EXISTENCE OF QUADRATIC HUBBARD TREES

HENK BRUIN, ALEXANDRA KAFFL, AND DIERK SCHLEICHER

ABSTRACT. A (quadratic) Hubbard tree is an invariant tree connecting the critical orbit within the Julia set of a postcritically finite (quadratic) polynomial. It is easy to read of the kneading sequences from a quadratic Hubbard tree; the result in this paper handles the converse direction.

Not every sequence on two symbols is realized as the kneading sequence of a real or complex quadratic polynomial. Milnor and Thurston classified all real-admissible sequences, and we give a classification of all complex-admissible sequences in [BS]. In order to do this, we show here that every periodic or preperiodic sequence is realized by a unique abstract Hubbard tree. Real or complex admissibility of the sequence depends on whether this abstract tree can be embedded into the real line or complex plane so that the dynamics preserves the embedded, and this can be studied in terms of branch points of the abstract Hubbard tree.

1. INTRODUCTION

Many properties of complex dynamics, both of the dynamical plane and the parameter space, can be described conveniently with symbolic dynamics, see [D1, M] for general introductions. In fact, one of the chief reasons why polynomial dynamics is far better understood that rational dynamics is because we have much better tools available for symbolic dynamics, see [D1, M] for general introductions. Hubbard trees, external angles and kneading theory, together with their recodings as internal addresses [LS, S1], are fundamental concepts within this symbolic description.

- A Hubbard tree is a sketch of the essential topological features of a polynomial Julia set from which many other topological or combinatorial data can be reconstructed quite easily. They were introduced by Douady and Hubbard in their Orsay Notes [DH1] in order to describe postcritically finite polynomials; this program was later carried through in [P, BFH] using a fundamental theorem of Thurston [DH2].

- External angles and external parameter angles help to organize Julia sets and the space of quadratic polynomials respectively, see [S4].
  - In the dynamical space of a fixed polynomial $f : z \mapsto z^2 + c_1$ (with c_1 in the Mandelbrot set $\mathcal{M}$), the angle doubling map on the circle $\mathbb{S}^1 = \mathbb{R}/\mathbb{Z}$ of external angles factorizes over the dynamics on the Julia set of f by identifying angles whose external rays land at the same point.
  - Similarly, every angle in $\mathbb{S}^1$ occurs as external parameter angle for some quadratic polynomial, and it is usually easy to tell if different angles correspond to the same polynomial (*i.e.,* if their external parameter rays land at the same $c \in \partial\mathcal{M}$. [1])

  Thus dynamical and parameter external angles give rise to topological models "from the outside" of Julia sets and the Mandelbrot set respectively, see [S3, T].
- The dynamics of a dendrite quadratic Julia set (or a Hubbard tree) can be described by itineraries using symbols 0 and 1 for the parts on either side of the critical point, and $\star$ for the critical point itself. The itinerary of the critical value c_1 is called the *kneading sequence*. Kneading sequences, especially in their human-readable form of internal addresses [LS, S1], allow to construct a model of the Mandelbrot set "from the inside". They describe which polynomials are topologically conjugate, and which external angles correspond to rays that land together.

There are many algorithms known to turn one of these three concepts into another. From a Hubbard tree one can directly read off the associated kneading sequence and the external angle(s), provided the Hubbard tree comes (as is traditionally the case) with an embedding into the plane, see [D2, BKS] and Example 2.7. Similarly, every external angle easily specifies the associated kneading sequence. It is much harder to find a Hubbard tree or an external angle from the kneading sequence. This is what we do in the present paper.

Part of the difficulty of finding a Hubbard tree for a given kneading sequence comes from the fact that some kneading sequences are not realized by complex polynomials and thus by Hubbard trees embedded in the plane. Other kneading sequences are realized by several such trees. It turns out that both difficulties disappear when turning attention to abstract (non-embedded) Hubbard trees. Our main result is the following, see Theorem 2.5:

---

[1]We are only involved with rational angles, so their external parameter rays land.

> *Every periodic or preperiodic kneading sequence is realized by*
> *a unique abstract Hubbard tree.*

This is a fundamental and long overdue step in a coherent description of symbolic dynamics of quadratic polynomials. We will specify the precise meanings of periodic kneading sequence, abstract Hubbard tree and uniqueness in Definitions 2.4, and 2.2 and the discussions following them. All our Hubbard trees are abstract non-embedded trees (even though we will usually omit the article "abstract" from now on); this distinguishes them from traditional Hubbard trees that come with an embedding into the plane.

Figure 1 gives an example of an abstract Hubbard tree that is not associated to any quadratic polynomial. In [BS], we discuss branch points of these abstract Hubbard trees and show how their properties are determined explicitly in terms of the kneading sequence alone. (Branch points in parameter space, *i.e.,* the Mandelbrot set, are closely related to branch points in Hubbard trees; see [S2].) This allows us to specify an explicit *admissibility condition* determining which kneading sequences are realized by complex quadratic polynomials (the corresponding question for real polynomials has been answered by Milnor and Thurston [MT] in the late 1970s).
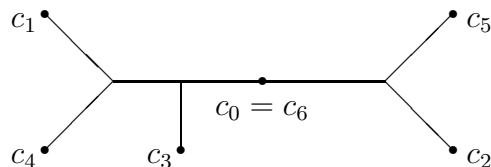

FIGURE 1. The abstract Hubbard tree for $\nu = \overline{10110\star}$. The three branch points form a period 3 orbit, but the third iterate permutes the arms in a non-cyclic way. As a result, this tree is not the Hubbard tree of any quadratic Julia set.

Further extensions of the theory (as covered in [BKS]) include (i) explicit algorithms to convert kneading sequences into Hubbard trees and vice versa, (ii) further algorithmic relations between various concepts of symbolic dynamics, (iii) the construction of abstract Julia sets as limit of backward iterates of the Hubbard tree, (iv) the construction of Hubbard dendrites, when ν is non-(pre)periodic, as projective limit of a sequence of Hubbard trees, (v) the construction of abstract Mandelbrot set as bifurcation diagram of the space of Hubbard trees, (vi) an analysis of biaccessible points both in Julia sets and the Mandelbrot sets, and the corresponding "biaccessible" external angles, and (vii) the proof that the set of complex-admissible kneading sequences has positive measure in the space of all sequences over two symbols.

The paper is organized as follows. In Section 2 we give the definitions of abstract Hubbard trees and symbolic dynamics, and we give a formal statement of our Main Theorem 2.5. We also give basic properties of a Hubbard tree including the existence of the α-fixed point. In Section 3, we introduce the triod argument which leads easily to the uniqueness of the abstract Hubbard tree, up to some natural equivalence relation. In Section 4 we define the ρ-function, which gives precise information about the relative positions of precritical points in the Hubbard tree. With regard to the proof of the main theorem, it enables us to show, purely symbolically, that c_1 is indeed an endpoint of T. This is easy to see once T is known to exist, see Lemma 2.3, but as a priori knowledge, this is a crucial step in the construction of T. In Section 5, we finally present the full construction of the tree. We first select its marked points by (i) taking shifts of ν and (ii) applying triod operations on the former. Gradually the neighbor-relations of marked points, and the dynamic properties of the tree will be verified.

2. HUBBARD TREES

In this section, we define Hubbard trees as abstract trees with dynamics and show their most fundamental properties. Our trees do not necessarily come with an embedding into the complex plane.

2.1. Definition (Trees, Arms, Branch Points and Endpoints)
A tree *T is a finite connected graph without loops. For a point $x \in T$, the (global) arms of x are the connected components of $T \setminus \{x\}$. A local arm at x is an intersection of a global arm with a neighborhood of x in T (where we use arc-length along the tree as metric). The point x is an* endpoint *of T if it has only one arm; it is a* branch point *if it has at least three arms.*

Between any two points x, y in a tree, there exists a unique closed arc connecting x and y; we denote it by $[x, y]$ and its interior by (x, y).

2.2. Definition (The Hubbard Tree)
A Hubbard tree is a tree T equipped with a map $f : T \to T$ and a distinguished point, the critical point, *satisfying the following conditions:*

(1) *$f : T \to T$ is continuous and surjective;*
(2) *every point in T has at most two inverse images under f;*
(3) *at every point other than the critical point, the map f is a local homeomorphism onto its image;*
(4) *all endpoints of T are on the critical orbit;*
(5) *the critical point is periodic or preperiodic, but not fixed;*

(6) *(expansivity) if x and y with $x \neq y$ are branch points or points on the critical orbit, then there is an $n \geq 0$ such that $f^{\circ n}([x, y])$ contains the critical point.*

We denote the critical point by $c_0 = 0$ and its orbit by $\mathrm{orb}_f(0) = \{0, c_1, c_2, \dots\}$. The *critical value c_1* is the image of the critical point. We use a standing assumption that $c_1 \neq 0$ in order to avoid having to deal with counterexamples when the entire tree is a single point. The branch points and the points on the critical orbit (starting with c_0) will be called *marked points*. Not all marked point need necessarily be endpoints or branchpoints (some post-critical point c_k could lie on the interior of arcs), but the set of marked points is finite and forward invariant, because the number of arms at any point can decrease under f only at the critical point.

Two Hubbard trees (T, f) and (T', f') are *equivalent* if there is a bijection between their marked points which is respected by the dynamics, and if the edges of the tree connect the same marked points. This is weaker than a topological conjugation. In particular, we do not care about details of the dynamics between marked points; there may be intervals of periodic points, attracting periodic points, and so on. (This is related to an equivalence class of branched covers in the sense of Thurston as in [DH2, HS].)

2.3. Lemma (Basic Properties of the Hubbard Tree)

The critical value c_1 is an endpoint, and the critical point 0 divides the tree into at most two parts. Each branch point is periodic or preperiodic, it never maps onto the critical point, and the number of arms is constant along the periodic part of its orbit. Any arc which does not contain the critical point in its interior maps homeomorphically onto its image.

PROOF. Suppose that c_1 has at least two arms. The points c_2, c_3, ... also have at least two arms as long as f is a local homeomorphism near this orbit. If this is no longer the case at some point, then the orbit has reached the critical point, and the next image is c_1 again. In any case, all points on the critical orbit have at least two arms. This contradicts the assumption that all endpoints of a Hubbard tree are on the critical orbit. Hence c_1 has exactly one arm, and 0 has at most two arms (or its image would not be an endpoint).

Since near every non-critical point, the dynamics is a local homeomorphism onto the image, every branch point maps onto a branch point with at least as many arms. Since the critical point has at most two arms, it can never be the image of a branch point. The tree and thus the number of branch points is finite, so every branch point is preperiodic or periodic and

its entire orbit consists of branch points; the number of arms is constant along the periodic part of the orbit.

Let γ be an arc within the tree. Since f cannot be constant on γ and there is no loop in the tree, the subtree $f(\gamma)$ has at least two endpoints. If an endpoint of $f(\gamma)$ is not the image of an endpoint of γ, then it must be the image of the critical point since f is a local homeomorphism elsewhere, and the critical point 0 must be in the interior of γ.   □

We have seen that $T \setminus \{0\}$ consists of at most two components. Let us denote them by T_0 and T_1 so that $c_1 \in T_1$ (with $c_1 \neq 0$ by definition); T_0 may be empty.

