})$ are called *equivalent* if they are equivalent as exponential dynamical trees and, in addition, the homeomorphism $\varphi: H \rightarrow \tilde{H}$ from Definition 4.2 can be chosen to preserve the cyclic order at marked points.

Let H be a homotopy Hubbard tree for the post-singularly finite exponential map E_λ . In §4, we have seen that for any choice of induced self-map $f: H \rightarrow H$, the triple (H, f, B_{k_i}) is an exponential dynamical tree. We have also seen that if \tilde{H} is another homotopy Hubbard tree for E_λ with choice of induced self-map $\tilde{f}: \tilde{H} \rightarrow \tilde{H}$, the exponential dynamical trees (H, f, B_{k_i}) and $(\tilde{H}, \tilde{f}, \tilde{B}_{k_i})$ are equivalent. As H is embedded into the complex plane, the branches of H at a marked point $v \in V_f \setminus \{v_T\}$ come equipped with a natural cyclic order \angle_v (of course this also holds for v_T , but the cyclic order of edges at v_T is already contained in the sector information). Also, the trees H and \tilde{H} are homotopic in $\mathbb{C} \text{ rel } P(E_\lambda)$, and such a homotopy preserves the cyclic order of branches at marked points. It follows that there is a well-defined map

$$\mathcal{F}: \{\text{psf exponential maps}\} \rightarrow \{\text{abstract exponential Hubbard trees}\},$$

where we consider abstract exponential Hubbard trees up to equivalence. In this section, we prove that \mathcal{F} is a bijection.

The major tool that has made classification results of this kind possible for certain classes of rational functions is Thurston’s topological characterization of rational maps. An analogous result for exponential maps has been established in [HSS]. To state this result precisely, we need to introduce some terminology.

Convention 6.2. In the following, \mathbb{S}^2 denotes an oriented topological 2-sphere with two distinguished points 0 and ∞ . All homeomorphisms and coverings will be understood to be orientation preserving. We also write $\mathbb{R}^2 = \mathbb{S}^2 \setminus \{\infty\}$.

Definition 6.3. (Topological exponential maps) A covering map $g: \mathbb{R}^2 \rightarrow \mathbb{R}^2 \setminus \{0\}$ is called a *topological exponential map*. It is called *post-singularly finite* if the orbit of 0 is finite, hence pre-periodic. The *post-singular set* is $P(g) := \bigcup_{n \geq 0} g^{on}(0)$.

Definition 6.4. (Thurston equivalence) Two post-singularly finite topological exponential maps f and g with post-singular sets $P(f)$ and $P(g)$ are called *Thurston equivalent* if there are two homeomorphisms $\varphi_0, \varphi_1: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ satisfying $\varphi_0|_{P(f)} = \varphi_1|_{P(f)}$ and $P(g) = \varphi_0(P(f)) = \varphi_1(P(f))$ such that the diagram

$$\begin{CD} (\mathbb{R}^2, P(f)) @>\varphi_1>> (\mathbb{R}^2, P(g)) \\ @VfVV @VVgV \\ (\mathbb{R}^2, P(f)) @>\varphi_0>> (\mathbb{R}^2, P(g)) \end{CD}$$

commutes, and φ_0 is homotopic to φ_1 relative to $P(f)$.

The main result of [HSS] is a criterion for a topological exponential map to be Thurston equivalent to a (necessarily post-singularly finite) holomorphic exponential map. There exists an equivalent holomorphic map if and only if the following obstruction does not occur; see below.

Definition 6.5. (Essential curves and Levy cycles) Let g be a post-singularly finite topological exponential map. A simple closed curve $\gamma \subset \mathbb{S}^2 \setminus (P(g) \cup \{\infty\})$ is called *essential* if both connected components of $\mathbb{S}^2 \setminus \gamma$ contain at least two points of $P(g) \cup \{\infty\}$. A *Levy cycle* of g is a finite sequence of disjoint essential simple closed curves $\gamma_0, \gamma_1, \dots, \gamma_{m-1}, \gamma_m = \gamma_0$ such that for $i = 0, 1, \dots, m - 1$, some component γ'_i of $g^{-1}(\gamma_{i+1})$ is homotopic to γ_i in \mathbb{R}^2 relative to $P(g)$ and $g: \gamma'_i \rightarrow \gamma_{i+1}$ is a homeomorphism.

Let U'_i be the bounded component of $\mathbb{R}^2 \setminus \gamma'_i$ and let U_i be the bounded component of $\mathbb{R}^2 \setminus \gamma_i$. If all restrictions $g: \overline{U'_i} \rightarrow \overline{U_{i+1}}$ are homeomorphisms, then the Levy cycle is called *degenerate*.

It is easy to see that for topological exponential maps, every Levy cycle is degenerate and that Levy cycles are preserved under Thurston equivalence. Furthermore, a routine hyperbolic contraction argument shows that a post-singularly finite exponential map does not have a Levy cycle. The following theorem is the main result of [HSS].

