Ergod. Th. & Dynam. Sys. (2013), **33**, 303–317 © *Cambridge University Press, 2011*
doi:10.1017/S0143385711000733

Hubbard forests

ALFREDO POIRIER

Departmento de Ciencias, Sección Matemáticas,
Pontificia Universidad Católica del Perú, Apartado 1761,
Lima 100, Perú
(e-mail: apoirie@pucp.edu.pe)

(*Received* 27 *April* 2010 *and accepted in revised form* 1 *August* 2011)

Abstract. The theory of Hubbard trees provides an effective classification of nonlinear postcritically finite polynomial maps in $\mathbb{C}$. This note extends the classification to maps from a finite union of copies of $\mathbb{C}$ to itself. Holomorphic proper maps on a finite union of copies of $\mathbb{C}$ which are postcritically finite and nowhere linear can be characterized by a 'forest' made up of one tree for each copy of the set of complex numbers.

1. *Introduction*

Let $\mathbb{C}^{\cup} = \bigcup \mathbb{C}_i$ be a cyclically arranged disjoint union of finitely many copies of the set of complex numbers. We consider holomorphic maps $F : \mathbb{C}^{\cup} \to \mathbb{C}^{\cup}$ whose restriction to each $\mathbb{C}_i$ is a polynomial F_i taking values in $\mathbb{C}_{i+1}$. To avoid trivial behavior, we assume that each restriction has degree at least two.

For such functions, the concepts of Julia set, filled Julia set, Fatou components and so on are naturally defined. For instance, the *filled Julia set* $K(F)$ is the collection of points whose orbit under iteration remains bounded. Note that the first return map to $\mathbb{C}_i$ is a polynomial P_i whose 'classical' filled Julia set $K(P_i)$ agrees with $K_i(F) = K(F) \cap \mathbb{C}_i$. As a consequence, when all critical points belong to $K(F)$, each $K_i(F) = K(P_i)$ is connected. This fact is implied by the polynomial case.

We are interested in the specific case where the orbit of every critical point is periodic or eventually periodic, and thus finite. Such maps are *postcritically finite*. For them, each $K_i(F)$, besides being connected, is locally connected (cf. [**3**], combined with the previous argument). The boundary $\partial K(F)$ of the filled Julia set, denoted by $J(F)$, is the *Julia set*, and its elements are *Julia points*. Its complement $\mathbb{C}^{\cup} - J(F)$ is the *Fatou set*. A periodic orbit with a critical member is a *critical cycle*. In the postcritically finite case, a periodic orbit belongs to the Fatou set of F if and only if it is a critical cycle. This implies that in this setting, the interior of $K(f)$ is non-empty if and only if there exists at least one critical cycle.

In this postcritically finite context, the dynamics admits a further decomposition. If the interior of $K(F)$ is non-empty, then F sends each bounded Fatou component (necessarily a simply connected set by the maximum modulus principle) onto some other component as a branched cover map. Moreover, Sullivan [**6**] proved that all components are eventually

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

periodic (see also [**3**, Theorem 16.4]). Each of these components can be uniformized so that in local charts F reads $z \mapsto z^m$ for some $m \geq 1$ (see [**3**, Theorem 9.1]). Additionally, if U is a periodic bounded Fatou component, then the first return map to U is conjugated to $z \mapsto z^k$, this time with $k \geq 2$. In particular, loops of components are in correspondence with critical cycles. Also, inside each (not necessarily periodic) Fatou component there is a unique point that sooner or later catches up with a critical point (precisely the one marked 0 in local coordinates), referred to later as the *center*.

We pick a finite invariant set M containing all critical points of F. In practice, this is the set we want to keep track of. Next, we write M_i for the intersection $M \cap \mathbb{C}_i$. Inside $K_i(F)$ we interconnect M_i by means of *regulated arcs*, that is, arcs subject to the extra condition that when a Fatou component is encountered, this intersection is formed by radial segments in the associated coordinate. This construction defines a finite topological tree $T(M_i)$ in each component, where the set M_i together with the intrinsic branching points are to be considered vertices (cf. [**1, 2, 5**]). In dynamical terms, this preferred set satisfies $F(T(M_i)) = F_i(T(M_i)) \subset T(M_{i+1})$. This generalization of Hubbard trees is what we call a *Hubbard forest*.

In order to grasp the essential combinatorial information, we first introduce some notation. An *angled tree* T (or tree for short) is a finite connected simplicial complex of dimension zero or one together with an *angle function*, which will be described shortly. By finite we mean that T is compact, the set of vertices V is a finite subset of T, and $T - V$ is a finite union of open 1-cells, the *edges*. The angle function $e, e' \mapsto \angle_v(e, e') \in \mathbb{Q}/\mathbb{Z}$ assigns a rational number modulo 1 to each pair of edges e and e' incident at $v \in V$. This *angle* $\angle_v(e, e')$ is skew-symmetric (i.e. it obeys $\angle_v(e, e') = -\angle_v(e', e)$), satisfies $\angle_v(e, e') = 0$ if and only if $e = e'$, and is such that $\angle_v(e, e'') = \angle_v(e, e') + \angle_v(e', e'')$ whenever three edges meet at v. Since this function induces a cyclic order among edges incident at each vertex, we have a preferred isotopy class of embeddings of T into the complex plane.

The *incidence* of a vertex v is the number of edges around v. *Endpoints* are vertices of incidence 1; *branching points* are vertices of incidence greater than or equal to 3; and vertices of incidence 2 are said to be *regular vertices*.

Example. In Figure 1, we have a tree T already embedded in $\mathbb{C}$. By definition, all endpoints (such as A) and branching points (such as B) of the tree must appear in the set of vertices V. All other points in the tree have incidence 2, and they may or may not be present in V; it depends on a contextual—but explicit—choice.

Sometimes we will need precise angles around a vertex, for example near B in Figure 1.

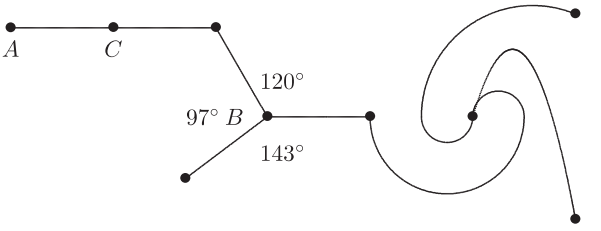


FIGURE 1. A tree embedded in $\mathbb{C}$. Strictly speaking, all vertices should have a name; however, this might be unpractical in drawings, so as a compromise we mark each vertex by a • to distinguish them from 'regular' points.

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

However, for many purposes, knowing the cyclic order of edges around a vertex is good enough.