2.4. Definition (Itinerary and Kneading Sequence)
The itinerary *of a point $z \in T$ on a Hubbard tree is the infinite sequence* $e(z) = e_1(z)e_2(z)e_3(z)\ldots$ *with*

$$e_i(z) = \begin{cases} 0 & \text{if } f^{\circ(i-1)}(z) \in T_0, \\ \star & \text{if } f^{\circ(i-1)}(z) = 0, \\ 1 & \text{if } f^{\circ(i-1)}(z) \in T_1. \end{cases}$$

The itinerary $e(c_1) =: \nu = \nu_1\nu_2\nu_3\ldots$ of c_1 is called the kneading sequence.

If $e_i(z) = \star$, then $e_{i+j}(z) = \nu_j$ for every $j \geq 1$. In particular, if there is n minimal such that $\nu_n = \star$, then $\nu = \overline{\nu_1\ldots\nu_{n-1}\star}$. Such kneading sequences are called $\star$-*periodic*.

We can now state our main theorem precisely.

2.5. Theorem (Existence and Uniqueness of Hubbard Trees)
Every $\star$-periodic or preperiodic kneading sequence is realized by a unique (up to equivalence) abstract Hubbard tree.

It is sometimes useful to extend a Hubbard tree to include finite set of (pre)periodic orbits. Such an *extended Hubbard tree* $(\hat{T}, \hat{f})$ satisfies Definition 2.2, except that $\hat{f} : \hat{T} \to \hat{T}$ need not be surjective and part (4) must be replaced by: all endpoints in $\hat{T}$ are (pre)periodic. Branch points and points on the critical orbit are still marked, but also the additional points with itineraries in $\hat{V}$ should now be marked.

2.6. Corollary (Existence and Uniqueness of Extended Hubbard Trees)
Given a $\star$-periodic or preperiodic kneading sequence ν and a set $\hat{V} \subset \{0, 1\}^{\mathbb{N}^*} \setminus \{0\nu, 1\nu\}$ such that $\hat{V} \cup \{\star\nu, \nu, \sigma(\nu), \sigma^{\circ 2}(\nu), \ldots\}$ is shift-invariant, there is an extended Hubbard tree $(\hat{T}, \hat{f})$ such that

- *for every $v \in \hat{V}$, there is $p \in \hat{T}$ with itinerary $e(p) = v$;*
- *$\hat{T}$ is minimal in the sense that every endpoint belongs to the critical orbit, or is a point with itinerary in $\hat{V}$.*

- $\hat{T}$ is unique up to equivalence.

2.7. Example (Extending Hubbard trees to compute external angles)

By extending a Hubbard tree with the β-fixed point and its preimage $-\beta$ (with itineraries $e(\beta) = \overline{0}$ and $e(-\beta) = 1\overline{0}$ respectively), one can reconstruct the external angle of the ray landing at some point in $x \in T$. The n-th coordinate in the binary expansion of the external angle of x is 0 if $f^{\circ n}(x)$ is above the *spine* $[\beta, -\beta]$ and 1 if $f^{\circ n}(x)$ is below the spine. In Figure 2 we illustrate how this works for the critical value of the Hubbard tree with kneading sequence $\nu = 110\overline{1}$. Since several points on the orbit of
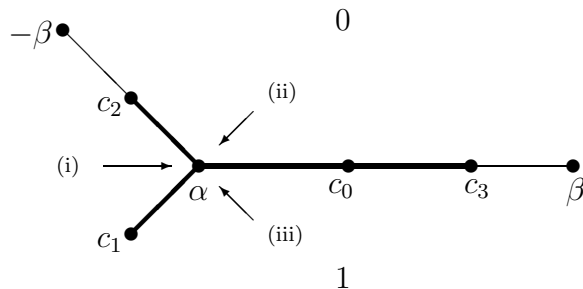

FIGURE 2. A Hubbard tree for $\nu = 110\overline{1}$ (in bold lines) and is extension. The third image of the critical value is the fixed point α which in this case has three external approaches, so c_1 has three external angles.

the critical value belong to the spine, c_1 has multiple (here three) external angles corresponding to its different external approaches. These angles are

$$(i):\quad 0.1101011011011011011\dots \text{ (binary) } = \frac{47}{56}$$

$$(ii):\quad 0.1100110110110110110\dots \text{ (binary) } = \frac{45}{56}$$

$$(iii):\quad 0.1011101101101101101\dots \text{ (binary) } = \frac{41}{56}$$

Note that this algorithm depends on the fact that T can be properly embedded into the plane. In fact, for $\nu = 110\overline{1}$, there is also another embedding, the mirror image of Figure 2, and the corresponding to external angles for c_1 are $\frac{15}{56}$, $\frac{11}{56}$ and $\frac{9}{56}$.

Throughout, we will write $\mathbb{N}^* = \{1, 2, 3, \dots\}$ and use the following notation for our symbol spaces:

$$
\begin{aligned}
\Sigma^1 &:= \{\nu \in \{0, 1\}^{\mathbb{N}^*} : \text{the first entry in } \nu \text{ is } 1\}, \\
\Sigma^\star &:= \Sigma^1 \cup \{\text{all } \star\text{-periodic sequences}\}, \\
\Sigma^{\star\star} &:= \{\nu \in \Sigma^1 : \nu \text{ is non-periodic}\} \cup \{\text{all } \star\text{-periodic sequences}\}.
\end{aligned}
$$

Obviously $e \circ f(z) = \sigma \circ e(z)$ where σ denotes the left shift. The expansivity condition of Definition 2.2 means that no two marked points have the same itinerary.

2.8. Lemma (Same Itinerary on Connected Subtree)

Suppose that z and z' are two points on a Hubbard tree such that $e_i(z) = e_i(z')$ for all $i < n$ (for some $n \leq \infty$). Then all $w \in [z, z']$ have $e_i(w) = e_i(z) = e_i(z')$ for $i < n$.

PROOF. We can assume that $z \neq z'$. Since 0 is the only point whose itinerary starts with $\star$, z and z' lie in the same component of $T \setminus \{0\}$. Therefore $[z, z']$ is mapped homeomorphically onto $f([z, z'])$ by f. Since $e_2(z) = e_2(z')$, the arc $f([z, z'])$ is contained in a single component of $T \setminus \{0\}$. By induction, we obtain that $f^{\circ i}([z, z'])$ is contained in a single component of $T \setminus \{0\}$ for each $i < n$. The claim follows. $\square$

2.9. Lemma (α-Fixed Point on Hubbard Tree)

There is a unique fixed point in T_1; it lies in $(0, c_1)$.

PROOF. Since c_1 is an endpoint, the intersection $[0, c_1] \cap f([0, c_1])$ is a non-degenerate arc $[c_1, x]$, i.e., $x \neq c_1$. If $x = 0$, then f maps $[0, c_1]$ over itself in an orientation reversing way, so there is a fixed point on $[0, c_1]$.

We may thus assume that $x \in (0, c_1)$. If $f([0, c_1]) \subset [0, c_1]$, then as above we have a fixed point in (c_1, x). Otherwise $f([0, c_1])$ branches off from $[0, c_1]$ at x. Let $y = f(x)$. Then y cannot be on $(0, x)$ because the path $f([0, c_1])$ starts at c_1 and branches off at x before reaching $(0, x)$. If $y \in [x, c_1]$, then f maps $[x, c_1]$ over $[y, c_2]$, and f has a fixed point in $[x, y]$.

The last possibility is that $y \in (x, c_2]$, and we show this does not occur. Let T' be the connected component of $T \setminus \{x\}$ containing y. The image $f(T') \not\ni x$, because one of the two inverse images of x is on $[0, x]$, and the other is separated from x by the critical point. Since x maps into T' and no point in T' maps onto x, all of T' maps strictly into itself under f. But this violates the expansivity condition: T' has an endpoint x' other than x, and the forward orbits of x and the endpoint are never separated by 0.

Now we have a fixed point in T_1; call it α. Suppose that it is not unique. Since f maps T_1 homeomorphically onto its image, f must fix a component

G of $T_1 \setminus \{\alpha\}$. This is not the component with 0 as boundary point, because α separates 0 from $c_1 = f(0)$. Let z be an endpoint of this fixed component G. Then α and z both are marked points with the same itinerary $\overline{1}$, and this contradicts the expansivity condition. $\square$

REMARK. The unique fixed point in T_1 is usually called α. The component T_0 can contain more fixed points, but by Lemma 2.8, they are all contained in a connected subtree of constant itinerary $\overline{0}$. If there is an endpoint with itinerary $\overline{0}$, it is called β; it exists on the Hubbard tree if and only if the kneading sequence terminates in an infinite string of symbols 0.

   A point $z \in T$ is *(pre)periodic* if $f^{ol}(z) = f^{o(l+m)}(z)$ for some $l \geq 0, m \geq 1$. We take l and m minimal with this property. Then m is the (exact) *period* of z and l the *preperiod*.

2.10. Lemma (Preperiod and Period)

The exact preperiod and period of any marked (pre)periodic point are equal to the exact preperiod and period of its itinerary.

PROOF. Suppose z is periodic with period m and let m' be the period of $e(z)$ (under the shift). Obviously m' divides m. If $m' \neq m$, then z and $f^{om'}(z)$ are different marked points with the same itinerary. This contradicts expansivity. The same argument works in the preperiodic case. $\square$

2.11. Lemma (Periodic Points and Itineraries)

If a Hubbard tree contains a point with periodic itinerary τ, then it contains a periodic point p with itinerary τ such that the exact periods of p and τ coincide.

PROOF. Let $T' \subset T$ be the set of all points with itinerary τ. By Lemma 2.8, T' is connected, so it is a connected subtree (possibly not closed).

   Let n be the period of τ. Then f^{on} maps T' homeomorphically onto its image in T'. Since marked points in T have different itineraries, T' can contain at most one branch point of T. If it contains one, then it must be fixed by f^{on}, so its exact period is n. Otherwise, T' is either a single point (and we are done), or T' is homeomorphic to an interval. If f^{on} sends T' to itself reversing the orientation, we get a unique fixed point in the interior of T', and we are done again.

   Now suppose that f^{on} preserves the orientation of T'. If $f^{on} \colon T' \to T'$ is not surjective, then for at least one endpoint x, say, $f^{on}(x)$ is in the interior of T'. If x is a branch point or an endpoint of T, then it is marked and we are done. Otherwise, x has a neighborhood in T which is homeomorphic to an open interval, but only a one-sided neighborhood has itinerary τ.

This implies that x maps to the critical point after finitely many iterations. Again T' contains a marked point, which must be fixed by $f^{\circ n}$.

The last case is that $f^{\circ n}$ maps T' homeomorphically onto itself, preserving the orientation and fixing both endpoints. Then the claim is satisfied by any endpoint of T' which does not hit the critical point on its forward orbit. If both endpoints do, say after k and m iterations with k and m minimal and $k < m$, then $f^{\circ(k+1)}(T')$ and $f^{\circ(m+1)}(T')$ are both intervals with c_1 as endpoints and not containing branch points of T, and $m - k < n$. Hence $f^{\circ(m-k)}$ must map T' onto itself reversing the orientation, so it fixes some point in T' which must have an itinerary with period dividing n. This is a contradiction. $\square$

3. Triods and Uniqueness of Hubbard Trees

In this section we introduce triods and their iteration in order to determine which marked points are branched and/or adjacent to each other. By analyzing the triods in T, the uniqueness of Hubbard trees, up to equivalence, follows as well.