THEOREM 6.6. (Topological characterization of exponential maps) *A post-singularly finite topological exponential map is Thurston equivalent to a post-singularly finite holomorphic exponential map if and only if it does not admit a degenerate Levy cycle. The holomorphic exponential map is unique up to conjugation with an affine map.*

Remark 6.7. In our parameterization $\lambda \mapsto \lambda \exp(z)$ of the exponential parameter space, with $\lambda \in \mathbb{C} \setminus \{0\}$, no two distinct E_λ and $E_{\tilde{\lambda}}$ are affinely conjugate. By Theorem 6.6, two post-singularly finite maps E_λ and $E_{\tilde{\lambda}}$ are never Thurston equivalent for different parameters $\lambda, \tilde{\lambda} \in \mathbb{C} \setminus \{0\}$.

We show that distinct psf exponential maps always yield abstract Hubbard trees that are not equivalent. For the proof, we need a result from [BFH, Corollary 6.6] about extensions of maps between embedded graphs.

LEMMA 6.8. (Extension of graph homeomorphisms) *Let $\Gamma_1, \Gamma_2 \subset \mathbb{S}^2$ be connected embedded graphs and let $f: \Gamma_1 \rightarrow \Gamma_2$ be a homeomorphism. Then f extends to an orientation-preserving homeomorphism $\hat{f}: \mathbb{S}^2 \rightarrow \mathbb{S}^2$ if and only if f preserves the cyclic order of the edges at all branch points of Γ_1 .*

THEOREM 6.9. (Different maps have non-equivalent trees) *Let E_λ and $E_{\tilde{\lambda}}$ be two post-singularly finite exponential maps. Let H be a homotopy Hubbard tree for E_λ and let \tilde{H} be a homotopy Hubbard tree for $E_{\tilde{\lambda}}$. If H and \tilde{H} yield the same abstract Hubbard tree, we have $\lambda = \tilde{\lambda}$.*

Proof. Choose dynamic rays g (respectively \tilde{g}) of E_λ (respectively $E_{\tilde{\lambda}}$) landing at the singular value. By Proposition 4.9, we can assume without loss of generality that H (respectively \tilde{H}) does not intersect $\text{Tr}(g)$ (respectively $\text{Tr}(\tilde{g})$). By hypothesis, there exists a

homeomorphism $\varphi: H \rightarrow \tilde{H}$ that restricts to a conjugation between E_λ and $E_{\tilde{\lambda}}$ on $P(E_\lambda)$, and preserves the cyclic order at each marked point of H . By Lemma 6.8, φ can be extended to a homeomorphism $\varphi_0: \mathbb{C} \rightarrow \mathbb{C}$ and we can choose φ_0 to satisfy $\varphi_0(\text{Tr}(g)) = \text{Tr}(\tilde{g})$. In analogy to Lemma 3.5, one shows that there exists a lift $\varphi_1: \mathbb{C} \rightarrow \mathbb{C}$ of φ_0 satisfying $\varphi_0 \circ E_\lambda = E_{\tilde{\lambda}} \circ \varphi_1$ and $\varphi_1(0) = 0$. We want to see that φ_0 is homotopic to $\varphi_1 \text{ rel } P(E_\lambda)$. This implies that E_λ and $E_{\tilde{\lambda}}$ are Thurston equivalent, so we have $\lambda = \tilde{\lambda}$ by Theorem 6.6 and the remark following it.

Let us first see that φ_1 restricts to a conjugation between E_λ and $E_{\tilde{\lambda}}$ on $P(E_\lambda)$, just as φ_0 does. Let \mathcal{D} be the dynamical partition for E_λ with respect to g , and let $\tilde{\mathcal{D}}$ be the dynamical partition for $E_{\tilde{\lambda}}$ with respect to \tilde{g} . It follows from $\varphi_0(\text{Tr}(g)) = \text{Tr}(\tilde{g})$ and $\varphi_1(0) = 0$ that the lifted map φ_1 maps the partition sector D_k homeomorphically onto the sector \tilde{D}_k of the same index. Hence, for all $n \geq 0$, φ_1 sends the unique pre-image of $E_\lambda^{on}(0)$ in the sector D_k to the unique pre-image of $E_{\tilde{\lambda}}^{on}(0)$ in \tilde{D}_k . As E_λ and $E_{\tilde{\lambda}}$ have the same kneading sequence (remember that the abstract Hubbard tree in particular contains the sector information for the post-singular points), it follows that $\varphi_1(E_\lambda^{o(n-1)}(0)) = E_{\tilde{\lambda}}^{o(n-1)}(0)$ for all $n > 0$.