For angled trees T_1 and T_2 with vertex sets V_1 and V_2, respectively, we need a suitable notion of morphism. Take a map $f : V_1 \to V_2$ subject to the following *non-degeneracy condition*: if v and v' are the endpoints of an edge $e \subset T_1$, then $f(v)$ and $f(v')$ are different vertices in T_2. Hence, we can extend f to the edge e so that it now covers homeomorphically the arc between $f(v)$ and $f(v')$. In other words, this function f between vertices can be identified with a map between trees $f : T_1 \to T_2$ that is a homeomorphism of each edge of T_1 onto its image.

Whenever there is a *local degree* $d : V_1 \to \{1, \ldots\}$ compatible with the dynamics in the sense that $\angle_{f(v)}(f(e), f(e')) = d(v)\angle_v(e, e')$ holds for all edges e and e' incident at $v \in V_1$, the pair (f, d) (or simply f with an abuse of notation) is an *angled (tree) map*. A vertex is said to be *critical* if $d(v) > 1$. We write Ω_f for the set of critical points. The quantity $d_f = 1 + \sum_{v\in\Omega_f}(d(v) - 1) = 1 + \sum_{v\in V_1}(d(v) - 1)$ is the *degree of* f.

More generally, a *cyclic angled graph* H (or cyclic graph for short) consists of a finite collection of angled trees labeled cyclically (say $H = \bigcup T_i$). A *cyclic angled graph dynamics* $f : H \to H$ by definition restricts in each constituent tree T_i to an angled map $f_i : T_i \to T_{i+1}$.

As usual, a vertex v is *periodic* if $f^{\circ k}(v) = v$ for some $k \geq 1$. The orbit of a periodic critical point is a *critical cycle*. A vertex is of *Fatou type* if it eventually maps to a critical cycle; otherwise it is of *Julia type* (or a *Julia vertex*).

Angles at Julia vertices are rather artificial, so we normalize them. If m edges $e_1, \ldots, e_m$ meet at a periodic Julia vertex v, then each $\angle_v(e_i, e_j)$ should be a multiple of $1/m$. (Thus, angles around a periodic Julia vertex convey no information beyond the cyclic order of the edges.) In this case, the graph is *normalized*.

Example. Hubbard trees and Hubbard forests as defined earlier determine abstract dynamical graphs. In this case, the degree function is given by the local degree. Elementary criteria show that a vertex is of Julia type if and only if it belongs to the Julia set (and, therefore, Fatou points are identified with points in the Fatou set). The angle function can be defined easily. Near a Fatou vertex (which is a center) we have a well-defined Böttcher coordinate: we measure angles there. Near a Julia vertex $v \in J(P_i) = \mathbb{C}_i \cap J(F)$ where m rays land, we define an angle between branches as a multiple of $1/m$ in such a way that this choice reflects the natural arrangement of the m components of $J(P_i) - \{v\}$. Of course, with this structure they are abstract Hubbard forests, as defined next.

If $v, v' \in V$ belong to the same component of H, the distance $\mathrm{dist}_H(v, v')$ counts the number of edges between v and v'. Otherwise, if v and v' are placed in different trees, we set $\mathrm{dist}_H(v, v') = \infty$. We say that H is *expanding* if for every edge e whose endpoints v and v' are Julia periodic vertices, there exists an $n \geq 1$ for which we have $\mathrm{dist}_H(f^{\circ n}(v), f^{\circ n}(v')) > 1$.

By an *abstract cyclic Hubbard forest* (or simply a cyclic Hubbard forest) we mean a normalized angled graph dynamics that obeys the expanding condition.

As a particular case, a *Hubbard tree* is a cyclic Hubbard forest with a single piece (cf. [**5**]). For Hubbard trees, the classification problem is already solved.

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

THEOREM 1.1. (Poirier [**5**]) *A normalized angled tree dynamics can be realized as the dynamics associated to a postcritically finite polynomial if and only if it is expanding. In other words, all Hubbard trees—and only them—can be realized. Such a realization is unique up to affine conjugation in the dynamical plane.*

As we will see in §4, for cyclic Hubbard forests conditions are similar. In other words, *a normalized cyclic angled graph dynamics can be realized as the dynamics associated to a postcritically finite proper holomorphic self-map of a cyclic collection of copies of* $\mathbb{C}$ *if and only if it is expanding.*

These ideas can be generalized. We can equally consider $\mathbb{C}^{\cup}$ as a disjoint union of finitely many (not necessarily cyclic) copies of $\mathbb{C}$, take F to be a holomorphic self-map of $\mathbb{C}^{\cup}$ whose restriction to each individual copy of $\mathbb{C}$ is a polynomial (taking values in a perhaps 'different' complex plane), and recover corresponding objects.

On the abstract side, we can take H to be a union of finitely many angled trees together with a local degree at each vertex and a compatible angled map $f : H \to H$ whose restriction to each tree has degree at least two. It should be clear here what we mean by Fatou or Julia points. A *Hubbard forest* is an expanding angled self-map on H with the previously explained normalization at periodic Julia vertices.

Now we are ready to state our main result.

THEOREM 1.2. *A normalized (not necessarily cyclic) angled graph dynamics can be realized as the dynamics associated to a postcritically finite proper holomorphic self-map (i.e. one that restricts to a polynomial in each connected component) of a finite union of copies of* $\mathbb{C}$ *if and only if it is expanding. In other words, all Hubbard forests—and only them—can be realized. Such a realization is unique up to affine conjugation in each component.*

We prove Theorem 1.2 only for cyclic forests. This certainly implies the theorem in the case where the forest is composed of a union of disjoint loops. Then, as eventually the dynamics is absorbed by the cyclic part, we proceed inductively on the number of steps needed to reach the cyclic part. In that case, our claim is an easy consequence of the following technical result (to be established in §2). This will complete the proof of Theorem 1.2.

PROPOSITION 1.3. *Let* $f : T_1 \to T_2$ *be an angled tree map of degree* $d_f \geq 2$. *Fix an embedding of* T_2 *into* $\mathbb{C}$. *Then there is an embedding of* T_1 *into* $\mathbb{C}$ *and a polynomial of degree* d_f *that realizes* $f : T_1 \to T_2$. *Once coordinates in the domain are chosen, this polynomial is uniquely determined. (Coordinates for the image were fixed at the beginning.)*

Throughout the paper we will write f for a global abstract function. Its restriction to the tree T_i will be denoted by f_i. In the cyclic case, we write p_i for the first return map to the abstract tree T_i. If any of the above is realized by a polynomial map, we will capitalize it as F, F_i or P_i, respectively.