3.1. Definition (Triod)
A triod *is a connected compact set homeomorphic to a subset of the letter* Y. *It is* degenerate *if it is homeomorphic to an arc or a point.*

For a sequence $\nu \in \Sigma^\star$, let $\star\nu$ be the symbol $\star$ followed by ν and define $S(\nu) := \{\star\nu, \nu, \sigma(\nu), \sigma^{\circ 2}(\nu), \dots\}$ (the orbit of $\star\nu$ under the shift). Then $\sigma(S(\nu)) \subset S(\nu)$. Note that $0\nu \in S(\nu)$ or $1\nu \in S(\nu)$ would imply that ν was periodic but not $\star$-periodic; thus $S(\nu) \cap \{0\nu, 1\nu\} = \emptyset$ for $\star$-periodic and preperiodic ν.

3.2. Definition (Formal Triod)
Any triple of pairwise different sequences $s, t, u \in S(\nu) \cup \{0, 1\}^{\mathbb{N}^*}$ *is called a formal triod* $[s, t, u]$.

3.3. Definition (The Formal Triod Map)
If $\{s, t, u\} \cap \{0\nu, 1\nu\} = \emptyset$, *then we define the formal triod map as follows:*

$$\varphi[s, t, u] := \begin{cases} [\sigma(s), \sigma(t), \sigma(u)] & \text{if } s_1 = t_1 = u_1 \in \{0, 1\}; \\ \texttt{stop} & \text{if } \{s_1, t_1, u_1\} = \{0, 1, \star\}; \\ [\sigma(s), \sigma(t), \nu] & \text{if } s_1 = t_1 \neq u_1 ; \\ [\sigma(s), \nu, \sigma(u)] & \text{if } s_1 = u_1 \neq t_1 ; \\ [\nu, \sigma(t), \sigma(u)] & \text{if } t_1 = u_1 \neq s_1 . \end{cases} \tag{1}$$

By construction, the only sequence which starts with $\star$ is $\star\nu$, so at most one of s, t, u can start with $\star$. If one of them does, then the other two sequences either have first entries which are different from each other (and

we are in line 2), or the other two first entries are equal and we are in lines 3–5. Therefore, the list covers all possible cases.

In all cases other than the `stop` case, $\varphi[s,t,u]$ returns three sequences in $S(\nu) \cup \{0,1\}^{\mathbb{N}^*}$. These form another formal triod, *i.e.*, all three image sequences are different: in line 1, this is clear; in the other lines, this follows because s,t,u are all different from 0ν and 1ν by assumption.

To use φ more freely in the rest of this paper, it helps to specify what is meant by majority and minority votes and chopping of arms of triods.

3.4. Definition (Majority/Minority Vote, Chopped Triods.)

The majority vote *of three sequences s,t,u is the symbol that occurs most frequently among their first symbols $\{s_1, t_1, u_1\}$. If two of these symbols are equal (called a* two-to-one majority*) then the value of the third symbol is the* minority vote.

Given a triod $[s,t,u]$ with two-to-one majority vote, the chopped off *version is the triod in which the sequence of the minority vote is replaced by $\star\nu$.*

Thus, in the last three lines of formula (1), we say that u, t, s (respectively) are chopped off from the triod and replaced by $\star\nu$. The image $\varphi[s,t,u]$ is the shift of this chopped triod. If one of s, t or u equals $\star\nu$, then this sequence is chopped off and replaced by itself, so formally the outcome is the same as it would be in line 1, but we record the chopping.

3.5. Proposition (Uniqueness of Hubbard Trees)

Any two Hubbard trees with the same $\star$-periodic or preperiodic kneading sequence are equivalent.

PROOF. Given three marked points x,y,z on a Hubbard tree, we denote the triod that they form by $[x,y,z]$, *i.e.,* $[x,y,z] = [x,y] \cup [y,z] \cup [z,x]$. For any two Hubbard trees T and T' with the same kneading sequence, we prove that any pair of triods $[c_k, c_l, c_m]$ and $[c'_k, c'_l, c'_m]$ are both non-degenerate or both degenerate in the same way.

We decide whether a triod is degenerate by iterating it. Assume that c_k, c_l and c_m are pairwise different.

(1) If the triod $[c_k, c_l, c_m]$ does not contain 0 in its interior, then it maps homeomorphically onto its image; we take the image.

(2) If 0 belongs to the interior of $[c_k, c_l, c_m]$ and $0 \notin \{c_k, c_l, c_m\}$, then we take the component of $[c_k, c_l, c_m] \setminus \{0\}$ containing two of the three points c_k, c_l, c_m, and take the closure of its image as new triod (we chop off the arc from 0 to the isolated endpoint of the triod and map only the rest).

(3) If 0 belongs to the interior of $[c_k, c_l, c_m]$, and 0 is equal to one of the three points, say $0 = c_k$, then the algorithm terminates. The triod is degenerate, and c_k is an interior point of $[c_k, c_l, c_m]$.

We iterate this procedure. Since the critical orbit is finite, the algorithm either terminates or eventually reaches a loop. If the algorithm never terminates, then at least two endpoints must be chopped off during the iteration of the triod. Otherwise, at least two endpoints must have identical itineraries (if ν is $\star$-periodic, then we must exclude the case that the triod iteration involves a triod with endpoint 0ν or 1ν; but this is clear). If each of the three points of the triod is chopped off at some step, the triod must be non-degenerate. If exactly one endpoint of the triod is never chopped off, then the triod is degenerate with the latter endpoint in the middle.

The key observation is that the type of the triod can be read off from the itineraries of its endpoints in terms of the form triod map. The triod $[c_k, c_l, c_m]$ is represented by the triple $(\sigma^{\circ(k-1)}(\nu), \sigma^{\circ(l-1)}(\nu), \sigma^{\circ(m-1)}(\nu))$, and the image triod has endpoints $\varphi[\sigma^{\circ(k-1)}(\nu), \sigma^{\circ(l-1)}(\nu), \sigma^{\circ(m-1)}(\nu)]$:

(1) If the first entries of $\sigma^{\circ(k-1)}(\nu), \sigma^{\circ(l-1)}(\nu), \sigma^{\circ(m-1)}(\nu)$ are the same, where $\star$ counts as 0 (resp. 1) if the other two first entries are 0 (resp. 1), then the shifted triple represents the image triod $[c_{k+1}, c_{l+1}, c_{m+1}]$.

(2) If the first entries of $\sigma^{\circ(k-1)}(\nu), \sigma^{\circ(l-1)}(\nu), \sigma^{\circ(m-1)}(\nu)$ are 0 (say twice) and 1 (say once), then we take the shift of the sequences starting with 0 and replace the remaining sequence by ν. This represents the chopping off one arm of the triod.

(3) If the first entries of $\sigma^{\circ(k-1)}(\nu), \sigma^{\circ(l-1)}(\nu), \sigma^{\circ(m-1)}(\nu)$ are $\{0, \star, 1\}$, then we do not define $\varphi[\sigma^{\circ(k-1)}(\nu), \sigma^{\circ(l-1)}(\nu), \sigma^{\circ(m-1)}(\nu)]$: the iteration terminates, the triod is degenerate and the sequence starting with $\star$ represents an interior point of the triod.

The kneading sequence fully describes the behavior of φ and thus determines which points on the critical orbit are between which others on the tree, and which are endpoints.

For any non-degenerate triod, the iteration of φ also encodes the itinerary of the interior branch point: this itinerary is constructed by majority vote from the first entries of the sequences of the triple at every step. The branch points have itineraries different from 0ν and 1ν because they are marked points, and their images are different from the critical value. Therefore, the argument above can also be applied to triods whose endpoints are arbitrary marked points (branch points or points on the critical orbit), and this implies that any two Hubbard trees with the same kneading sequence are equivalent. □

An independent proof of the same result is given in [BS]: there we prove that all endpoints and all branch points can be read off explicitly from the kneading sequence, and that corresponding points are connected by edges; and this exactly amounts to uniqueness of the Hubbard tree.

3.6. Corollary (Uniqueness of extended Hubbard trees)

Any extended Hubbard tree is unique up to equivalence.

PROOF. This proof is the same as the proof of Proposition 3.5, as long as we include the points with itinerary in $\hat{V}$. Since $\hat{V} \not\ni 0\nu, 1\nu$, triods can be iterated under φ without collapsing their endpoints. $\square$

4. THE ρ-FUNCTION AND ITS PROPERTIES

Next we introduce the ρ-function on $\mathbb{N}^*$, which depends on ν. It is a useful tool in revealing the properties of precritical and periodic points of T, see [BS], but in this paper ρ will be used to derive some fundamental facts about sequences in $\Sigma^\star$. In particular, we show symbolically that c_1 is always an endpoint of T.

4.1. Definition (ρ-Function and Internal Address)

For a sequence $\nu \in \Sigma^\star$, define

$$\rho_\nu : \mathbb{N}^* \to \mathbb{N}^* \cup \{\infty\}, \quad \rho_\nu(n) = \inf\{k > n : \nu_k \neq \nu_{k-n}\}.$$

We usually write ρ for ρ_ν. For $k \geq 1$, we call

$$\mathrm{orb}_\rho(k) := k \to \rho(k) \to \rho^{\circ 2}(k) \to \rho^{\circ 3}(k) \to \ldots$$

the ρ-orbit of k. The case $k = 1$ is the most important one; we call

$$\mathrm{orb}_\rho(1) = 1 \to \rho(1) \to \rho^{\circ 2}(1) \to \rho^{\circ 3}(1) \to \ldots$$

the internal address *of ν. For real unimodal maps, the numbers $\rho^{\circ k}(1)$ are known as the* cutting times *of the map. If $\rho^{\circ k+1}(1) = \infty$, then we say that the internal address is finite: $1 \to \rho(1) \to \ldots \to \rho^{\circ k}(1)$; as a result, the orbit orb_ρ is a finite or infinite sequence that never contains ∞.*

REMARK. The use of the ρ-function is explained by its relation to the *closest precritical points* on the Hubbard tree. Let $\zeta_1 := 0$ be the closest precritical point of STEP 1, because it takes one iterate to map to the critical value. Given closest precritical point ζ_k of STEP k, then next closest precritical point is $\zeta_l \in (\zeta_k, c_1]$ such that $f^{\circ l}(\zeta_l) = c_1$, and l is the minimal number with this property. It is not hard to verify, see [BKS], that $\mathrm{orb}_\rho(1)$ finds the sequence of closest precritical points in $[0, c_1]$ while $\mathrm{orb}_\rho(\rho(k) - k)$ finds the sequence of closest precritical points between the critical value c_1 and the postcritical point c_k.

The following technical lemma on the behavior of ρ-orbits for arbitrary kneading sequences establishes a combinatorial result which will imply that the critical value is an endpoint of the Hubbard tree. The lemma can be interpreted by saying that within the Hubbard tree, the arcs from c_1 to the critical point and to any other postcritical point c_k intersect, so c_1 is an endpoint of the tree[2].

4.2. Lemma (Critical Value is Endpoint (Combinatorial Version))

Let $\nu \in \Sigma^\star$ be arbitrary. Then for each $k \in \mathbb{N}^*$ such that $\rho(k) - k < \infty$, there exists an i with $\rho^{\circ i}(\rho(k) - k) \le \rho(k)$ such that $\rho^{\circ i}(\rho(k) - k) \in \mathrm{orb}_\rho(1)$.