Set $H' := [P(E_\lambda)]_{\widehat{E_\lambda}^{-1}(H)}$ and $\tilde{H}' := [P(E_{\tilde{\lambda}})]_{\widehat{E_{\tilde{\lambda}}}^{-1}(H)}$. The map φ_1 extends to a map $\widehat{\varphi}_1: \mathbb{C}_T \rightarrow \mathbb{C}_{\tilde{T}}$ by setting $\widehat{\varphi}_1(-\infty) := -\infty$. As $\varphi_1(P(E_\lambda)) = P(E_{\tilde{\lambda}})$, we have $\widehat{\varphi}_1(H') = \tilde{H}'$. By Lemma 3.14, there exists an embedded tree $H'' \subset \mathbb{C}$ such that H'' is homotopic to $H' \text{ rel } P(E_\lambda)$. Any homotopy between H' and $H'' \text{ rel } P(E_\lambda)$ is pushed forward by φ_1 to a homotopy between \tilde{H}' and an embedded tree $\tilde{H}'' := \varphi_1(H'') \subset \mathbb{C} \text{ rel } P(E_{\tilde{\lambda}})$. By Lemma 3.17, H'' is ambient isotopic to $H \text{ rel } P(E_\lambda)$ and \tilde{H}'' is ambient isotopic to $\tilde{H} \text{ rel } P(E_{\tilde{\lambda}})$. Stated differently, there exist homeomorphisms $\Psi_0, \Psi_1: \mathbb{C} \rightarrow \mathbb{C}$ such that Ψ_0 is isotopic to $\text{id rel } P(E_\lambda)$, Ψ_1 is isotopic to $\text{id rel } P(E_{\tilde{\lambda}})$, and $\varphi'_1 := \Psi_1 \circ \varphi_1 \circ \Psi_0$ satisfies $\varphi'_1(H) = \tilde{H}$. The composition φ'_1 is isotopic to $\varphi_1 \text{ rel } P(E_\lambda)$ and it is equal to φ_0 on the set of marked points of H . By [BFH, Corollary 6.3], there exists a homeomorphism $\Psi: \mathbb{C} \rightarrow \mathbb{C}$ isotopic to id relative to the set of marked points of H (and in particular relative to $P(E_\lambda)$) such that $\varphi'_1 \circ \Psi = \varphi_0$. However, $\varphi'_1 \circ \Psi$ is isotopic to $\varphi_1 \text{ rel } P(E_\lambda)$, so E_λ and $E_{\tilde{\lambda}}$ are Thurston equivalent. □

Next, we explain how to obtain a topological exponential map from an abstract exponential Hubbard tree. For convenience, we are going to embed abstract Hubbard trees into the complex plane \mathbb{C} (or suitable extensions of it). Note, however, that the complex structure of \mathbb{C} does not play any role; it only simplifies the construction. At first, we define a suitable embedding of the abstract tree into the complex plane. Afterwards, we show that the self-map of the embedded tree extends to a map on the plane and that this map is a topological exponential map. The constructions are inspired by and similar to those in [LSV, Ch. 5], so we skim some of them and refer to [LSV] for more details.

At first, we specify a general mapping layout for all of our topological exponential maps, which is independent of the abstract Hubbard tree that we want to realize. See Figure 4 for a sketch of the upcoming construction. Let $\gamma_h: (0, \infty) \rightarrow \mathbb{C}$, $\gamma_h(t) := t$ be a parameterization of the horizontal line of positive reals. For each $k \in \mathbb{Z}$, let $d_k: (-\infty, +\infty) \rightarrow \mathbb{C}$, and $d_k(t) := t + (2k - 1)\pi i$ be a parameterization of the straight horizontal line at constant imaginary part $(2k - 1)\pi$. We define a continuous map $g_0: \bigcup \text{Tr}(d_k) \rightarrow \text{Tr}(\gamma_h)$




FIGURE 4. Sketch of the embedded trees $H := \iota(H)$ and $H' := \iota(H)$. The two embedded trees agree on $\{z \in \mathbb{C} \mid \operatorname{Re}(z) \geq -1\}$.

by $g_0(d_k(t)) := \gamma_h(\exp(t))$. Denote by D_j the connected component of $\mathbb{C} \setminus \bigcup \operatorname{Tr}(d_k)$ bounded by $\operatorname{Tr}(d_j)$ and $\operatorname{Tr}(d_{j+1})$.

Let $(H, \mathfrak{f}, B_{\kappa_i}, \angle_{\mathfrak{v}_T})$ (we use a different font here to distinguish the abstract tree from the embedded tree) be an abstract Hubbard tree. The branches $B_{\kappa_0}, \dots, B_{\kappa_n}$ of H at \mathfrak{v}_T are indexed by distinct integers $\kappa_i \in \mathbb{Z}$, where $\kappa_0 = 0$ (compare Definition 4.2). Denote by $\mathfrak{v}_{\kappa_i}^{(1)}, \dots, \mathfrak{v}_{\kappa_i}^{(j_{\kappa_i})} \in V_{\mathfrak{f}} \cap B_{\kappa_i}$ the marked points of $(H, \mathfrak{f}, B_{\kappa_i}, \angle_{\mathfrak{v}_T})$ contained in the branch B_{κ_i} , where $\mathfrak{v}_{\kappa_i}^{(1)}$ is the unique marked point in B_{κ_i} incident to \mathfrak{v}_T . Choose auxiliary points $\mathfrak{v}_{\kappa_i}^{(0)} \in (\mathfrak{v}_T, \mathfrak{v}_{\kappa_i}^{(1)})$. We define an embedding $\iota: H \rightarrow \mathbb{C}$ in the following way.