To finish this section, we give an overview of how the paper is organized. In §2 we introduce the concept of branched (or homogeneous) cover and show how any angled

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

map can be naturally turned into one of these. Once this is achieved, the proof of Proposition 1.3 becomes routine. In §3 we analyze the first return map p_i to each tree T_i and show how this defines a Hubbard tree. By Theorem 1.1, each of those Hubbard trees can be realized by a polynomial P_i. Therefore, we have an 'honest' filled Julia set $K_1 = K(P_1)$ together with an embedding of T_1 into K_1. In §4 we construct (by backward induction) a series of canonical embeddings of $T_k, \ldots, T_1$ into the complex plane, together with polynomials $F_k, \ldots, F_2, F_1$ that realize $f_k, \ldots, f_1$, respectively. At this point we are essentially done since the uniqueness part of Proposition 1.3 implies $P_1 = F_k \circ \cdots \circ F_2 \circ F_1$. As a consequence, the holomorphic map whose restriction to each copy of $\mathbb{C}$ is F_i is postcritically finite, and it only remains to fill some gaps. Finally, in §5, we present a small application to what can be called 'postcritically finite anti-holomorphic polynomial dynamics in $\mathbb{C}$'.

2. *Abstract and concrete covers*

In this section we define as precisely as possible the concept of cover between trees. Then we prove Proposition 1.3, which says that given a map between angled trees, once the target tree is embedded, there is a unique polynomial that realizes this map.

Suppose we have an angled map $f : T_1 \to T_2$. If we want to extend T_2 to a bigger angled tree, then this extension is set-theoretical in nature and the original data provides a well-behaved angled map. The relevant question is, therefore, by how much we can extend T_1 and still have a satisfactory morphism. With the aim of answering this question, we adapt without proof some results from [**5**, §§6 and 8]

PROPOSITION 2.1. *Let $f : T_1 \to T_2$ be an angled tree map of degree d_f. There are d_f subtrees $G_1, G_2, \ldots, G_{d_f} \subset T_1$ with union T_1 such that every vertex $v \in T_1$ belongs to exactly $d(v)$ of the subtrees G_i, every edge belongs to exactly one subtree, and f restricted to each G_i is injective.*

Proof. The exact reference is [**5**, Lemma 6.2]. □

Example. See Figure 2. Asterisks denote points of local degree two; all angles are faithfully depicted. The vertex dynamics of the angled tree is given by $A_1 \to F_2$,

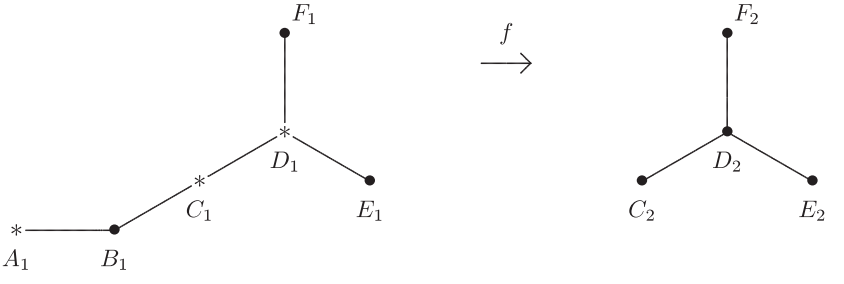


FIGURE 2. Example of an angled tree map $f : T_1 \to T_2$. The symbol $*$ denotes a point of local degree two.

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

$B_1 \to D_2$, $C_1 \to E_2$, $D_1 \to D_2$, $E_1 \to C_2$, $F_1 \to F_2$. A possible decomposition of the source tree is given in Figure 3.

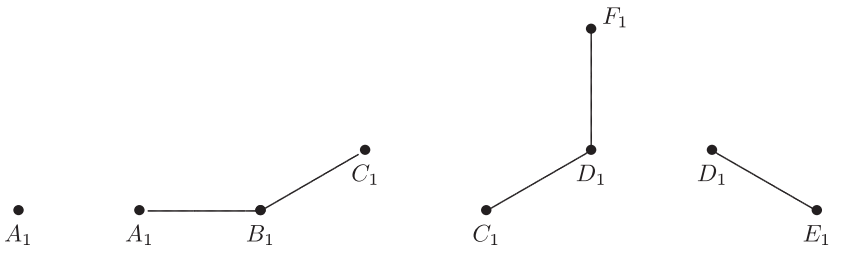


FIGURE 3. A decomposition of the source tree in Figure 2. The trick is to carefully 'unglue' a critical point v into $d(v)$ copies. Each 'decriticed' copy of v will form a bouquet that picks branches stemming from a wedge whose width corresponds to $1/d(v)$ turns of a circle.

COROLLARY 2.2. *For an angled tree map of degree d_f, every point in the image tree has at most d_f inverses, counting multiplicity.*

A degree d_f map $f : T_1 \to T_2$ (with sets of vertices V_1 and V_2, respectively) is called a *branched cover* (or simply a *cover*) if every $v \in V_2$ has d_f preimages in V_1, counting multiplicity. In other words, every vertex in V_2 has a full set of inverses in V_1, counting multiplicity (cf. Corollary 2.2). Sometimes, for the sake of emphasis, we will also use the term *homogeneous*. A standard count yields the following property.

LEMMA 2.3. *Let $|V_i|$ denote the cardinality of V_i. An angled tree map of degree d_f is a branched cover if and only if it satisfies $|V_1| - 1 = d_f(|V_2| - 1)$.*

This last lemma allows us to recognize a cover among tree maps by counting vertices. The relevance of these functions is made clear by the following results.

PROPOSITION 2.4. *Let $f : T_1 \to T_2$ be a branched cover of degree d_f. There is a decomposition of T_1 into d subtrees $G_1, G_2, \ldots G_{d_f}$ (as in Proposition 2.1) such that every vertex $v \in T_1$ belongs to exactly $d(v)$ subtrees among the G_i and each $f : G_i \to T_2$ is a homeomorphism.*

Proof. The partition in Proposition 2.1 is good enough. We already know that f restricted to G_i is injective and therefore a homeomorphism onto its image. If $f(G_i)$ were a proper subtree of T_2, then we could find an endpoint of T_2, and hence a vertex $v \in V_2$, with no preimage in G_i. But this is impossible since v has exactly d_f preimages, with at most one in each G_j. □

COROLLARY 2.5. *Let $f : T_1 \to T_2$ be a branched cover. If e is an edge in T_1 with endpoints v and v', then $f(e)$ is an edge in T_2 with endpoints $f(v)$ and $f(v')$.*

Proof. Suppose this were not true. Then there would be a vertex $w \in V_2$ standing between $f(v)$ and $f(v')$. As f restricted to e is a homeomorphism, there would be a point $q \in e$ that maps to w. However, by definition, w must be in V_1, since otherwise w would have more than d_f preimages, a situation that was rendered impossible by Corollary 2.2. □

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

Next, we turn our attention to the actual existence of branched covers. Of course, the degree at any potential new vertex (not present in the original source tree) should be one, so as not to alter the global degree.