PROOF. If ν is $\star$-periodic, say of period N, then $m := \rho(k) - k$ is either infinite or at most N (depending on whether or not k is divisible by N). In the finite case, $N \in \mathrm{orb}_\rho(m) \cap \mathrm{orb}_\rho(1)$ inevitably. Therefore we only need to consider $\nu \in \Sigma^1$. We argue by induction on m, using the induction hypothesis IH[m]:

> **IH[m]:** For every $\nu \in \Sigma^1$ and corresponding ρ-function for which there exists a k such that $\rho(k) - k = m$, the orbits $\mathrm{orb}_\rho(1)$ and $\mathrm{orb}_\rho(m)$ intersect at the latest at $\rho(k)$.

REMARK. IH[m] does not imply that $\mathrm{orb}_\rho(1) \cap \mathrm{orb}_\rho(m)$ contains $\rho(k)$, not even if k is minimal such that $\rho(k) - k = m$. For example, if

$$\nu = 1011001101101\ldots$$

with $m = 6$ and $k = 7$, then $m \in \mathrm{orb}_\rho(1)$, but $\rho(m) > k > m$.

The induction hypothesis is trivially true for $m = 1$. So assume that IH[m'] holds for all $m' < m$. Take $\nu \in \Sigma^1$ arbitrary and k minimal such that $\rho(k) - k = m$. If no such k exists, then IH[m] is true for this ν by default. Let $m_0 \in \mathrm{orb}_\rho(m)$ be maximal such that $m_0 \le \rho(k)$; thus $\rho(m_0) > \rho(k)$. We distinguish two cases:

Case I: $m_0 < \rho(k)$. If $m_0 \le k$, then $\rho(m_0) > \rho(k)$ implies $m_0 < k$ and

$$\nu_1 \ldots \nu_{k-m_0+1} \ldots \nu_{\rho(k)-m_0} = \nu_{m_0+1} \ldots \nu_{k+1} \ldots \nu_{\rho(k)} \ ,$$

hence $\rho(k - m_0) - (k - m_0) = \rho(k) - k = m$, contradicting minimality of k. Therefore $k < m_0 < \rho(k)$. Since $\rho(m_0) > \rho(k)$ and

$$\nu_{k+1} \ldots \nu_{m_0+1} \ldots \nu_{\rho(k)} = \nu_1 \ldots \nu_{m_0-k+1} \ldots \nu'_m$$

(where ν'_m is the opposite symbol of ν_m), we have $\rho(m_0 - k) = m$. Consider $\tilde{\nu} := \nu_1 \ldots \nu_{m_0-1} \nu'_{m_0} \nu_{m_0+1} \ldots$ (with arbitrary continuation) with associated function $\tilde{\rho}$. Then $\tilde{\rho}(k) = m_0$.

---

[2]By the same token, the number of arms of c_1 *in the Julia set* is the number disjoint ρ-orbits in $\mathbb{N}^*$.

(i) If $m_0 = m$, then the fact that $\rho(m_0 - k) = m$ implies $\tilde{\rho}(m_0 - k) > m_0$, so $\tilde{\rho}(k) \notin \mathrm{orb}_{\tilde{\rho}}(m_0 - k)$.

(ii) If $m_0 > m$, then $\tilde{\rho}(k) = m_0 \notin \mathrm{orb}_{\tilde{\rho}}(m)$, and $\tilde{\rho}(m_0 - k) = \rho(m_0 - k) = m < m_0$ again implies $\tilde{\rho}(k) \notin \mathrm{orb}_{\tilde{\rho}}(m_0 - k)$.

So in both cases $\tilde{\rho}(k) \notin \mathrm{orb}_{\tilde{\rho}}(m_0 - k)$. Now $\tilde{\rho}(k) - k = m_0 - k < \rho(k) - k = m$, so by the induction hypothesis $\mathrm{IH}[m_0 - k]$, $\mathrm{orb}_{\tilde{\rho}}(1)$ and $\mathrm{orb}_{\tilde{\rho}}(m_0 - k)$ meet at or before $\tilde{\rho}(k)$; since $\tilde{\rho}(k) \notin \mathrm{orb}_{\tilde{\rho}}(m_0 - k)$, they meet before $\tilde{\rho}(k) = m_0$.

As a result, also $\mathrm{orb}_{\rho}(1)$ and $\mathrm{orb}_{\rho}(m_0 - k)$ meet before $m_0 < \rho(k)$, and since $\rho(m_0 - k) = m$, it follows that $\mathrm{orb}_{\rho}(1)$ and $\mathrm{orb}_{\rho}(m)$ meet before $\rho(k)$.

Case II: $m_0 = \rho(k)$. In this case $\rho(k) \in \mathrm{orb}_{\rho}(m)$. Let $n_0 \in \mathrm{orb}_{\rho}(1)$ be maximal such that $n_0 \le \rho(k)$, hence $\rho(n_0) > \rho(k)$. If $n_0 = \rho(k)$ there is nothing to prove, so assume that $n_0 < \rho(k) < \rho(n_0)$. As in Case I (by minimality of k), we only need to consider the case that $k < n_0 < \rho(k) < \rho(n_0)$. Since $\nu_{k+1} \ldots \nu_{\rho(k)} = \nu_1 \ldots \nu'_m$, we have $\rho(n_0 - k) = m$ (similarly as above). Set $\tilde{\nu} := \nu_1 \ldots \nu'_{n_0} \ldots$ with associated function $\tilde{\rho}$. Then $\tilde{\rho}(k) = n_0 < \rho(k)$ and by $\mathrm{IH}[n_0 - k]$, $\mathrm{orb}_{\tilde{\rho}}(1)$ and $\mathrm{orb}_{\tilde{\rho}}(n_0 - k)$ meet at the latest at $\tilde{\rho}(k) = n_0$.

(i) If $m < n_0$, then $\tilde{\rho}(n_0 - k) = \rho(n_0 - k) = m$, so $\mathrm{orb}_{\tilde{\rho}}(1)$ and $\mathrm{orb}_{\tilde{\rho}}(m)$ meet at the latest at n_0. But $n_0 \notin \mathrm{orb}_{\tilde{\rho}}(1)$, so in fact $\mathrm{orb}_{\tilde{\rho}}(1)$ and $\mathrm{orb}_{\tilde{\rho}}(m)$ meet before n_0. But then $\mathrm{orb}_{\rho}(1)$ and $\mathrm{orb}_{\rho}(m)$ also meet before $n_0 < \rho(k)$.

(ii) If $m = n_0$, then $\mathrm{orb}_{\rho}(1)$ and $\mathrm{orb}_{\rho}(m)$ obviously meet at $n_0 < \rho(k)$.

(iii) The case $m > n_0$ is impossible: We have $\rho(k) - k = m > n_0 > k$, so $\rho(k) > 2k$. Since $\rho(n_0) > \rho(k) = m + k > n_0 + k$, we find that $\nu_{k+1} \ldots \nu_{n_0+1} \ldots \nu_{\rho(k)-1} = \nu_1 \ldots \nu_{n_0-k+1} \ldots \nu_{\rho(k)-k-1}$, hence $\rho(n_0 - k) \ge \rho(k) - k > n_0$. For the sequence $\tilde{\nu}$ this means that $\tilde{\rho}(n_0 - k) = n_0$, while $n_0 \notin \mathrm{orb}_{\tilde{\rho}}(1)$. Therefore $\mathrm{orb}_{\tilde{\rho}}(1)$ and $\mathrm{orb}_{\tilde{\rho}}(n_0 - k)$ do not meet at or before n_0; since $\tilde{\rho}(k) - k = n_0 - k$, this contradicts $\mathrm{IH}[n_0 - k]$.

This completes Case II and proves that $\mathrm{orb}_{\rho}(1)$ and $\mathrm{orb}_{\rho}(m)$ intersect at the latest at $\rho(k)$, where k is minimal with the property that $\rho(k) - k = m$. For an arbitrary k with $\rho(k) - k = m$, let k' minimal with this property. Then the two orbits meet at the latest at $\rho(k') = m + k' \le m + k = \rho(k)$, so the statement holds for arbitrary k. This proves $\mathrm{IH}[m]$.

To prove the lemma for $\nu \in \Sigma^1$ and arbitrary k, set $m := \rho(k) - k$. Then $\mathrm{orb}_{\rho}(m)$ and $\mathrm{orb}_{\rho}(1)$ meet at the latest at $\rho(k)$. This proves the first statement.

Now for the second statement, assume that k' is the largest multiple of k less than $\rho(k)$, so $\rho(k') = \rho(k)$. Since k failing the admissibility means that $k \in \mathrm{orb}_{\rho}(\rho(k') - k')$ but $k \notin \mathrm{orb}_{\rho}(1)$, the second statement is an immediate consequence. $\square$

The following combinatorial lemma will be used to locate the images of certain closest precritical points in Hubbard trees.

4.3. Lemma (Combinatorics of ρ-Orbits)
Let $\nu \in \Sigma^1$ (not containing a $\star$) and let m belong to the internal address of ν.

(1) *If s is such that $s < m < \rho(s)$, then $\mathrm{orb}_\rho(\rho(m-s) - (m-s)) \ni m$.*

(2) *If $\rho(m) > 2m$, then for every $s \in \{1, \ldots, m\}$, either m or $2m$ belongs to $\mathrm{orb}_\rho(s)$.*

(3) *If $\rho(m) = \infty$, then m is the exact period of ν.*

PROOF. Let $\nu = \nu_1\nu_2 \ldots$ be the kneading sequence. Using Lemma 4.2 we prove two claims.

Claim 1: *If $k \geq 1$ and $l \in \mathrm{orb}_\rho(k) \cap \mathrm{orb}_\rho(1) \setminus \{k\}$, then $l \in \mathrm{orb}_\rho(\rho(k) - k)$.*

Assume by contradiction that $l \notin \mathrm{orb}_\rho(\rho(k) - k)$. Consider $\tilde{\nu} = \overline{\nu_1 \ldots \nu_l}$ and let $\tilde{\rho}$ be the corresponding ρ-function. Note that $\rho(k) = \tilde{\rho}(k)$ because $l \geq \rho(k)$. We have $l = \max \mathrm{orb}_{\tilde{\rho}}(1)$, but $l \notin \mathrm{orb}_{\tilde{\rho}}(\rho(k) - k)$. This contradicts Lemma 4.2.

Claim 2: *If $k < m < \rho(k)$ and $m' \in \mathrm{orb}_\rho(1)$ is such that $\rho(m') = m$, then $m' \in \mathrm{orb}_\rho(m - k)$.*

Consider $\tilde{\nu} = \overline{\nu_1 \ldots \nu_{m'}}$. Then $m' = \max \mathrm{orb}_{\tilde{\rho}}(1)$ and $\tilde{\rho}(k) = m$ (this follows directly from $\rho(k) > m$ and $\rho(m') = m$). Therefore Lemma 4.2 implies that $m' \in \mathrm{orb}_{\tilde{\rho}}(m - k)$, and Claim 2 is proved.