- For $i \in \{0, \dots, n\}$, set $\iota(\mathfrak{v}_{\kappa_i}^{(0)}) := -1 + 2k_i\pi i =: v_{k_i}^{(0)}$.
- Choose an arbitrary point $v_T \in \{z \in \mathbb{C} \mid \operatorname{Re}(z) < -1\}$ and choose pairwise disjoint arcs $\gamma_i: [0, 1] \rightarrow \mathbb{C}$ connecting v_T to $v_{k_i}^{(0)}$ such that $\gamma_i((0, 1)) \subset \{z \in \mathbb{C} \mid \operatorname{Re}(z) < -1\}$. Let $\iota: [\mathfrak{v}_T, \mathfrak{v}_{\kappa_i}^{(0)}] \rightarrow \operatorname{Tr}(\gamma_i)$ to be a homeomorphism satisfying $\iota(\mathfrak{v}_T) = v_T$ and $\iota(\mathfrak{v}_{\kappa_i}^{(0)}) = v_{k_i}^{(0)}$.
- For $i \in \{1, \dots, n\}$, choose distinct points $v_{k_i}^{(1)}, \dots, v_{k_i}^{(j_{\kappa_i})} \in D_{k_i}$ satisfying $\operatorname{Re}(v_{k_i}^{(j)}) > -1$. Define $\iota: B_{\kappa_i} \setminus [\mathfrak{v}_T, \mathfrak{v}_{\kappa_i}^{(0)}] \rightarrow D_{k_i} \cap \{z \in \mathbb{C} \mid \operatorname{Re}(z) > -1\}$ to be an embedding satisfying $\iota(\mathfrak{v}_{\kappa_i}^{(j)}) = v_{k_i}^{(j)}$ such that the cyclic order of branches at $v_{k_i}^{(j)}$ coincides with the cyclic order of the corresponding branches of the abstract tree at $\mathfrak{v}_{\kappa_i}^{(j)}$. For the embedding of B_0 , we require $\iota(\mathfrak{f}(\mathfrak{v}_T)) = 0$ and $\iota(B_0) \cap \operatorname{Tr}(\gamma_h) = \emptyset$ in addition. Note that every oriented topological tree is planar, so there always exists such an embedding.

The image $H := \iota(H) \subset \mathbb{C}$ is an embedded tree. Next, we define a slightly different embedding ι' of H whose image will be the subset H' of the pre-image tree of H spanned by $P(g)$ under the yet to be constructed topological exponential map g .

Let $-\infty$ and $+\infty$ be abstract points not contained in \mathbb{C} and let $\mathbb{C}_\infty := \mathbb{C} \cup \{-\infty, +\infty\}$ denote the extension of the complex plane by these abstract points. We turn \mathbb{C}_∞ into a topological space by defining $U_n := \{z \in \mathbb{C} \mid \operatorname{Re}(z) < -n\} \cup \{-\infty\}$ to be a neighborhood

basis of $-\infty$ and $V_n := \{z \in \mathbb{C} \mid \operatorname{Re}(z) > n\} \cup \{+\infty\}$ to be a neighborhood basis of $+\infty$. For every $i \in \{0, \dots, n\}$, let $\tilde{\gamma}_i: [0, 1] \rightarrow \mathbb{C}_\infty$ be an arc connecting $-\infty$ to $v_{k_i}^{(0)}$ and satisfying $\gamma_i((0, 1)) \subset D_{k_i} \cap \{z \in \mathbb{C} \mid \operatorname{Re}(z) < -1\}$. Let $\iota'|_{[v_T, v_{k_i}^{(0)}]}: [v_T, v_{k_i}^{(0)}] \rightarrow \operatorname{Tr}(\tilde{\gamma}_i)$ be a homeomorphism, where $\iota'(v_T) = -\infty$ and $\iota'(v_{k_i}^{(0)}) = v_{k_i}^{(0)}$. On the complement of $[v_T, v_{k_i}^{(0)}]$, define ι' to be equal to ι and set $H' := \iota'(H) \subset \mathbb{C}_\infty$.

The next step is to define a map from H' to H . From now on, let $d_i: [-\infty, +\infty] \rightarrow \mathbb{C}_\infty$ denote the extended curve, where $d_i(-\infty) = -\infty$ and $d_i(+\infty) = +\infty$. We also extend γ_h to a curve $\gamma_h: [0, +\infty] \rightarrow \overline{\mathbb{C}}$, where $\gamma_h(0) = 0$ and $\gamma_h(+\infty) = \infty$. Then we define an extended map $g_0: \bigcup \operatorname{Tr}(d_i) \rightarrow \operatorname{Tr}(\gamma)$ by setting $g_0(-\infty) = 0$ and $g_0(+\infty) = \infty$. We define a graph map

$$\tilde{g}: H' \cup \left(\bigcup \operatorname{Tr}(d_i) \right) \rightarrow H \cup \operatorname{Tr}(\gamma_h), \quad \tilde{g}(z) := \begin{cases} g_0(z) & \text{if } z \in \bigcup \operatorname{Tr}(d_i), \\ \iota \circ \mathfrak{f} \circ (\iota')^{-1}(z) & \text{if } z \in H'. \end{cases}$$