PROPOSITION 2.6. *Given an angled tree map $f : T_1 \to T_2$, we can always extend T_1 and f to get a branched cover. This extension is unique up to homeomorphims in the domain of f.*

Proof. The construction is identical to that given in [**5**, Proposition 8.3]. □

Whenever we embed a tree into the plane and travel counterclockwise around it, we will observe a pattern of edges followed by vertices followed by edges (and so on) that repeats itself cyclically upon return to the starting point. The following definition formalizes this idea in a language suitable for our purposes.

If e and e' are consecutive edges around a vertex $v \in V$, then the triplet $\mathcal{A} = (v, e, e')$ is *a pseudoaccess based at* v (see Figure 4). (Accesses are related to the landing points of rays, which is why we avoid using this term; cf. [**5**].) With the notation above, let v' be the second endpoint of the edge e'. Near v', let e'' follow e' in the natural cyclic order (if v' is an endpoint of the tree, then $e'' = e'$). Then the pseudoaccess (v', e', e'') is the *successor* of (v, e, e').

LEMMA 2.7. *In a tree with k edges, the $2k$ pseudoaccesses are cyclically arranged.*

Proof. Both properties are trivial to show and can be checked by induction on k. See also [**5**, §9] for extra details. □

LEMMA 2.8. *Let $f : T_1 \to T_2$ be a branched cover of degree d_f. Then f induces a d_f-to-1 orientation-preserving cover between the pseudoaccesses of T_1 and T_2.*

Proof. If e and e' are edges that are consecutive at v, then $f(e)$ and $f(e')$ (which are actually edges of T_2 by Corollary 2.5) are consecutive at $f(v)$ because f is a cover. As f is homogeneous, f clearly maps consecutive pseudoaccesses in T_1 to consecutive pseudoaccesses in T_2, again by Corollary 2.5. Finally, we calculate the degree. If there are k edges in T_2, then there are $d_f k$ edges in T_1, since f is a cover. Therefore there are $2d_f k$ pseudoaccesses in T_1 but only $2k$ in T_2. It follows that the induced cyclic cover is d_f to 1. □

Whenever we have a finite cyclic set, we can always imagine it as being immersed in the unit circle $\mathbb{R}/\mathbb{Z}$. Of course, pseudoaccesses are no exception, and there is some information to be gained by doing so.

Suppose we have assigned to the cyclic list of pseudoaccesses $\mathcal{A}_1, \ldots, \mathcal{A}_{2k}$ the angles $\theta_1, \ldots, \theta_{2k}$, respectively. The first issue we want to address is what interpretation we can give to the oriented segment (θ_i, θ_{i+1}) when θ_i and θ_{i+1} are consecutive as pseudoaccesses. In this case, these values really represent triplets (v_i, e_i, e_i') and $(v_{i+1}, e_{i+1}, e_{i+1}')$, where e_j' follows e_j in the natural order around v_j (here $j = i$ or $i + 1$). Also, as θ_{i+1} is the successor of θ_i, by definition we have the equality $e_i' = e_{i+1}$, where this edge goes from the base point of $\mathcal{A}_i$ to the base point of $\mathcal{A}_{i+1}$. A traveler around the tree can identify the

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

segment (θ_i, θ_{i+1}) with the oriented edge $e'_i = e_{i+1}$ that runs between the mileposts θ_i and θ_{i+1} standing at the vertices v_i and v_{i+1}.

Now we may ask the opposite question: how does an edge e appear in this circle representation? The answer is again fairly straightforward: it appears twice. Let us see why. An edge e has two endpoints v and $\tilde{v}$, say. Near v (respectively, $\tilde{v}$), let the edge a (respectively, $\tilde{a}$) be immediately anterior to e, and let p (respectively, $\tilde{p}$) be immediately posterior to e. In this setup, by definition, the only possible accesses that involve e are (v, a, e) and (v, e, p) based at v, and $(\tilde{v}, \tilde{a}, e)$ and $(\tilde{v}, e, \tilde{p})$ based at $\tilde{v}$. Also note that $(\tilde{v}, e, \tilde{p})$ follows (v, a, e) while (v, e, p) follows $(\tilde{v}, \tilde{a}, e)$ in the pseudoaccess order. Once we assign arguments $\theta_{(v,a,e)}$, $\theta_{(v,e,p)}$, $\theta_{(\tilde{v},\tilde{a},e)}$ and $\theta_{(\tilde{v},e,\tilde{p})}$ to them, we conclude that e shows up in $\mathbb{R}/\mathbb{Z}$ as the intervals $(\theta_{(v,a,e)}, \theta_{(\tilde{v},e,\tilde{p})})$ and $(\theta_{(\tilde{v},\tilde{a},e)}, \theta_{(v,e,p)})$. For future reference, we remark that both of these arc segments are homeomorphic to e. Using e as an intermediary, we see that the map between $(\theta_{(v,a,e)}, \theta_{(\tilde{v},e,\tilde{p})})$ and $(\theta_{(\tilde{v},\tilde{a},e)}, \theta_{(v,e,p)})$ is orientation reversing. Also note that when we paste together all such k pairs of arcs, we recover the original tree.



FIGURE 4. Examples of pseudoaccesses. Even $e = a = p$ may occur.

Let $f : T_1 \to T_2$ be a branched cover of degree d_f. Any 'coordinate system' for the pseudoaccesses of T_2 can be pulled back to the pseudoaccesses of T_1 in such a manner that the map induced by f becomes multiplication by d_f. (In fact, this procedure can be performed in d_f ways.) If, with an abuse of notation, we identify a pseudoaccess with the argument assigned to it, we can deduce several trivial set-theoretic properties. First, when θ can be recognized among the pseudoaccesses of T_1, then $d_f\theta$ is a pseudoaccess in T_2; and, as a consequence, $\theta + j/d_f$, for $j = 0, \ldots, d_f - 1$, are also pseudoaccess in T_1. That $d_f\theta$ is a pseudoaccess in T_2 follows from the construction of the coordinates in T_1. Also, as $d_f\theta$ must have d_f preimages (owing to Lemma 2.8), they can only be the ones suggested above. From the previous discussion we deduce that if θ_1 and θ_2 are consecutive in T_1, then $\theta_1 + j/d_f$ and $\theta_2 + j/d_f$ are also consecutive in the pseudoaccess order of T_1, for all j. Furthermore, in such a situation $d_f\theta_1$ and $d_f\theta_2$ are consecutive in T_2. This essentially reminds us of the fact that the 'edge' (θ_1, θ_2) (in T_1) gets mapped to the 'edge' $(d_f\theta_1, d_f\theta_2)$ (in T_2) according to Corollary 2.5. See also Figure 5.