To prove the first assertion, let m' be as in Claim 2. We have $m' \in \mathrm{orb}_\rho(m-s)$. Hence $m = \rho(m') \geq \rho(m-s)$. By Claim 1, $m \in \mathrm{orb}_\rho(\rho(m-s) - (m-s))$.

Now we prove the second assertion. If $m \in \mathrm{orb}_\rho(s)$ there is nothing to prove. Hence we may assume without loss of generality that $s < m < \rho(s)$.

Assume first that $\rho(s) < \infty$. Assertion (1) implies $\rho(m-s) - (m-s) \leq m$, hence $\rho(m-s) \leq 2m - s < \rho(m) - s$. If $\rho(s) \geq \rho(m)$, then

$$\nu_1 \ldots \nu_{m-s}\nu_{m-s+1} \ldots \nu_{\rho(m)-s-1} = \nu_{s+1} \ldots \nu_m\nu_{m+1} \ldots \nu_{\rho(m)-1}$$
$$= \nu_{s+1} \ldots \nu_m\nu_1 \ldots \nu_{\rho(m)-m-1},$$

hence $\rho(m-s) \geq \rho(m) - s$, a contradiction. Hence $\rho(m) > \rho(s)$ and, similarly as above,

$$\nu_{s+1} \ldots \nu_m\nu_1 \ldots \nu_{\rho(s)-m} = \nu_{s+1} \ldots \nu_m\nu_{m+1} \ldots \nu_{\rho(s)}$$
$$= \nu_1 \ldots \nu_{m-s}\nu_{m-s+1} \ldots \nu'_{\rho(s)-s},$$

where ν'_j denotes the opposite symbol of ν_j. Therefore $\rho(m-s) = \rho(s) - s$. By Assertion (1), $m \in \mathrm{orb}_\rho(\rho(m-s) - (m-s)) = \mathrm{orb}_\rho(\rho(s) - m)$. (Note

that this implies $\rho(s) \leq 2m$.) Since $\nu = \nu_1 \ldots \nu_m \nu_1 \ldots \nu_m \ldots$, we then have $2m \in \mathrm{orb}_\rho(\rho(s))$ as claimed (we simply start m entries later).

If $\rho(s) = \infty$, we first change the entry ν_k for some $k > 2m$. Then still $\rho(m) > 2m$ and $\rho(s) = k > 2m$, and we can use the above argument.

For the third assertion, consider $\tilde{\nu} = \nu_1 \ldots \nu_m \nu_1 \ldots \nu_m \nu_1' \ldots$ with ρ-function $\tilde{\rho}$. Then $\tilde{\rho}(m) = 2m + 1 > 2m$. If $s < m$ is the exact period of ν, then $\rho(s) = \infty$ and $\tilde{\rho}(s) = 2m + 1$. Hence both m and $2m \notin \mathrm{orb}_{\tilde{\rho}}(s)$, contradicting the second assertion. $\qquad\square$

The following lemma is rather trivial, but helpful to refer to in longer arguments.

4.4. Lemma (Translation Property of ρ)
If $\rho(m) > km$ for $k \geq 2$, then $\rho(km) = \rho(m)$.

PROOF. Let ν be a kneading sequence associated to ρ. Then $\rho(m) > km$ says that the first m entries in ν repeat at least k times, and $\rho(m)$ finds the first position where this pattern is broken. By definition, $\rho(km)$ does the same, omitting the first k periods. $\qquad\square$

4.5. Lemma (ν is Always Endpoint)
Fix a sequence $\nu \in \Sigma^\star$. If the triod $[\sigma^{\circ k}(\nu), \sigma^{\circ l}(\nu), \nu]$ (with $k, l \geq 1$) can be iterated infinitely often, then the sequence ν will be chopped off eventually. If the stop *case is reached, the initial $\star$ does not occur in the last sequence. Hence ν cannot be an interior point of such a triod.*

PROOF. Suppose the triod can be iterated infinitely often but ν is never chopped off. By the definition of the ρ-function, $\sigma^{\circ k}(\nu)$ differs from ν differ at entry $\rho(k)$, so one of the two is chopped off at STEP $\rho(k) - k$. If the third sequence is chopped off, then we are done, so we may assume that it is the first sequence which is chopped off. Therefore,

$$\varphi^{\circ(\rho(k)-k)}\left([\sigma^{\circ k}(\nu), \sigma^{\circ l}(\nu), \nu]\right) = [\nu, \cdots, \sigma^{\circ(\rho(k)-k)}(\nu)].$$

(The second entry depends on whether or not $\sigma^{\circ l}(\nu)$ has been chopped off in the meantime.) The next time that the first and third sequences differ is at iterate $\rho(\rho(k) - k)$, and again we may assume the first sequence is chopped off, yielding

$$\varphi^{\circ \rho(\rho(k)-k)}\left([\sigma^{\circ k}(\nu), \sigma^{\circ l}(\nu), \nu]\right) = [\nu, \cdots, \sigma^{\circ \rho(\rho(k)-k)}(\nu)] \ ,$$

and in general the first sequence is chopped off at steps in $\mathrm{orb}_\rho(\rho(k) - k)$.

Similarly, the second sequence is chopped off at iteration steps in $\mathrm{orb}_\rho(\rho(l) - l)$. By Lemma 4.2, there is an iteration step $s \in \mathrm{orb}_\rho(\rho(k)-k) \cap \mathrm{orb}_\rho(\rho(l)-l)$.

Take s minimal with this property. At this step, the first and second sequences both differ from the third, so in this step the sequence ν must be chopped off (or the `stop` case is reached, but this is excluded by hypothesis). This settles the first claim.

For the second claim, the `stop` case can be reached only if ν is $\star$-periodic, say of period n. We can assume by Lemma 4.4 that $1 \leq k < l < n$. By Lemma 4.2, $\mathrm{orb}_\rho(\rho(k) - k)$ and $\mathrm{orb}_\rho(\rho(l) - l)$ intersect $\mathrm{orb}_\rho(1)$ at an entry $s \leq \max\{\rho(k), \rho(l)\}$, and since ν is $\star$-periodic, $\max\{\rho(k), \rho(l)\} \leq n$. The above argument implies that ν is chopped off at some iterate $s' \leq s$. If $s' = n$, then $\rho(k) = n$ or $\rho(l) = n$. If $\rho(l) \geq n$, then the `stop` case is reached after $n - l$ iterates of φ, and the $\star$ is in the second sequence. If $\rho(k) \geq n$, then the `stop` case is reached after $n - k$ iterates of φ, and the $\star$ is in the first sequence. Finally, if $s < n$, then the triod after s iterates will be $\varphi^{\circ s}([\sigma^{\circ k}(\nu), \sigma^{\circ l}(\nu), \nu]) = [\sigma^{\circ k'}(\nu), \sigma^{\circ l'}(\nu), \nu]$ for some $k', l' \leq n$, and we can repeat the argument. $\qquad\square$

5. Existence of Hubbard Trees

In this section we prove the existence of Hubbard trees for $\star$-periodic and preperiodic kneading sequences. Starting with the set $\{\star\nu, \nu, \sigma(\nu), \dots\}$ we use the triod map to find the branch points and therefore all marked points (represented by itineraries) of the tree. Further triod arguments then determine how to connect marked points by edges, and allow for verifying the properties of the dynamics on the tree.

Fix a $\star$-periodic or preperiodic sequence ν and set

$$S(\nu) = \{\star\nu, \nu, \sigma(\nu), \sigma^{\circ 2}(\nu), \dots\}.$$

Clearly, $\sigma(S(\nu)) \subset S(\nu)$ and $S(\nu) \cap \{0\nu, 1\nu\} = \emptyset$.

In order to introduce the other marked points of the Hubbard tree, we need to analyze the behavior of the triod map of Definition 3.2 in more detail.

Branch Points $\mathfrak{b}(s, t, u)$. Fix a $\star$-periodic or preperiodic ν. Take $s, t, u \in S(\nu) \cup \{0, 1\}^{\mathbb{N}^*}$ and consider the triod $[s, t, u]$. If each of these sequences is chopped off infinitely often under iteration of φ, then $[s, t, u]$ is called a *branched* triod, and it has a "branch point" called $\mathfrak{b}(s, t, u)$: that is the sequence which is defined by majority vote of the three sequences in each iteration step. More precisely, the n-th entry of this sequence is defined by majority vote among the first entries in $\varphi^{\circ(n-1)}(s, t, u)$. Clearly, $\mathfrak{b}(s, t, u) \in \{0, 1\}^{\mathbb{N}^*}$. If s, t, u are (pre)periodic, then $\mathfrak{b}(s, t, u)$ is necessarily (pre)periodic too. Moreover, $\mathfrak{b}(s, t, u)$ differs from all sequences s, t, and u:

otherwise, if say $\mathfrak{b}(s,t,u) = s$, then because of the majority construction, the first sequence in the triod $[s,t,u]$ always coincides with at least one other sequence, so it is never chopped off, contrary to our assumption.

Types of Triods. Based on equation (1) in Definition 3.3, we can distinguish five types of triods:

(1) the triod can be iterated indefinitely so that all three sequences are chopped off infinitely often (this implies all three sequences remain distinct under iteration and the `stop` case is never reached); in this case we call the triod *branched*. The sequence $\mathfrak{b}(s,t,u) \in \{0,1\}^{\mathbb{N}^*}$ obtained by majority vote is the *branch point* of the triod.

(2) the triod can either be iterated indefinitely and precisely two sequences are chopped off infinitely often whereas the remaining sequence is never chopped, or the iteration reaches the `stop` case so that the sequence that lands on $\star$ at the `stop` case has never been chopped off before; in this case we call the triod *flat*. The sequence which reaches the $\star$ in the `stop` case or which is never chopped off is called the *middle point* of the flat triod.

(3) not all three sequences remain distinct during the iteration, *i.e.*, two or three sequences become identical at some iterate of φ; in this case we say that φ has *collapsing* sequences. Note that collapsing sequences means that one of these sequences must have been equal to 0ν or 1ν the iterate before collapsing. Note also that if φ can be iterated indefinitely, but only one sequence ever gets chopped off infinitely often, then the remaining sequences must collapse.

(4) the iteration of φ reaches the `stop` case so that the $\star$ is in a sequence that had been chopped off before.

(5) the iteration of φ can be carried on forever without collapsing, and some sequence is chopped off at least once, but not infinitely often.

For us, the most important cases are the first two because we will show in Lemma 5.1 that (3)–(5) do not occur when s,t,u are distinct and forward shifts of the kneading sequence ν.

Note that all five cases are mutually distinct and cover all possible cases: if two or sequences collapse, then the `stop` case cannot be reached and at most one sequence can be chopped off infinitely often, and we are exactly in case (3). Otherwise, there is no collapsing. If we can iterate forever without collapsing, then either some sequence is chopped off finitely many times and we are exactly in case (5); or every sequence is chopped off either infinitely many times or not at all, and the number of chopped sequences must be 2 or 3, so we are in cases (2) or (1). The last possibility is that the `stop` case

is reached; depending on whether the sequence landing on $\star$ had or had not been chopped off before, we are in case (4) or (2).