Our next goal is to show that the graph map \tilde{g} can be extended to a post-singularly finite topological exponential map. Let us see how to use Lemma 6.8 in our setting. There exist continuous maps $\varphi_1: \mathbb{D} \rightarrow \operatorname{cl}_{\mathbb{C}_\infty}(D_i)$ and $\varphi_2: \mathbb{D} \rightarrow \operatorname{cl}_{\overline{\mathbb{C}}}(\mathbb{C} \setminus \operatorname{Tr}(\gamma_h))$ that restrict to orientation-preserving homeomorphisms from \mathbb{D} onto D_i and $\mathbb{C} \setminus \operatorname{Tr}(\gamma_h)$, respectively. The map φ_1 is a homeomorphism, while every interior point of $\operatorname{Tr}(\gamma_h)$ has two pre-images under φ_2 . We choose φ_1 to satisfy $\varphi_1(1) = +\infty$, $\varphi_1(-1) = -\infty$, and φ_2 to satisfy $\varphi_2(-1) = 0$, $\varphi_2(1) = \infty$. Then, there exist inverse branches $\psi_u: \operatorname{Tr}(\gamma_h) \rightarrow \partial\mathbb{D} \cap \{z \in \mathbb{C} \mid \operatorname{Im}(z) \geq 0\}$ and $\psi_l: \operatorname{Tr}(\gamma_h) \rightarrow \partial\mathbb{D} \cap \{z \in \mathbb{C} \mid \operatorname{Im}(z) \leq 0\}$ of φ_2 . We define a homeomorphism $f: \partial\mathbb{D} \rightarrow \partial\mathbb{D}$ by

$$f(z) := \begin{cases} \psi_u(g_0(\varphi_1(z))) & \text{if } \operatorname{Im}(z) \geq 0, \\ \psi_l(g_0(\varphi_1(z))) & \text{if } \operatorname{Im}(z) \leq 0. \end{cases}$$

We view $\mathbb{D} \subset \overline{\mathbb{C}}$ as a subset of the Riemann sphere. Setting $\Gamma_1 := \varphi_1^{-1}(\partial D_i \cup (H' \cap D_i))$, $\Gamma_2 := \varphi_2^{-1}(\tilde{g}(H' \cap D_i) \cup \operatorname{Tr}(\gamma_h))$, the $\Gamma_i \subset \overline{\mathbb{C}}$ are connected embedded graphs. We extend f to a homeomorphism $f: \Gamma_1 \rightarrow \Gamma_2$ by setting $f(z) = \varphi_2^{-1}(\tilde{g}(\varphi_1(z)))$ for $z \in \Gamma_1 \setminus \partial\mathbb{D}$. Then f preserves the circular order at branch points of Γ_1 because \tilde{g} preserves the circular order at branch points. By Lemma 6.8, there exists an orientation-preserving homeomorphism $\hat{f}: \overline{\mathbb{C}} \rightarrow \overline{\mathbb{C}}$ extending f and satisfying $\hat{f}(\mathbb{D}) = \mathbb{D}$. We define $g(z) := \varphi_2 \circ \hat{f} \circ \varphi_1^{-1}$ for $z \in D_i$.

Pasting these extensions together, we obtain a globally defined map $g: \mathbb{C} \rightarrow \mathbb{C} \setminus \{0\}$. As it is a covering map, it is a post-singularly finite topological exponential map. To see that g is Thurston equivalent to a holomorphic map, we have to show that g does not admit a Levy cycle. We can assume without loss of generality that γ_h is pre-periodic (as a set) under the dynamics of g and the forward images of γ_h are pairwise disjoint (except possibly for their landing points). By construction, the boundary of the partition $(D_j)_{j \in \mathbb{Z}}$ consists of the pre-images of γ_h , and by expansivity of the abstract Hubbard Tree, different post-singular points have different itineraries with respect to $(D_j)_{j \in \mathbb{Z}}$.

LEMMA 6.10. *The map g does not admit a Levy cycle.*

Proof. Assume to the contrary that g admits a Levy cycle $\Lambda = \{\delta_0, \delta_1, \dots, \delta_k = \delta_0\}$. Set $\gamma_0 := \gamma_h$ and denote by $\gamma_i := g^{oi}(\gamma_0)$ the forward iterates of γ_0 . By construction, γ_0 is pre-periodic as a set. For curves δ_j and γ_i , let

$$|[\delta_j] \cap [\gamma_i]| := \min_{\delta' \text{ isotopic to } \delta_j \text{ rel } P(g) \cup \{\infty\}} |\text{Tr}(\delta') \cap \text{Tr}(\gamma_i)|$$

denote the minimal intersection number of δ_j and γ_i . It follows from the periodicity of the curves γ_i and δ_j that $|[\delta_j] \cap [\gamma_0]| = 0$ for all $j \in \{0, \dots, k - 1\}$. See [BFH, Lemma 8.7] for a proof of this fact. Therefore, the curves δ_j of the Levy cycle do not intersect the partition boundary (up to homotopy), so all post-singular points surrounded by the same curve δ_j have equal itineraries. This contradicts the fact that different post-singular points have different itineraries with respect to $(D_j)_{j \in \mathbb{Z}}$. See [LSV, Lemma 5.5] for a detailed proof. □

By Theorem 6.6 and Lemma 6.10, the topological exponential map g is Thurston equivalent to a holomorphic exponential map E_{λ_g} . We are in the position to close the classification cycle.