Now we are in a position to prove Proposition 1.3. In practice, this result turns branched covers between trees into branched covers between embedded trees.

Proof of existence in the analytic case. Thanks to Proposition 2.6, we may assume that $f : T_1 \to T_2$ (of degree $d = d_f$) is a branched cover (as we will be realizing a more restrictive object).

So suppose that the edges of T_2 appear in $\mathbb{C}$ as analytic arcs. Because the edges of T_2 are analytic, the uniformizing Riemann map of the complement of T_2 can be extended

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

inside the unit circle everywhere except for a finite number of points, which we mark. This Riemann map allows us to reconstruct $\mathbb{C}$ by pairwise identification of $2k$ circular arcs, with k equal to the number of edges of T_2. Also, by construction, we can identify several unit elements $e^{2\pi i\theta_1}, \ldots, e^{2\pi i\theta_{2k}}$ with the pseudoaccesses of T_2. Working only with the arguments $\theta_1, \ldots \theta_{2k}$, we have just assigned coordinates to the pseudoaccesses of T_2.

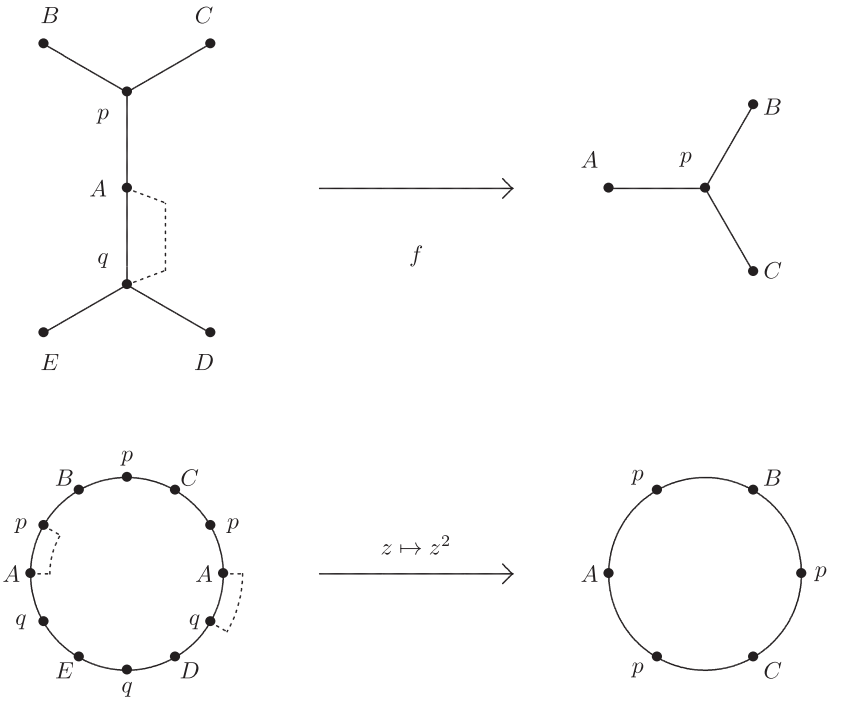


FIGURE 5. The pattern of T_1 can always be read inside the preimage of the pattern of T_2 at the correct angles.

Let us think of $z \mapsto z^d$ as a degree d cover of the complement of the unit disk, construed as a marked surface with boundary. The domain of this cover is an unbounded punctured disk whose boundary is composed of $2kd$ consecutive segments, arcs that start and finish at appropriate preimages of $e^{2\pi i\theta_1}, \ldots, e^{2\pi i\theta_{2k}}$. We make a pairwise identification of arcs between these points using the pattern of T_1 (see the previous discussion), and obtain a new Riemann surface S (where T_1 is embedded). But this identification is compatible with $z \mapsto z^d$, so it induces a degree d map $F : S \to \mathbb{C}$. As this last map is holomorphic and proper, it extends to the respective one-point compactifications, namely $\hat{S}$ and $\hat{\mathbb{C}}$. If we use T_1 and T_2 as part of a triangulation of these surfaces, then the Riemann–Hurwitz formula (see, e.g., [**3**, §7.2]) shows that $\hat{S}$ is a compact Riemann surface whose Euler characteristic is 2, and hence is isomorphic to $\hat{\mathbb{C}}$, the Riemann sphere. With this identification, F is clearly a degree d polynomial. □

Proof of existence in the general case. Once T_2 lies inside $\mathbb{C}$, we consider a homotopic embedding $\tilde{T}_2$ relative to the vertices but with analytical edges. Let F be the polynomial that realizes f with target tree $\tilde{T}_2$. As the critical points of F belong to $F^{-1}(\tilde{T}_2 \cap T_2)$, by

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

construction $F^{-1}(T_2)$ is homotopic to $F^{-1}(\tilde{T}_2)$. Define $T_1 = F^{-1}(T_2)$ and employ F as cover. □

Proof of uniqueness. For simplicity we still denote by T_2 the embedded image of T_2 in $\mathbb{C}$. Let $f : T_1 \to T_2$ be the abstract map we want to realize. We write $\tilde{f} : \tilde{T}_1 \to T_2$ for the homogeneous extension of f. Suppose now that $F^1 : T_1^1 \to T_2$ and $F^2 : T_1^2 \to T_2$ are two realizations of f. Let $\hat{T}_2 \subset \mathbb{C}$ be any tree homotopic to T_2 (modulo vertices) formed by analytic arcs. Since the critical values of both polynomials lie inside T_2, we have that $F^1 : (F^1)^{-1}(\hat{T}_2) \to \hat{T}_2$ and $F^2 : (F^2)^{-1}(\hat{T}_2) \to \hat{T}_2$ are both polynomial realizations of $\tilde{f} : \tilde{T}_1 \to \hat{T}_2$. It is an easy exercise to verify that by picking the correct branch of the inverses, the analytic map $(F^1)^{-1} \circ F^2$ from one complement of '$\tilde{T}_1$' to the other extends continuously even to the trees. As analytic embedded finite trees are obviously removable singularities, the composition $(F^1)^{-1} \circ F^2$ is an entire function. By reversing the roles, we can conclude the same for $(F^2)^{-1} \circ F^1$, its inverse. The only invertible entire functions are affine maps, so we are done. □

3. *The first return map*

Take a cyclic Hubbard forest

$$T_1 \xrightarrow{f_1} T_2 \xrightarrow{f_2} \cdots \xrightarrow{f_{k-1}} T_k \xrightarrow{f_k} T_1$$

of length k. To assume that each f_i is a branched cover presents a minor problem: the T_1 in the left, from now on written as T_1^L, turns out to be much bigger, in general, than the T_1 on the right. However, it is clear that this small conflict can easily be overcome: take T_1^L to be an extension of T_1 that fits into the diagram, and do not forget about the inclusion $T_1 \subset T_1^L$.