5.1. Lemma (If All Three Sequences are Chopped)

Let $\nu \in \Sigma^\star$ be $\star$-periodic or preperiodic. If $[\sigma^{\circ k}(\nu), \sigma^{\circ l}(\nu), \sigma^{\circ m}(\nu)]$ is such that each of the three sequences is chopped off at least once, then each sequence is chopped off infinitely often, and the triod can be iterated forever without reaching the stop *case.*

PROOF. First note that $\sigma^{\circ s}(\nu) \neq 0\nu$ or 1ν for all s. Therefore the triod can be iterated until the stop case is reached, if ever.

Suppose the first chopping occurs after s iterations, and it is the third sequence which is chopped off. The resulting triod is $[\sigma^{\circ(k+s)}(\nu), \sigma^{\circ(l+s)}(\nu), \nu]$. By Lemma 4.5, the stop case is never reached with the $\star$ at the third position.

By assumption, there are iteration times when the first and second sequences are chopped off, and it follows similarly that the stop case can never be reached with the $\star$ at any position. The triod can thus be iterated infinitely often. Lemma 4.5 again gives that the third sequence will be chopped off, and repeating this argument, we see that indeed all three sequences will be chopped off infinitely often. □

The structure of triods in $S(\nu)$: Every triod in $S(\nu)$ is either flat or branched: no triod can collapse, and it follows from Lemma 4.5 that every sequence which gets chopped off once will never be the center of a flat triod when the stop case is reached. Therefore if a triod cannot be iterated forever then it reaches the stop case and is flat. If a triod can be iterated forever, then at least two sequences get chopped off infinitely often, and if all three sequences get chopped off at least once, then by Lemma 5.1 the triod is branched.

5.2. Lemma (Branch Points of Branched Triods)

Fix a $\star$-periodic or preperiodic ν. Suppose that $s, t, u \in S(\nu) \cup \{0,1\}^{\mathbb{N}^*}$ are such that $[s, t, u]$ is a branched triod; set $v := \mathfrak{b}(s, t, u)$. If $\sigma^{\circ k}(w) \notin \{0\nu, 1\nu\}$ for all $w \in \{s, t, u, v\}$ and all $k \geq 0$, then $[s, t, v]$, $[s, u, v]$ and $[t, u, v]$ are flat with v in the middle.

PROOF. By assumption, the triod map can be iterated forever on $[s, t, u]$. If the iteration stops for $[s, t, v]$ after finitely many steps, then either the iteration reaches one of the sequences 0ν and 1ν (which is excluded by hypothesis), or the stop case in (1) is reached; but since $v \in \{0,1\}^{\mathbb{N}^*}$ is constructed by majority vote among the sequences s, t, u, this can never occur. The majority vote also assures that v can never be chopped off along

the iteration, so $[s, t, v]$ is flat with v in the middle. The reasoning for the other two sequences $[s, u, v]$ and $[t, u, v]$ is the same.   □

Marked points of the tree: Set

$$V := S(\nu) \cup \bigcup_{[s,t,u]} \{\mathfrak{b}(s, t, u)\},$$

where the union runs over all branched triods $[s, t, u]$ with endpoints in $S(\nu)$, so their branch points $\mathfrak{b}(s, t, u)$ are well-defined. The set V is σ-invariant because $S(\nu)$ is and $\sigma(\mathfrak{b}(s, t, u))$ equals $\mathfrak{b}(\varphi[s, t, u])$.

The triod map can be iterated: Triods with endpoints in V can be branched and flat (types (1) and (2)). In order for φ to be well-defined, we need to ensure that types (3)-(5) does not occur. In addition, we want to know that for $s, t, u \in V$, the triod map can only result in elements of V. This is the contents of the following lemma.

5.3. Lemma (V is Closed under Taking Triods)

For each triple of different sequences $s, t, u \in V$, $\mathfrak{b}(s, t, v) \in V$. In particular, no triod $[s, t, u]$ is of type (3)-(5).

PROOF. Take $s, t, u \in V$ arbitrary but distinct. Let us first show that $V \cap \{0\nu, 1\nu\} = \emptyset$, so $[s, t, u]$ is not of type (3). To see why this is true, suppose that $0\nu \in V$ (the case $1\nu \in V$ is analogous). Since $0\nu \notin S(\nu)$, this implies that $0\nu = \mathfrak{b}(s, t, u)$ for a branched triod $[s, t, u]$ with $s, t, u \in S(\nu)$. We may suppose that $s_1 = t_1 = 0$ and $u_1 \in \{0, \star, 1\}$. If $u_1 = 0$, then $\varphi[s, t, u] = [\sigma(s), \sigma(t), \sigma(u)]$ is a triod with endpoints in $S(\nu)$ and $\sigma(\mathfrak{b}(s, t, u)) = \nu$ in the middle. This contradicts that ν is an endpoint (Lemma 4.5). If $u_1 \in \{\star, 1\}$, then $\varphi[s, t, u] = [\sigma(s), \sigma(t), \nu]$ is a triod with middle point ν. This contradicts that ν is an endpoint (Lemma 4.5) once more. As a result, every triod in V can be iterated forever without sequences collapsing, unless the `stop` case is reached.

The next step is to find $s', t', u' \in S(\nu)$ such that $\mathfrak{b}(s', t', u') = \mathfrak{b}(s, t, u)$, showing that $\mathfrak{b}(s, t, u) \in V$. As s, t, u are taken distinct, at least two of them are chopped off under iteration of φ. Assume that s is chopped off first and t second. If $s \in S(\nu)$, then put $s' = s$; otherwise, $s = \mathfrak{b}(s^1, s^2, s^3)$ for some $s^1, s^2, s^3 \in S(\nu)$. We iterate the φ-map on $[s, t, u]$, and keep track what happens to $[s^1, s^2, s^3]$.

- As long as $[s, t, u]$ has unanimous vote (*i.e.*, $s_1 = t_1 = u_1$), we refrain from making a selection among s^1, s^2, s^3 and take $\varphi[s^1, s^2, s^3]$ for the next iterate.

- If s has minority vote (*i.e.*, $s_1 \neq t_1 = u_1$), then select $s' = s^k$ for any $k = 1, 2, 3$ such that s^k shares vote with s. We have $\varphi[s, t, u] = \varphi[s', t, u]$.

We iterate this algorithm until s has reached minority vote, which by the choice of s must happen eventually. It follows that $\mathfrak{b}(s, t, u) = \mathfrak{b}(s', t, u)$.

Now we repeat the argument with $[s', t, u]$. If $t \in S(\nu)$ then let $t' = t$; otherwise $t = \mathfrak{b}(t^1, t^2, t^3)$ for some $t^1, t^2, t^3 \in S(\nu)$.

- As long as $[s', t, u]$ has unanimous vote, we refrain from making a selection among t^1, t^2, t^3 and take $\varphi[t^1, t^2, t^3]$ for the next iterate.
- If t has two-to-one majority vote (*i.e.*, $s'_i \neq t_i = u_i$), disqualify the t^k (if any) with $t^k_i = s'_i$ from being selected. As both t^k and s' will be replaced by ν by the action of φ, t^k and s' will take the same vote ever after, so no other $t^{k'}$ can later share minority vote with s'.
- If t has minority vote eventually, then select $t' = t^{k'}$ for an undisqualified $k' = 1, 2, 3$ such that t and $t^{k'}$ share vote. We have $\varphi[s', t', u] = \varphi[s', t, u]$.

We iterate this algorithm until t has reached minority vote, which by the choice of t must happen eventually. It follows that $\mathfrak{b}(s, t, u) = \mathfrak{b}(s', t', u)$.

Finally we turn to u. If $u \in S(\nu)$ then take $u' = u$. Otherwise $u = \mathfrak{b}(u^1, u^2, u^3)$ for some $u^1, u^2, u^3 \in S(\nu)$. If u is never chopped off, then $[s', t', u]$ is flat with u in the middle, and no new branch point is created. If u is chopped off eventually, then we follow an algorithm as before.

- As long as $[s', t', u]$ has unanimous vote we refrain from making a selection among u^1, u^2, u^3 and take $\varphi[u^1, u^2, u^3]$ for the next iterate.
- If s' has minority vote, disqualify the u^k (if any) voting with s' from being selected. As both u^k and s' will be replaced by ν by the action of φ, u^k and s' will take the same vote ever after, so no other $u^{k'}$ can later share minority vote with s'.
- If t' has minority vote, disqualify the $u^{k'}$ (if any) voting with t' from being selected. Note that $k \neq k'$ because u^k shares vote with s'. As both $u^{k'}$ and t' will be replaced by ν by the action of φ, $t^{k'}$ and t' will take the same vote ever after.
- If u has minority vote eventually, let u' be the remaining undisqualified $u^{k''}$ which shares vote with u.

In this way we obtain $s', t', u' \in S(\nu)$ such that $\mathfrak{b}(s, t, u) = \mathfrak{b}(s', t', u')$.

Since $[s, t, u]$ behaves as $[s', t', u']$ under φ and whenever s, t or u is chopped off, this sequence is replaced by ν just as s', t' or u' is replaced by ν, it follows from Lemma 5.1 that if all three sequences s, t, u are chopped off once, they will be chopped off infinitely often. If one sequence, say s, is

never chopped off, and another sequence, say t is chopped a finite number of times only, some iterate $\varphi^{\circ k}[s, t, u]$ takes the form $[\sigma^{\circ k}, \nu, \tilde{u}]$ form some $\tilde{u} \in V$. In this triod, $\sigma^{\circ k}$ and ν are never chopped off, so they are equal. But this contradicts that $0\nu, 1\nu \notin V$. This proves that $[s, t, u]$ is not of type (5).

Finally, to prove that $[s, t, u]$ is not of type (4), assume by contradiction that the first sequence is chopped off, and later on reaches the stop case with the $\star$ in that sequence. Since $[s', t', u']$ has the same behavior under φ, it would reach stop case with the $\star$ in the first sequence as well. This contradicts Lemma 4.5. $\qquad\square$

Marked points along arcs: For sequences $s, t \in V$, set

$$E(s, t) := \{u \in V : \text{the triod } [s, t, u] \text{ is flat with } u \text{ in the middle}\} \cup \{s, t\} \ .$$

We call s and t *adjacent* if $E(s, t) = \{s, t\}$. If $u, u' \in E(s, t) \setminus \{s\}$ are different sequences, then at least one of u and u' must be chopped off the triod $[s, u, u']$ under iteration of φ. We write $u \succ u'$ if u' is chopped off, and $u' \succ u$ otherwise. If $[s, u, u']$ reaches the stop case with u in the middle, then we write $u \succ u'$ as well. (It follows from Lemma 5.5 item (1) below that the sequence u or u' which is not chopped off, is indeed the middle point of the flat triod $[s, u, u']$.) By convention we take $s \succ u$ for all $u \in E(s, t) \setminus \{s\}$.

5.4. Lemma ($\succ$ is Transitive and Linear)
The order $\succ$ is transitive and linear on $E(s, t)$.

PROOF. For unity of exposition, we will say that u is chopped off from $[s, u, u']$ by φ also if $[s, t, u]$ reaches the stop case with another sequence than u in the middle. As we will not iterate φ further in this proof once the order of sequences in $[s, t, u]$ has been established, this abuse of terminology is inconsequential.