THEOREM 6.11. (Classification of post-singularly finite exponential maps by abstract Hubbard trees) *The map*

$$\mathcal{F}: \{\text{psf exponential maps}\} \rightarrow \{\text{abstract exponential Hubbard trees}\}$$

is a bijection between post-singularly finite exponential maps (up to affine conjugation) and abstract exponential Hubbard trees (up to equivalence).

Proof. We have already shown that \mathcal{F} is well defined. We choose a representative for every equivalence class of abstract exponential Hubbard trees and perform the construction above to obtain a post-singularly finite exponential map associated with it. This sets up a map

$$\mathcal{G}: \{\text{abstract exponential Hubbard trees}\} \rightarrow \{\text{psf exponential maps}\}.$$

Let $(H, \mathfrak{f}, B_{\kappa_{\pm}}, \angle_{\nu})$ be the chosen representative for a given equivalence class of abstract Hubbard trees and let $E_{\lambda_g} = \mathcal{G}([(H, \mathfrak{f}, B_{\kappa_{\pm}}, \angle_{\nu})])$ be the corresponding holomorphic exponential map obtained by the construction above. We want to see that $\mathcal{F}(E_{\lambda_g}) = [(H, \mathfrak{f}, B_{\kappa_{\pm}}, \angle_{\nu})]$, that is, $\mathcal{F} \circ \mathcal{G} = \text{id}$. As we have already seen in Theorem 6.9 that \mathcal{F} is injective, this shows that \mathcal{F} is indeed a bijection.

By the definition of Thurston equivalence, there exist two homeomorphisms $\varphi_0, \varphi_1 : \mathbb{C} \rightarrow \mathbb{C}$ such that the diagram

$$\begin{CD} (\mathbb{C}, P(g)) @>\varphi_1>> (\mathbb{C}, P(E_{\lambda_g})) \\ @VgVV @VV E_{\lambda_g} V \\ (\mathbb{C}, P(g)) @>\varphi_0>> (\mathbb{C}, P(E_{\lambda_g})) \end{CD}$$

commutes and φ_0 and φ_1 are homotopic rel $P(g)$. We want to see that $H_0 := \varphi_0(H)$ is a homotopy Hubbard tree for E_{λ_g} and that it yields the abstract exponential Hubbard tree with which we have started. Obviously, H_0 is a finite embedded tree spanned by $P(E_{\lambda_g})$.

We want to see that H_0 is invariant up to homotopy rel $P(E_{\lambda_g})$. Set $H_1 := \varphi_1(H)$. By Proposition 3.12, φ_0 and φ_1 are isotopic rel $P(g)$, and hence the embedded trees H_0 and H_1 are ambient isotopic. More precisely, the map $\varphi := \varphi_1 \circ \varphi_0^{-1} : \mathbb{C} \rightarrow \mathbb{C}$ is isotopic to id rel $P(E_{\lambda_g})$ such that $\varphi(H_0) = H_1$. We want to see that $H'_0 = [P(E_{\lambda_g})]_{\widehat{E_{\lambda_g}}^{-1}(H_0)}$ is homotopic to H_0 in \mathbb{C}_T rel $P(E_{\lambda_g})$.

By construction, there exists a homotopy $I : H \times [0, 1] \rightarrow \mathbb{C} \cup \{-\infty\}$ between H and H' rel $P(g)$ satisfying $I^1 = \iota' \circ \iota^{-1}$. The homeomorphism φ_1 extends to a homeomorphism $\widehat{\varphi}_1 : \mathbb{C} \cup \{-\infty\} \rightarrow \mathbb{C}_T$ by setting $\widehat{\varphi}_1(-\infty) := -\infty$. By commutativity of the Thurston diagram, we have $H'_0 = \widehat{\varphi}_1(H')$. The homotopy I is pushed forward through $\widehat{\varphi}_1$ to a homotopy $\tilde{I} := \widehat{\varphi}_1 \circ I$ between H_1 and H'_0 rel $P(E_{\lambda_g})$. We have shown above that H_1 is homotopic to H_0 rel $P(E_{\lambda_g})$, so H_0 is homotopic to H'_0 rel $P(E_{\lambda_g})$.

The identification of H_0 with H'_0 yielded by the constructed homotopy is given by $\psi := \widehat{\varphi}_1 \circ \iota' \circ \iota^{-1} \circ \varphi_0^{-1} : H_0 \rightarrow H'_0$. It remains to show that the induced self-map $f := E_{\lambda_g} \circ \tilde{\psi} : H_0 \rightarrow H_0$ is expansive. Looking carefully at the definition of the involved maps, we see that $f = (\varphi_0 \circ \iota) \circ \mathfrak{f} \circ (\varphi_0 \circ \iota)^{-1}$, that is, the self-map of H_0 is conjugate to the self-map \mathfrak{f} of the abstract Hubbard tree H . Hence, f is expansive by the expansivity of \mathfrak{f} .