To be formal, when we change from the above diagram to

$$\begin{array}{ccccccccc} T_1^L & \xrightarrow{f_1} & T_2 & \xrightarrow{f_2} & \cdots & \xrightarrow{f_{k-1}} & T_k & \xrightarrow{f_k} & T_1, \\ \multicolumn{9}{c}{\underbrace{\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad}_{p_1}} \end{array}$$

we presuppose that some work has been performed in the middle. First, we have to extend T_k to $\bar{T}_k$ and $f_k : T_k \to T_1$ to $\bar{f}_k : \bar{T}_k \to T_1$ so that the latter is now branched (cf. Proposition 2.6). Then we extend T_{k-1} to $\bar{T}_{k-1}$ so that $\bar{f}_k : T_{k-1} \to \bar{T}_k$ is a branched extension of $f_{k-1} : T_{k-1} \to \bar{T}_k$, which in turn is the set-theoretical extension of $f_{k-1} : T_{k-1} \to T_k$. We continue inductively up to when $\bar{f}_1 : \bar{T}_1 \to \bar{T}_2$ is branched and extends $f_1 : T_1 \to T_2$. Finally, we recycle names. The original T_1 will still be called T_1. As $\bar{T}_2, \ldots, \bar{T}_k$ are the extensions of $T_2, \ldots, T_k$, we will need to simplify notation: we shall keep calling them $T_2, \ldots, T_k$; similarly for the $\bar{f}_i$: we will use f_i. We cannot, however, do the same for $\bar{T}_1$, because T_1 stands for the original T_1 while $\bar{T}_1$ is usually bigger. The remedy is to replace T_1 by T_1^L in the diagram.

If for every vertex v in T_1^L we write $d(v) = d_1(v)d_2(v) \cdots d_k(v)$ and $p_1(v) = f_k \circ \cdots \circ f_1(v)$, we get an angled self-map of T_1^L.

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

LEMMA 3.1. *The tree map p_1 on T_1^L is expanding and branched over its image. As a consequence, $p_1 : T_1^L \to T_1^L$ defines a homogeneous Hubbard tree of degree $d_f = d_{f_1} \cdots d_{f_k}$.*

Proof. Let e be an edge in T_1^L with endpoints v and v'. The angled map $f_1 : T_1^L \to T_2$ is a cover, so Corollary 2.5 tells us that $f_1(e)$ is an edge with endpoints $f_1(v)$ and $f_1(v')$ (in T_2). Continuing in the same fashion, we observe that $p_1(e)$ is the edge in T_1 between $p_1(v)$ and $p_1(v')$. In particular, $p_1(v)$ and $p_1(v')$ are different and thus p_1 satisfies the non-degeneracy condition. If $p_1(v)$ and $p_1(v')$ were non-contiguous in T_1^L, then expansiveness between v and v' would be granted for free. Otherwise, assume that for all $n \geq 1$ the vertices $p_1^{\circ n}(v)$ and $p_1^{\circ n}(v')$ are next to each other in T_1^L. Then, by the same argument as above, all iterates of v and v' (and *a fortiori* those of $p_1(v)$ and $p_1(v')$) under the maps f_i would persistently be next to each other. However, by construction, $p_1(v)$ and $p_1(v')$ belong to T_1, one of the original constituent trees of the cyclic Hubbard forest. But a Hubbard forest is expanding, and as such the f_i spread apart adjacent Julia vertices. The only way out is to recognize that at least one of v and v' is of Fatou type.

The homogeneity and the degree of p_1 are handled by counting vertices or edges (see Lemma 2.3). □

As the indices are cyclically arranged, this construction can be started from any tree. In particular, the first return map p_i defines the structure of a Hubbard tree in each T_i^L. As such, the p_i can be realized by postcritically finite polynomials P_i, all of degree $d_f = d_{f_1} \cdots d_{f_k}$ (the degree can easily be calculated by counting vertices). Hence, there are associated filled Julia sets $K_i = K(P_i)$ where we can find the trees T_i naturally embedded; it will turn out that one maps to the next.

Remark. If we restrict p_1 to the original T_1, we still have a Hubbard tree. In fact, T_1 contains the orbit of the critical set. As this Hubbard tree is a restriction of the one introduced in the lemma, by uniqueness it has the same polynomial assigned as before.

Another result of interest is the following.

LEMMA 3.2. *A cyclic Hubbard forest as above can be extended so that T_1 grows into a bigger tree $\hat{T}_1$ if and only if the underlying topological tree of the first return Hubbard tree can be grown into $\hat{T}_1$.*

Proof. If the Hubbard forest can be extended so that T_1 grows into $\hat{T}_1$, then the new first return Hubbard tree has the original first return Hubbard tree as a restriction.

Conversely, write, as usual, $p_1 = f_k \circ f_{k-1} \circ \cdots \circ f_1$ for the first return map to T_1. Consider the set inclusion $T_1 \subset \hat{T}_1$. Applying the existence part of Proposition 2.6 to $f_k : T_k \to \hat{T}_1$, we get a tree $\tilde{T}_k \supset T_k$ and a cover $\tilde{f}_k : \tilde{T}_k \to \hat{T}_1$. Proceeding inductively, we obtain trees $\tilde{T}_{k-1} \supset T_{k-1}, \ldots, \tilde{T}_1 \supset T_1$ and covers $\tilde{f}_{k-1} : \tilde{T}_{k-1} \to \tilde{T}_k, \ldots, \tilde{f}_1 : \tilde{T}_1 \to \tilde{T}_2$ that extend $f_{k-1}, \ldots, f_1$, respectively. The composition of all these maps gives rise to $\tilde{f}_k \circ \cdots \circ \tilde{f}_1 : \tilde{T}_1 \to \hat{T}_1$, a homogeneous cover. On the other hand, by hypothesis $\hat{T}_1$ is the underlying topological tree of the Hubbard tree with structural map p_1. Hence, we can find an extension tree $\hat{T}_1^L \supset \hat{T}_1$ so that p_1 extends to a branched cover $p_1 : \hat{T}_1^L \to \hat{T}_1$ (here there is no need to refer to [**5**]: we just apply Proposition 2.6 once more). The bottom line is