Suppose that $u, u', u'' \in E(s, t) \setminus \{s\}$ with $u \succ u'$ and $u' \succ u''$, say u'' is chopped off from $[s, u', u'']$ at iterate l and u' is chopped off from $[s, u, u']$ at iterate k. Iterate φ on the triod $[s, u, u'']$; as long as the second or third sequence is not chopped off from $[s, u, u'']$, the chopping off of the first sequence happens at exactly the same times as for $[s, u, u']$ and $[s, u', u'']$. So let us wait until one of u and u'' is chopped off. Assume by contradiction that u is chopped off first; by the choice of k, this can only happen at an iterate $\geq k$. Since $u_i = u'_i$ for $i < k$, our assumption implies that also $u'_i = u''_i$. By definition of k, the first entries of the first and second sequence of $\varphi^{\circ k}([s, u, u'])$ agree, but the first entry of the third sequence (namely u'_k) is different. Since u' is the middle point of $[s, u', u'']$, $u'_k = u''_k$. But this

means that u'' is chopped off from $[s, u, u'']$ at this iterate, and we have a contradiction.

This shows that $u \succ u''$; hence the order $\succ$ is indeed transitive. This shows that the set $E(s, t)$ is linearly ordered by $\succ$, with maximal sequence s and minimal sequence t.

Note that $E(t, s)$ is the same set of marked points, but the ordering goes in the other direction. □

5.5. Lemma (Properties of $\succ$)

(1) If $u, u' \in E(s, t)$ are such that $s \succ u \succ u'$, then $[s, u, u']$ is a flat triod with u as middle point.

(2) If s, t, u, v are such that $u \in E(s, v)$ and $t \in E(s, u)$, then $t \in E(s, v)$ and $t \succeq u$ in the $E(s, v)$-order.

(3) If s, t, u are different sequences in V and $E(s, u)$ and $E(t, u)$ intersect only in u, then u is the middle point of the flat triod $[s, t, u]$.

(4) For all $v \in V$, there exists $s \in S(\nu)$ such that $v \in E(s, \star\nu)$.

PROOF.   In this proof, we will iterate triods $[u, u', u'']$ with sequences in $V \setminus S(\nu)$. We have seen in Lemma 5.3 that such a triod is not of type (3)-(5). We always have that $u, u', u'' \in E(s, t)$ for some $s, t \in S(\nu)$, say $u \succ u' \succ u''$ in the $E(s, t)$-order. Therefore iterating φ on $[u, u', u'']$ mimics iterating φ on $[s, u', t]$, and once u and u'' are chopped off, the two triods become identical. Therefore $[u, u', u'']$ is flat.

(1) Let k be the first iterate that u' is chopped off from $[s, u, u']$. At this iterate, $\varphi^{\circ k}([s, u, u']) = \varphi^{\circ k}([s, u, t])$, and because u is the middle point of $[s, u, t]$, u will never be chopped off neither from $[s, u, t]$ nor from $[s, u, u']$.

(2) Let l be the first time that u is chopped off from the triod $[s, t, u]$. The most work goes in proving that $t \in E(s, v)$. Suppose by contradiction that at some iterate k, the second sequence of $[s, t, v]$ is chopped off under iteration of φ. There are three cases:

- If $k > l$. Then $\varphi^{\circ l}([s, t, u]) = \varphi^{\circ l}([s, t, v])$, and since t is the middle point of $[s, t, u]$, it will not be chopped off at all.
- $k = l$. Then the first and second sequence of $\varphi^{\circ k-1}([s, t, u])$ start with the same symbol, the first and second sequence of $\varphi^{\circ k-1}([s, t, v])$ start with different symbols. Yet for both triods, these first symbols are the same, so we have a contradiction.
- $k < l$. Then the first and third sequence of $\varphi^{\circ k-1}([s, t, v])$ start with the same symbol, but the first symbol of the second sequence is different. Also the first symbols of the second and third sequence of $\varphi^{\circ k-1}([s, t, u])$ are the same. This implies that the first symbols

of the first and third sequence of $\varphi^{\circ k-1}([s, u, v])$ agree, but disagree from the first symbol of the second sequence. But this implies that u is chopped off from $[s, u, v]$, contradicting our assumption.

This shows that t is never chopped off from $[s, t, v]$, hence $t \in E(s, v)$. Since t is the middle point of $[s, t, u]$, u will be chopped off before t is, so $t \succ u$ in the $E(s, v)$ order.

(3) We iterate the triod $[s, t, u]$. If the triod is branched, then by Lemma 5.2, the branch point $v = \mathfrak{b}(s, t, u)$ belongs to both $E(s, u) \cap E(t, u)$. If $[s, t, u]$ is flat, with s in the middle, then $s \in E(s, u) \cap E(t, u)$. If $[s, t, u]$ is flat, with t in the middle, then $t \in E(s, u) \cap E(t, u)$. So in all three cases, $E(s, u) \cap E(t, u)$ contains more than one point. The remaining possibility is that u is the middle point of $[s, t, u]$.

(4) This is trivial if $v \in S(\nu)$, so assume that $v = \mathfrak{b}(s, t, u) \in V \setminus S(\nu)$ for some $s, t, u \in S(\nu)$. By Lemma 5.2, $v \in E(s, t) \cap E(t, u) \cap E(u, s)$. If $\star\nu$ equals one of s, t, u, any other sequence of s, t, u fulfills the assertion. So we can assume that s, t, u have a common first entry. If $t \in E(s, \star\nu)$, then by item (2) applied to the quadruple $s, v, t, \star\nu$, also $v \in E(s, \star\nu)$ and we are finished. Similarly, if $s \in E(t, \star\nu)$, then $v \in E(t, \star\nu)$ and again we are finished. The remaining case is that $[s, t, \star\nu]$ is a branched triod, say $v' = \mathfrak{b}(s, t, \star\nu)$. If $v = v'$, then we are finished again. If $v \succeq v'$ in the $E(s, t)$-order, then $v \in E(s, \star\nu)$ by item (2) applied to the quadruple $s, v, v', \star\nu$. Similarly, if $v' \succeq v$, then $v \in E(t, \star\nu)$. This proves item (4). $\square$

Edges of the tree: For any $s \in V$, let $E(\star\nu, s) =: \{s^0, s^1, , \ldots, s^{k-1}, s^k\}$ be as above, in decreasing $E(\star\nu, s)$-order, with $s^0 = \star\nu$ and $s^k = s$. Then attach edges $[s^i, s^{i+1}]$ to the tree, for $i = 0, 1, \ldots, k-1$. Take the union of such edges for all $s \in V$, omitting repetitions (so that every pair of marked points in V is joined by at most one edge).

5.6. Proposition (The Union of Edges is a Tree)
The union T of edges is a tree, and every endpoint of T belongs to $S(\nu)$.

PROOF. Since V is finite, T is finite. By construction, each $s \in S(\nu)$ is connected to $\star\nu$, and hence T connects all $s \in S(\nu)$. If $v \in V \setminus S(\nu)$, then by Lemma 5.5 item (4), $v \in E(s, \star\nu)$ for some $s \in S(\nu)$. Therefore T is connected.

Let us prove that T contains no loops. Since we constructed T by attaching strings of edges $E(\star\nu, s)$, T can only have a loop if the following occurs:

> There are $s, s' \in S(\nu)$ such that $t \in E(\star\nu, s) \cap E(\star\nu, s')$, but there is $u \in E(\star\nu, s) \setminus E(\star\nu, s')$ such that $\star\nu \succ u \succ t$ in the $E(\star\nu, s)$-order.

We show that the above cannot happen. Indeed, by Lemma 5.5 item (1), u is the middle point in the flat triod $[\star\nu, u, t]$. Therefore $u \in E(\star\nu, t)$. Next apply Lemma 5.5 item (2) to the quadruple $\star\nu, u, t, s'$ to conclude that $u \in E(\star\nu, s')$, but this contradicts that $u \in E(\star\nu, s) \setminus E(\star\nu, s')$. Hence T contains no loop.

Finally, each $v \in V \setminus S(\nu)$ is obtained as a branch point of a branched triod, and therefore cannot be an endpoint. Hence all the endpoints of T belong to $S(\nu)$.   □

5.7. Lemma (Components of the Tree)

There are at most two edges in T with $\star\nu$ as endpoint, and $T \setminus \{\star\nu\}$ consist of at most two trees: the sequences starting with 0 form one tree, and the sequences starting with 1 form the other.

PROOF.   If $\star\nu$ has the endpoint of at least three edges, then there are $s, t, u \in S(\nu)$ such that $E(\star\nu, s)$, $E(\star\nu, t)$ and $E(\star\nu, u)$ are pairwise disjoint, except for the common endpoint $\star\nu$. At least two, say s and t, share the first symbol. Therefore $\varphi[s, t, u] = [\sigma(s), \sigma(t), \sigma(u)]$ or $[\sigma(s), \sigma(t), \nu]$. Furthermore, $\sigma(E(\star\nu, s)) = E(\nu, \sigma(s))$ is disjoint from $\sigma(E(\star\nu, t)) = E(\nu, \sigma(t))$, except for the common endpoint ν. Lemma 5.5 item (4) states that ν is the middle point of $[\sigma(s), \sigma(t), \nu]$, but this contradicts Lemma 4.5. Therefore $\star\nu$ has at most two arms, and $T \setminus \{\star\nu\}$ consists of at most two components, each of which is connected and contain no loops.

For the last statement, recall that $E(\star\nu, s) \setminus \{\star\nu, s\}$ contains all sequences in $v \in V$ that are the middle point of a flat triod $[\star\nu, v, s]$. Applying φ to it does not result in the stop case, so s and v share the first symbol. Therefore the components of $T \setminus \{\star\nu\}$ are $\cup_{s \in S(\nu), s_1=0} E(\star\nu, s) \setminus \{\star\nu\}$ and $\cup_{s \in S(\nu), s_1=1} E(\star\nu, s) \setminus \{\star\nu\}$, and all the sequences in one component have the same the first symbol.   □

Dynamics of the tree: In order to define a map $f \colon T \to T$, set $f(s) := \sigma(s)$ for $s \in V$. For any edge $[s, t]$ between marked points $s, t \in V$, define the map $f|_{[s,t]} \colon [s, t] \to [f(s), f(t)] \subset T$ to be an orientation preserving homeomorphism. Since T is a tree, the map $f|_{[s,t]}$ is unique up to homotopy.

5.8. Lemma (Dynamics Locally Injective)

For every connected subtree $T' \subset T$ such that $\star\nu$ does not disconnect T', the restriction of f to T' is injective.

PROOF. Since T and T' are connected trees, the fact that $f(x) = f(y)$ for $x, y \in T'$ would imply that f was not locally injective on every point in $[x, y] \subset T'$. It therefore suffices to prove that f is locally injective for every $x \in T' \setminus \{\star\nu\}$.

By construction, f is locally injective on every interior point of every edge. It thus suffices to show that f is locally injective at every marked point $s \neq \star\nu$. If this was not the case, then there would be marked points $t, u \in T'$ such that $[s, t]$ and $[s, u]$ were disjoint except for the common endpoint s, while $[f(s), f(t)]$ and $[f(s), f(u)]$ had more points in common than $f(s)$. But Lemma 5.5 item (3) implies that the triod $[s, t, u]$ is flat with s in the middle, and the triod $[f(s), f(t), f(u)] = \varphi[s, t, u]$ is flat with $f(s)$ in the middle, so $f(s) \in [f(t), f(u)]$. This is a contradiction. $\qquad\square$

We are now ready to prove Theorem 2.5.