This also shows that the abstract Hubbard trees $(H_0, f, B_{k_i}, \angle_v)$ and $(H, \mathfrak{f}, B_{k_i}, \angle_v)$ are equivalent through $\varphi_0 \circ \iota$ except for the fact that the sector information of H and H_0 is consistent. For a post-singular point $\mathfrak{p} \in B_{k_i}$, we have $\iota(\mathfrak{p}) \in D_{k_i}$, where $(D_i)_{i \in \mathbb{Z}}$ is the partition for g defined above. As $\text{Tr}(\gamma_h) \cap H = \{0\}$, we have $H_0 \cap \text{Tr}(\gamma) = \{0\}$, where $\gamma := \varphi_0 \circ \gamma_h$. By Proposition 4.9 and Lemma 3.17, the arc γ is ambient isotopic rel $P(E_{\lambda_g})$ to a dynamic ray $g_{\underline{s}}$ landing at the singular value. Let $(\tilde{D}_i)_{i \in \mathbb{Z}}$ be the dynamical partition for E_{λ_g} with respect to $g_{\underline{s}}$. As φ_0 restricts to a conjugation between $P(g)$ and $P(E_{\lambda_g})$, and as the ambient isotopy between γ and $g_{\underline{s}}$ does not move post-singular points, we have $\varphi_0 \circ \iota(\mathfrak{p}) \in \tilde{D}_{k_i}$. Hence, the sector information of H and H_0 is indeed consistent. \square

This concludes the classification of psf exponential maps in terms of abstract exponential Hubbard trees.

7. Outlook

The key tools for our construction of homotopy Hubbard trees are dynamic rays, their landing properties, and the combinatorial encoding of both. It is natural to ask whether more general maps also have homotopy Hubbard trees. One putative obstacle is that the escaping set of an arbitrary post-singularly finite entire function need not consist of dynamic rays. Recently, the concept of *dreadlocks* has been introduced in [BR] as a generalization of dynamic rays. For every post-singularly finite entire function, the escaping set naturally decomposes into dreadlocks and the collection of dreadlocks can be encoded in terms of external addresses. Every repelling periodic point is the landing point of at least one and at most finitely many dreadlocks, all of which are periodic of the same period. In short, dreadlocks possess all the important properties needed to study the branching of the Julia set in terms of the escaping set. Using these findings, the existence and uniqueness of homotopy Hubbard trees for all post-singularly finite transcendental entire functions has been established based on a study of branch points of the Julia set in

[Pf]. These results will be published in a series of forthcoming articles, the first of which is [PPS].

Another interesting question that has not been answered in [Pf] is which post-singularly finite transcendental entire function has a Hubbard tree in the classical sense, that is, a compact embedded tree that is forward invariant as a set, not just forward invariant up to homotopy. By Theorem 3.4, psf exponential maps do not have a Hubbard tree essentially because of the existence of an asymptotic value. We expect this to be the only obstacle to the existence of Hubbard trees.

Acknowledgements. We gratefully acknowledge that this project was partially supported by a grant from the Deutsche Forschungsgemeinschaft, as well as by the Advanced Grant HOLOGRAM of the European Research Council. Major parts of this work were carried out at Jacobs University Bremen; other parts happened during a visit to Cornell University in the spring of 2018, for which we are grateful as well. We would like to thank Mikhail Hlushchanka, Russell Lodge, Dzmitry Dudko, Bernhard Reinke, John Hubbard, and Lasse Rempe for helpful discussions and comments. We would like to thank the anonymous referee for some helpful comments regarding the presentation.