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

that we are left with two branched covers that extend $f_k \circ f_{k-1} \circ \cdots \circ f_1 : T_1 \to T_1$ with target tree $\hat{T}_1$. By Proposition 2.6, the two agree, i.e. we can take $\tilde{T}_1 = \hat{T}_1^L$. As $\hat{T}_1^L$ contains $\hat{T}_1$ by construction, as soon as we restrict $\tilde{f}_1$ to $\hat{T}_1$, then $\tilde{f}_2$ to $\tilde{f}_1(\hat{T}_1)$ and so on, we get the required extension of the cyclic forest. □

4. *Proof of Theorem 1.2 in the cyclic case*

When we realize the Hubbard tree associated with the first return map to T_1^L by a polynomial P_1, among other things we have fixed an embedding of T_1 into a copy $\mathbb{C}_1$ of the complex plane. According to Proposition 1.3, we automatically get an embedding of T_k into $\mathbb{C}_k$ and a polynomial $F_k : \mathbb{C}_k \to \mathbb{C}_1$ that realizes the cover $f_k : T_k \to T_1$. We can continue this process inductively for $i = k-1, \ldots, 2$ and get embeddings of the T_i into, say, $\mathbb{C}_i$ and polynomials F_i that realize all covers $f_i : T_i \to T_{i+1}$. Even for $i = 1$ this can be achieved, but we must be careful: what we get is an embedding of T_1^L and a polynomial F_1 that realizes $f_1 : T_1^L \to T_2$.

On the other hand, we have two realizations of $p_1 : T_1^L \to T_1$. One is given by the Hubbard tree realization P_1, while the other is given by $F_k \circ \cdots \circ F_2 \circ F_1$. According to Proposition 1.3, there is an affine map φ such that $P_1 = F_k \circ \cdots \circ F_2 \circ F_1 \circ \varphi$. In other words, after replacing F_1 by $F_1 \circ \varphi^{-1}$ we can assume directly that $P_1 = F_k \circ \cdots \circ F_2 \circ F_1$ holds and that the embeddings of T_1 agree.

Let F be this cyclic collection of maps. At this point it should be clear that F is postcritically finite. Also, this collection of cyclic maps is the only one that has any chance of realizing the original data f. In fact, the embedding of T_1 is determined by P_1, while the embedding of T_k and the polynomial F_k are, at the same time, determined by the embedding of T_1 and the angled map f_k, and so on. It only remains to prove that the embedded trees are regulated in the sense that their intersection with the interior of the filled Julia set is composed of radial segments in the Böttcher coordinate.

As the filled Julia set is completely invariant, the obvious equality $K(P_1) = K(F) \cap \mathbb{C}_1$ readily implies $K(F) \cap \mathbb{C}_i = F_i^{-1} \circ \cdots \circ F_k^{-1}(K(P_1))$; write K_i for this set. Now let U be a component of the interior of K_i for which $U \cap T_i$ is non-empty. Since T_i maps to T_1 under $F_k \circ \cdots \circ F_i$, there is a Fatou component V of $K(P_1)$ such that $T_1 \cap V = F_k \circ \cdots \circ F_i(T_i \cap U)$. However, T_1 is an actual Hubbard tree, and therefore $T_1 \cap V$ is formed by radial segments in the appropriate coordinate. Because F is postcritically finite, it follows immediately that $T_i \cap U$ is also made up of radial segments. This completes the proof of Theorem 1.2.

5. *An application: anti-Hubbard trees*

We present a straightforward application of our results to postcritically finite anti-holomorphic polynomial maps.

Fix a polynomial F of degree n. Consider the dynamical system determined by $z \mapsto Q(z) = F(\bar{z})$. For this kind of function, the Julia set, filled Julia set, Fatou components and so on are naturally defined due to the fact that $Q \circ Q = \bar{F} \circ F$ is a polynomial map, here $\bar{F}(z) = \overline{F(\bar{z})}$ (cf. [**4**]). Also, as complex conjugation is a homeomorphism, the critical points of Q agree with those of $\overline{Q(z)} = \overline{F(\bar{z})} = \bar{F}(z)$, which we may find by

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

taking derivatives. Hence, these maps can be called *postcritically finite anti-holomorphic polynomials* if the critical points have finite orbits.

In this setup, we imitate the process that led us to Hubbard trees: we pick a finite invariant set M containing all critical points of Q, and inside $K(Q)$ we interconnect M by means of *regulated arcs* (as defined in the introduction) to form a tree $T(M)$. This tree is invariant under Q, and the vertices obey the expanding condition. Therefore, we have a generalization of Hubbard trees, the difference being that we are now dealing with an orientation-reversing map. We call this dynamical system the *anti-Hubbard tree of* Q.

Let us turn briefly to the abstract counterpart. Let $f : T_1 \to T_2$ be a non-degenerate map between angled trees (we refer to the introduction for definitions). If the local degree $d : V_1 \to \{1, \ldots\}$ satisfies $\angle_{f(v)}(f(e), f(e')) = -d(v)\angle_v(e, e')$ for all edges e and e' incident at $v \in V_1$, then the pair (f, d) (or just f for simplicity) is an *orientation-reversing angled (tree) map*.

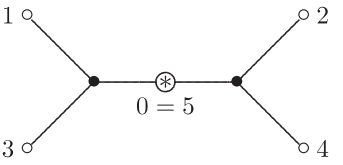


FIGURE 6. A critically periodic degree-two anti-Hubbard tree.

Given an angled tree T, we can define $\overline{T}$, the *flip* of T, as the angled tree whose underlying topological tree is the same graph T but with a new angle structure $\bar{\angle}_v(e, e') = -\angle_v(e, e')$. This procedure reverses the orientation of T but does not tamper with the topology of the tree.

We fix an angled map $f : T_1 \to T_2$. Then its *post-flip* $f^{\sim} : T_1 \to \overline{T}_2$ and its *pre-flip* ${}^{\sim}f : \overline{T}_1 \to T_2$ are defined pointwise by the formula $f^{\sim}(x) = f(x) = {}^{\sim}f(x)$. Even though from the set-theoretical viewpoint they do nothing to f, both maps turn orientation-preserving maps into orientation-reversing maps and vice versa.

The *conjugate map* $\overline{f} : \overline{T}_1 \to \overline{T}_2$, where $\bar{f}(x) = f(x)$, has the same orientation as f in view of the relation $\bar{f} = {}^{\sim}f^{\sim}$.