PROOF. We start with the existence proof. Set $c_0 := \star\nu$ (the critical point) and $c_k := \sigma^{\circ k}(\star\nu)$ (the critical orbit). Construct a tree T with dynamics as in Proposition 5.6. We check the six properties of a Hubbard tree:

(1) The map f is clearly continuous on T, and it is surjective on $S(\nu) \setminus \{\star\nu\}$. If $\nu = \overline{1^k\star}$ for some k, then f permutes (and hence is surjective on) $S(\nu)$, whereas if $\nu = 1^k 0 \ldots$, then $\star\nu$ is obviously the middle point of the triod $(\nu, \star\nu, \sigma^{\circ k}(\nu))$, and hence not an endpoint. By Proposition 5.6, this means that f is surjective on the set of endpoints of T, hence surjective onto all of T.

(2) In $T \setminus \{c_0\}$, any two marked points whose itineraries start with the same entry are in the same connected component by construction of T. Therefore, $T \setminus \{c_0\}$ consists of at most two connected components, and f is injective on each of the by Lemma 5.8. Therefore, each point in T has at most two preimages.

(3) It follows from Lemma 5.8 that f is locally injective at every $s \in T \setminus \{c_0\}$. It is continuous by definition, hence a local homeomorphism everywhere except at the critical point.

(4) It has been shown in Proposition 5.6 that every endpoint of T is in $S(\nu)$, i.e., on the critical orbit.

(5) The critical point is obviously periodic or preperiodic because ν is.

Before proving that the last condition, it is worthwhile to interpret the sequences in V as itineraries. By Lemma 5.7, $T \setminus \{c_0\}$ consists of at most two connected components such that two sequences $s, t \in V \setminus \{\star\nu\}$ are in the same component if and only if their initial entries coincide. Therefore, every marked point $s \in V$ (which is a sequence in $\{0, 1, \star\}$) encodes the itinerary of its own dynamics with respect to the usual partition of T induced by removing $\star\nu$.

(6) Expansivity is now trivial: marked points are in V, and these are distinguished by their itineraries. If the itineraries of s and t first differ in the k-th position, then the arc $[f^{\circ k}(s), f^{\circ k}(t)]$ contains the critical point $\star\nu$.

Uniqueness has already been proved in Proposition 3.5. $\qquad\square$

PROOF OF COROLLARY 2.6. The uniqueness is proven in Corollary 3.6. For the existence, we start by building the Hubbard tree (T, f) for ν, and then add points from $\hat{V}$ and arcs where necessary in an inductive procedure. We describe this procedure in detail, using points on the tree and their symbolic itinerary interchangeably.

For $w \in \hat{V}$, we iterate the triod map φ on $[w, s, t]$ for all marked points $s, t \in V$, to decide how to attach w to the tree. Recall that marked points $s, t \in V$ are adjacent if the arc $E(s, t) = \{s, t\}$. There are three possibilities:

(a) There are adjacent $s, t \in T$ such that $[w, s, t]$ is flat with w in the middle. Put w as a new marked point on the arc $[s, t]$.

(b) We can find $s \in V$ such that for all $t \in V$ adjacent to s, the triod $[w, s, t]$ is flat with s in the middle. In this case, attach an arc $[w, s]$ to s. (The points w and s are adjacent, as long as we don't have to add new marked point on the open arc (w, s) later in the process.)

(c) If neither (a) nor (b) hold, then there is at least one pair of adjacent marked points $s, t \in V$ such that $[w, s, t]$ is a branched triod. Let $b := \mathfrak{b}(w, s, t)$. Note that Lemma 4.5 no longer holds for $V \cup \hat{V}$, and it can indeed happen that $b = b_1 b_2 \ldots b_k 0\nu$ or $b_1 b_2 \ldots b_k 1\nu$ for some (possibly empty) word $b_1 b_2 \ldots b_k \in \{0, 1\}^k$. In this case, recode b to $b_1 b_2 \ldots b_k \star\nu$ and attach an arc $[w, b_1 b_2 \ldots b_k \star\nu]$ to $b_1 b_2 \ldots b_k \star\nu$. Put b on the arc $[s, t]$ and attach an arc $[w, b]$ to b. (The points w and b are adjacent, as long as we don't have to add new marked point on the open arc (w, b) later in the process.)

Choose the next $w' \in \hat{V}$ and repeat the whole process with the tree with dynamics and marked points created so far. After all of $\hat{V}$ is treated, we check that the resulting graph is indeed a proper extended Hubbard tree.

(i) $(\hat{T}, \hat{f})$ is closed under taking triods, *i.e.*, for any choice of marked points $s, t, u \in \hat{T}$, the branch point $\mathfrak{b}(s, t, u)$ already exists in $\hat{T}$. This is the same proof as in Lemma 5.3, but without the burden to verify that the sequences 0ν and 1ν do not appear among the marked points. This is because in the construction of extended Hubbard trees, we replace 0ν and 1ν with $\star\nu$ as under case (c) above. Note also that if $b = \mathfrak{b}(s, t, u)$ for some $s, t, u \in V \cup \hat{V}$, then $\sigma(b) = \mathfrak{b}(\varphi[s, t, u])$, and the three sequence of $\varphi[s, t, u]$ belong to $V \cup \hat{V}$ as well. Therefore the set of marked points is σ-invariant.

(ii) $\hat{T}$ has no loops. This is immediate because T has no loops, and only arcs are attached to T in the process of creating $\hat{T}$. As each of these attached arcs has a point in $\hat{V}$ as endpoint, each endpoint of $\hat{T}$ belongs to $V \cup \hat{V}$.

(iii) We define the dynamics of $\hat{f}$ on the (σ-invariant set of) marked points by σ. This shows that $\hat{f}$ is at most two-to-one on the marked points, and (since 0ν, 1ν and $\star\nu$ are identified in case (c)) also that $\hat{f}$ is expansive. An

arc $[s, t]$ between two adjacent marked points will be mapped homeomorphically onto the arc $[\sigma(s), \sigma(t)]$, precise as explained above Lemma 5.8. Lemma 5.8 itself then shows that $\hat{f}$ is locally injective, and in fact that every point in $\hat{T}$ has at most two preimages. This verifies all the conditions of an extended Hubbard tree. $\square$

REMARK.   Note that if ν is $\star$-periodic, then c_0 can get any number of arms. For example, if $\nu = \overline{10\star}$, then $T = [c_1, c_2]$ is an arc and when $\hat{V} = \{\sigma^{\circ i}(w) : i \geq 0\}$ for

$$w = \underbrace{110110110\ldots110}_{110 \text{ repeated } n \text{ times}}0\star\nu,$$

then in the extended tree $(\hat{T}, \hat{f})$, the critical point c_0 has $n + 2$ arms and c_1 and c_2 both have $n + 1$ arms, see Figure 3.
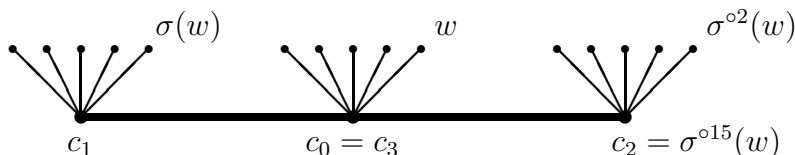

FIGURE 3. The extended Hubbard tree for $\nu = \overline{10\star}$ and $\hat{V} = \{\sigma^{\circ i}(w) : i \geq 0\}$ with $w = (110)^n 0\star\nu$ for $n = 5$. The Hubbard tree is in bold lines.

REFERENCES

[BFH]   Ben Bielefeld, Yuval Fisher, John Hubbard: *The classification of critically prepe-riodic polynomials as dynamical systems*, Journal AMS **5** 4 (1992) 721–762.

[BKS]   Henk Bruin, Alexandra Kaffl, Dierk Schleicher, *Symbolic dynamics of quadratic polynomials,* In progress. A preliminary version (Preprint 2002) is available on the web-pages of the Mittag-Leffler institute: http://www.ml.kva.se/preprints/archive2001-2002.php

[BS]   Henk Bruin, Dierk Schleicher, *Admissibility of kneading sequences and structure of Hubbard trees for quadratic polynomials*, Journal of the LMS (2008) doi: 10.1112/jlms/jdn033.

[D1]   Adrien Douady, *Julia sets and the Mandelbrot set*, in: H.-O. Peitgen, P. Richter: The beauty of fractals, Springer-Verlag, New York (1986) 161–173.

[D2]   Adrien Douady, *Algorithms for computing angles in the Mandelbrot set,* in: Chaotic dynamics and fractals Acad. Press 1986.

[DH1]   Adrien Douady, John Hubbard, *Études dynamique des polynômes complexes I & II*, Publ. Math. Orsay. (1984-85) *(The Orsay notes)*.

[DH2]   Adrien Douady, John Hubbard, *A proof of Thurston's topological characterization of rational functions*, Acta Mathematica **171** (1993) 263–297.

[HS]   John Hubbard, Dierk Schleicher, *The spider algorithm,* Proc. of Symp. Appl. Math. **49** *Complex dynamical systems: the mathematics behind the Mandelbrot and Julia sets* (1994) 155–180.

[LS]   Eike Lau, Dierk Schleicher, *Internal addresses in the Mandelbrot set and irreducibility of polynomials,* Stony Brook Preprint #**19** (1994).

[M]   John Milnor, *Dynamics in one complex variable,* Third edition. Annals of Mathematics Studies, **160**. Princeton University Press, Princeton, NJ, 2006.

[MT]   John Milnor, William Thurston, *On iterated maps of the interval,* LNM **1342** (J. C. Alexander ed.). 465–563 (1988).

[P]   Alfredo Poirier, *On postcritically finite polynomials, part two: Hubbard trees,* Stony Brook Preprint **#17** (1993).

[S1]   Dierk Schleicher, *Internal addresses in the Mandelbrot set and irreducibility of polynomials,* Submitted (2007). ArXiv math.DS/9411238.

[S2]   Dierk Schleicher, *Rational parameter rays of the Mandelbrot set,* Asterisque **261** (2000), 409–447.

[S3]   Dierk Schleicher, *Laminations, Julia Sets, and the Mandelbrot Set,* Appendix to [T]. In: Complex dynamics, families and friends, D. Schleicher (ed), A.K. Peters, Wellesley(MA), (2009), 111–137. ¿

[S4]   Dierk Schleicher, *On fibers and local connectivity of Mandelbrot and Multibrot-Sets,* In: A Mandelbrot Jubilee, Michel Lapidus (ed.), Proceedings of Symposia in Pure Mathematics Vol. 72, 1 (2004) 477-507.

[T]   William P. Thurston, *On the geometry and dynamics of iterated rational maps,* In: Complex dynamics, families and friends, D. Schleicher (ed), A.K. Peters, Wellesley(MA), (2009), 1–110.

Department of Mathematics, University of Surrey, Guildford GU2 7XH, United Kingdom
   E-mail address: H.Bruin@surrey.ac.uk

School of Engineering and Science, Jacobs University Bremen, P.O. Box 750 561, D-28725 Bremen, Germany
   E-mail address: Alexandra Kaffl <a.kaffl@web.de>

School of Engineering and Science, Jacobs University Bremen, P.O. Box 750 561, D-28725 Bremen, Germany
   E-mail address: dierk@jacobs-university.de