REFERENCES

- [B] H. Mihaljević-Brandt. Topological dynamics of transcendental entire functions. *PhD Thesis*, University of Liverpool, 2009.
- [BFH] B. Bielefeld, Y. Fisher and J. Hubbard. The classification of critically preperiodic polynomials as dynamical systems. *J. Amer. Math. Soc.* **5**(4) (1992), 721–762.
- [BFP] A. M. Benini, J. E. Fornæss and H. Peters. Entropy of transcendental entire functions. *Ergod. Th. & Dynam. Sys.* **41**(2) (2021), 338–348.
- [BKS] H. Bruin, A. Kaffl and D. Schleicher. Existence of quadratic Hubbard trees. *Fund. Math.* **202**(3) (2009), 251–279.
- [BR] A. M. Benini and L. Rempe. A landing theorem for entire functions with bounded post-singular sets. *Geom. Funct. Anal.* **30**(6) (2020), 1465–1530.
- [DGH] C. Bodelón, R. L. Devaney, M. Hayes, G. Roberts, L. R. Goldberg and J. H. Hubbard. Dynamical convergence of polynomials to the exponential. *J. Difference Equ. Appl.* **6**(3) (2000), 275–307.
- [DH1] A. Douady and J. H. Hubbard. *Etude Dynamique des Polynômes Complexes (The Orsay Notes) (Publications Mathématiques d’Orsay, 84-02/85-04)*. Département de Mathématique, Université de Paris-Sud, Orsay, France, 1984/1985.
- [DH2] A. Douady and J. H. Hubbard. A proof of Thurston’s topological characterization of rational functions. *Acta Math.* **171**(2) (1993), 263–297.
- [DS] D. Dudko and D. Schleicher. Core entropy of quadratic polynomials. *Arnold Math. J.* **6**(3–4) (2020), 333–385, with an appendix by Wolf Jung.
- [E1] A. Eremenko. On the iteration of entire functions. *Dynamical Systems and Ergodic Theory (Warsaw, 1986) (Banach Center Publications, 23)*. PWN, Warsaw, 1989, pp. 339–345.
- [E2] A. Eremenko. Singularities of inverse functions. *Presented at the ICMS Conference ‘The Role of Complex Analysis in Complex Dynamics’, 2013*.
- [FM] B. Farb and D. Margalit. *A Primer on Mapping Class Groups (Princeton Mathematical Series, 49)*. Princeton University Press, Princeton, NJ, 2012.
- [FS] M. Förster and D. Schleicher. Parameter rays in the space of exponential maps. *Ergod. Th. & Dynam. Sys.* **29**(2) (2009), 515–544.
- [GK] L. R. Goldberg and L. Keen. A finiteness theorem for a dynamical class of entire functions. *Ergod. Th. & Dynam. Sys.* **6**(2) (1986), 183–192.
- [Ha] M. Hassler. Core entropy of unicritical maps. *Bachelor’s Thesis*, Jacobs University Bremen, 2020.
- [HSS] J. Hubbard, D. Schleicher and M. Shishikura. Exponential Thurston maps and limits of quadratic differentials. *J. Amer. Math. Soc.* **22**(1) (2009), 77–117.

- [LMS1] R. Lodge, Y. Mikulich and D. Schleicher. Combinatorial properties of Newton maps. *Indiana Univ. Math. J.*, to appear. Preprint [arXiv:1510.02761](https://arxiv.org/abs/1510.02761).
- [LMS2] R. Lodge, Y. Mikulich and D. Schleicher. A classification of postcritically finite Newton maps. Preprint, 2015, [arXiv:1510.02771](https://arxiv.org/abs/1510.02771).
- [LSV] B. Laubner, D. Schleicher and V. Vicol. A combinatorial classification of postsingularly finite complex exponential maps. *Discrete Contin. Dyn. Syst.* **22**(3) (2008), 663–682.
- [M] J. Milnor. *Dynamics in One Complex Variable (Annals of Mathematics Studies, 160)*, 3rd edn. Princeton University Press, Princeton, NJ, 2006.
- [P] A. Poirier. Hubbard trees. *Fund. Math.* **208**(3) (2010), 193–248.
- [Pf] D. Pfrang. Homotopy Hubbard trees for post-singularly finite transcendental entire maps. *PhD Thesis*, IRC-Library, Information Resource Center Jacobs University Bremen, 2019. <http://nbn-resolving.org/urn:nbn:de:gbv:579-opus-1008936>.
- [PPS] D. Pfrang, S. Petrat and D. Schleicher. Dreadlock pairs and dynamic partitions for post-singularly finite entire functions. *Manuscript*, in preparation, July 2021.
- [R] M. Rothgang. Hubbard trees for post-singularly finite exponential maps. *Bachelor's Thesis*, Jacobs University Bremen, 2016.
- [RRRS] G. Rottenfusser, J. Rückert, L. Rempe and D. Schleicher. Dynamic rays of bounded-type entire functions. *Ann. of Math. (2)* **173**(1) (2011), 77–125.
- [S1] D. Schleicher. On fibers and local connectivity of Mandelbrot and Multibrot sets. *Fractal Geometry and Applications: A Jubilee of Benoît Mandelbrot. Part 1 (Proceedings of Symposia in Pure Mathematics, 72)*. American Mathematical Society, Providence, RI, 2004, pp. 477–517.
- [S2] D. Schleicher. Dynamics of entire functions. *Holomorphic Dynamical Systems (Lecture Notes in Mathematics, 1998)*. Springer, Berlin, 2010, pp. 295–339.
- [SZ1] D. Schleicher and J. Zimmer. Escaping points of exponential maps. *J. Lond. Math. Soc. (2)* **67**(2) (2003), 380–400.
- [SZ2] D. Schleicher and J. Zimmer. Periodic points and dynamic rays of exponential maps. *Ann. Acad. Sci. Fenn. Math.* **28**(2) (2003), 327–354.
- [T] W. P. Thurston. Entropy in dimension one. *Frontiers in Complex Dynamics: In Celebration of John Milnor's 80th Birthday*. Eds. A. Bonifant, M. Lyubich and S. Sutherland. Princeton University Press, Princeton, NJ, 2014, pp. 339–384.
- [Ti] G. Tiozzo. Continuity of core entropy of quadratic polynomials. *Invent. Math.* **203**(3) (2016), 891–921.