If $\varphi(T)$ is an embedding of T into $\mathbb{C}$, then $\overline{\varphi(T)}$ (here the bar denotes complex conjugation) is an embedding of $\overline{T}$. This fact has several consequences. Let $F : \mathbb{C} \to \mathbb{C}$ be a holomorphic realization of $f : T_1 \to T_2$. Then the post-flip $f^{\sim}$ can be realized by $\overline{F(z)}$, the pre-flip ${}^{\sim}f$ by $F(\bar{z})$ and the conjugate $\overline{f}$ by $\overline{F}(z) = \overline{F(\bar{z})}$. These facts are trivial, and the details of their proof are left to the reader. Other useful relations are $f \circ f = {}^{\sim}f \circ f^{\sim}$ and $\overline{f \circ f} = \bar{f} \circ \bar{f} = f^{\sim} \circ {}^{\sim}f$.

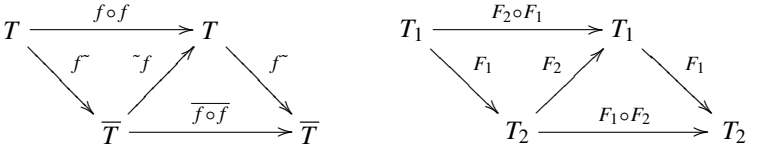
An orientation-reversing self-map on an angled tree is *normalized* if whenever m edges $e_1, \ldots, e_m$ meet at a periodic Julia vertex v, each $\angle_v(e_i, e_j)$ is a multiple of $1/m$. An *(abstract) anti-Hubbard tree* is a normalized orientation-reversing angled tree map that satisfies the expanding condition for all adjacent periodic Julia vertices.

The criterion for realization of postcritically finite anti-holomorphic polynomial dynamics is given next.

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

THEOREM 5.1. *A normalized orientation-reversing angled tree dynamics can be realized as the dynamics associated to a postcritically finite anti-holomorphic polynomial map if and only if it is expanding. In other words, all anti-Hubbard trees—and only them—can be realized. Such a realization is unique up to affine conjugation in the dynamical plane.*

Proof. Let $f : T \to T$ be the anti-Hubbard tree we want to realize. Define $T_1 = T$ and $T_2 = \overline{T}$. Then the two maps $f^{\sim} : T_1 \mapsto T_2$ and ${}^{\sim}f : T_2 \to T_1$ give rise to a cyclic Hubbard forest with two components. By Theorem 1.2, they can be realized by a pair of polynomials F_1, F_2. The angled graphs maps and the actual polynomial realizations fit into the following commutative diagrams:



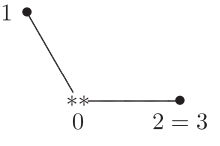
(here we are keeping the names T_1 and T_2 for the embedded trees in the respective copy of the complex numbers). When we fix the copy of $\mathbb{C}$ in which T_1 is embedded (that is, the domain of F_1), everything is uniquely determined up to an affine change of coordinates in the intermediate copy of $\mathbb{C}$ where T_2 is embedded.

First, we find a linear change of coordinates φ so that $\overline{\varphi(T_2)}$ agrees with T_1. In fact, we have in sight two realizations of the angled map $\overline{f \circ f}$, namely $F_1 \circ F_2 : T_2 \to T_2$, already exhibited in the diagram above, and $\overline{F_2 \circ F_1} : \overline{T_1} \to \overline{T_1}$. The uniqueness part of Proposition 1.3 provides us with an affine linear transformation φ subject to $\varphi(T_2) = \overline{T_1}$ (or, alternatively, $T_1 = \overline{\varphi(T_2)}$) and $\varphi \circ F_1 \circ F_2 \circ \varphi^{-1} = \overline{F_2 \circ F_1}$.

Notice that the anti-polynomial defined by the formula $Q(z) = \overline{\varphi(F_1(z))}$ maps T_1 to itself and, in fact, realizes the anti-Hubbard tree $f : T \to T$. (Of course, the same is true for $\hat{Q}(z) = F_2(\varphi^{-1}(\bar{z}))$.) This settles the existence.

Uniqueness is easily handled. Our preparatory work with the auxiliary Hubbard forest implies that there is, modulo coordinates, a unique way to embed T_1. If T has $k > \deg(f)$ vertices, this implies uniqueness of the realization of f, since an anti-polynomial of degree $\deg(f)$ is determined by $\deg(f) + 1$ values. Otherwise, using Lemma 3.2, we extend the auxiliary Hubbard forest so that it now carries more than $\deg(f)$ vertices; then we proceed as before. □

Example. The dynamical tree map



of degree three can be interpreted both as a Hubbard tree and as an anti-Hubbard tree. (Here 0 is the only critical point, of local degree three.) However, the realizations in the two cases look quite different: as a Hubbard tree, the map is realized by $P(z) = z^3 - 0.340625019 - 1.271229878i$ (with postcritical fixed point at $1.27123 + 0.340625i$);

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press

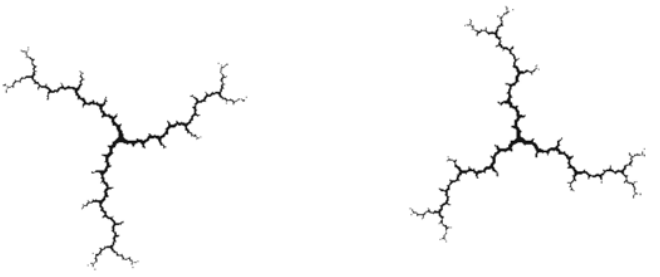


FIGURE 7. The Hubbard tree is shown on the left and the anti-Hubbard tree on the right.

as an anti-Hubbard tree, it is realized by $Q(z) = \bar{z}^3 - 0.801175097 - 1.044111714i$ (with postcritical fixed point at $1.30481 - 0.171782i$). Their Julia sets are shown in Figure 7.

REFERENCES

[1] A. Douady. Descriptions of compact sets in $\mathbb{C}$. *Topological Methods in Modern Mathematics: A Symposium in Honor of John Milnor's Sixtieth Birthday*. Publish or Perish, Inc., Houston, 1993, pp. 429–465.

[2] A. Douady and J. Hubbard. Étude dynamique des polynômes complexes, I, II. *Publ. Math. Orsay* (1984–1985).

[3] J. Milnor. *Dynamics in One Complex Variable: Introductory Lectures*. Vieweg, Wiesbaden, 2000.

[4] S. Nakane and D. Schleicher. On multicorns and unicorns I: antiholomorphic dynamics, hyperbolic components and real cubic polynomials. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.* **13**(10) (2003), 2825–2844.

[5] A. Poirier. Hubbard trees. *Fund. Math.* **208**(3) (2010), 193–248.

[6] D. Sullivan. Quasiconformal homeomorphisms and dynamics I: solution to the Fatou–Julia problem on wandering domains. *Ann. of Math.* (2) **122** (1985), 401–418.

https://doi.org/10.1017/S0143385711000733 Published online by Cambridge University Press