NOTES ON THE DYNAMICS OF LINEAR OPERATORS

JOEL H. SHAPIRO

Abstract. These are notes for a series of lectures that illustrate how continuous linear transformations of infinite dimensional topological vector spaces can have interesting dynamical properties, the study of which forges new links between the theories of dynamical systems, linear operators, and analytic functions.

1. Introduction to transitivity

Our story begins in a separable complete metric space X on which acts a (not always continuous) mapping T. We are interested in the behavior of the sequence

$$I,\ T,\ T^2,\ T^3, \cdots$$

of *iterates* of T, where T^n denotes the composition of T with itself n times, and we will be particularly interested in knowing when there exists a point $x \in X$ for which the T-*orbit*

$$\operatorname{orb}(T, x) = \operatorname{orb}(x) := \{x, Tx, T^2x\ \ldots\}$$

of x under T is *dense* in X. When this happens we call T *(topologically) transitive* and refer to x as a *transitive point* for T. Before long, we will be considering T to be a linear transformation on a metrizable topological vector space, in which case the word "hypercyclic" will be used instead of "transitive" (more on the reason for this terminology shift later).

Note that if x is a transitive point for T then the same is true of every point in $\operatorname{orb}(T, x)$, so once there is one transitive point there is a dense set of them.

Even for simple metric spaces like the unit interval or the unit circle, the study of transitive maps can be fascinating. We begin with some elementary examples that illustrate this point.

1.1. **Irrational translation modulo one.** The metric space here is the closed unit interval $I = [0, 1]$. Fix an irrational real number α and define and the mapping $T : I \to I$ by $Tx = x + \alpha$ (mod 1). That is, Tx is what's left when you subtract off the integer part of $x + \alpha$. Now T is not a continuous map of I, but if this causes you problems, then identify

Date: May 27, 2001.

the endpoints 0 and 1 to make I into a circle, and regard T to be rotation of the this circle through the angel $2\pi\alpha$, a continuous map. More precisely, let $E(x) = e^{2\pi i x}$, which maps I continuously onto the unit circle $\mathbb{T}$. Let $R : \mathbb{T} \to \mathbb{T}$ be the mapping of rotation through the angle $2\pi\alpha$, i.e., $Tz = e^{2\pi i\alpha}z$. Then R is a continuous mapping of T onto itself, and $E \circ T = R \circ E$, which, because E is "almost one-to-one" exhibits the action of T on I as being "essentially the same as" that of R on $\mathbb{T}$. In particular, to show that T is transitive it's enough to do this for R. In fact we'll do more:

1.2. **Proposition.** *Every orbit of R (hence every orbit of T) is dense.*

Proof. Let $\omega := e^{2\pi i\alpha}$ and note that the irrationality of α guarantees that ω is not a root of unity. Thus

$$\operatorname{orb}(R, 1) = \{1,\ \omega,\ \omega^2,\ \ldots\}$$

is an infinite subset of $\mathbb{T}$, and therefore (because $\mathbb{T}$ is compact) has a limit point in $\mathbb{T}$. In particular, there is a strictly increasing sequence $\{n_k\}$ of positive integers such that

$$0 = \lim_k |\omega^{n_{k+1}} - \omega^{n_k}| = \lim_k |\omega^{n_{k+1}-n_k} - 1| = \lim_k |R^{n_{k+1}-n_k}(1) - 1|,$$

hence 1 is a limit point of its own R-orbit. Thus, given $\varepsilon > 0$ there is a point $\zeta \in \operatorname{orb}(R, 1)$ such that the arc between 1 and ζ has length $< \varepsilon$. The successive powers of ζ therefore partition the unit circle into non-overlapping arcs of length $< \varepsilon$, hence every point of the circle lies within ε of one of these powers, all of which belong to $\operatorname{orb}(R, 1)$. Thus $\operatorname{orb}(R, 1)$ is dense in $\mathbb{T}$, establishing the transitivity of R.

To see that *every* R-orbit is dense, it's enough to notice that for $\zeta \in \mathbb{T}$:

$$\operatorname{orb}(R, \zeta) = \zeta \cdot \operatorname{orb}(R, 1),$$

and to observe that the map $z \to \zeta z$ is a homeomorphism of $\mathbb{T}$, which therefore takes dense sets to dense sets and in particular transfers the density of $\operatorname{orb}(R, 1)$ to $\operatorname{orb}(R, \zeta)$. □

Why transitivity?

Let's return to the general situation where X is a metric space on which acts a map T, and let's say that a subset E of X is *T-invariant* if $T(E) \subset E$.[1] For $x \in X$ the closure of $\operatorname{orb}(T, x)$ is the smallest closed T-invariant subset of X that contains x. Thus for maps like irrational translation mod one (or, equivalently, rotation of the unit circle through an

[1] *Warning:* This terminology is standard in operator theory, but not universal.

angle that is an irrational multiple of π), density of every orbit is equivalent to the fact that the only closed invariant sets are the empty set and the whole space. Such maps are called *minimal*.

The more modest concept of transitivity merely asserts that open invariant sets are either empty or dense. Perhaps it's time to give an example of a transitive mapping that is not minimal.

1.3. **The baker map.** This is the map $B : I \to I$ defined by $Bx = 2x \pmod 1$. The name comes from a strategy for kneading dough: take a strip of unit length, roll it out to double the length, cut it in half, translate the right half on top of the left half, knead the two halves together into one strip of unit length, and repeat the process.

Note that the baker map is *not* minimal—the origin is a fixed point! More generally, for each dyadic rational $x = q/2^n$ (q and n positive integers with $q \le 2^n$), $B^m x = 0$ for all $m \ge n$, so there is a dense set of points with finite orbits. Nevertheless:

1.4. **Proposition.** *The baker map is transitive.*

Proof. Represent each point x of I by a binary expansion, i.e., $x = .a_1a_2\dots$ ($a_i = 0$ or 1) means that $x = \sum_{k=1}^{\infty} a_k/2^k$. Binary rationals $q/2^n$ have two such expansions: one that is finitely nonzero and the other that is not. For these, choose the finite one. The binary expansion of every other point of I is unique. For x represented as above, $Bx = .a_2a_3\dots$, i.e., B performs a backward shift on binary expansions. Enumerate the countable dense set of binary rationals in I as $\{b_1, b_2, \dots\}$. Form the point $x \in I$ as follows: begin with the (finite, recall) binary expansion of b_1, follow it with a zero, after which you copy in the expansion of b_1 again, then two zeros, then the expansion of b_2, then three zeros, then back to b_1, four zeros, b_2, five zeros, b_3, six zeros, etc. The idea is to get the finite expansion of each binary rational copied into that of x infinitely often, each time followed by successively more zeros. This produces a point $x \in I$ with each binary rational a limit point of $\operatorname{orb}(B, x)$, hence x is a transitive point for B. $\square$

Characterizing transitivity. Suppose the mapping T of the separable metric space X is transitive. Let $\operatorname{trans}(T)$ denote the (necessarily dense) set of transitive points of T. To better describe $\operatorname{trans}(T)$, fix a countable basis $\{B_j\}$ of open sets for the topology of X. For example, this could be the collection of open balls of rational radius with centers in some fixed countable dense subset of X. Then it's easy to check that

$$\operatorname{trans}(T) = \bigcap_j \bigcup_n T^{-n}(B_j). \tag{1}$$

Since $\operatorname{trans}(T)$ is dense in X, so is the union on the right for each positive integer j, and this means that for every basic open set B_j and every nonempty open set V there is a non-negative integer n such that $T^{-n}(B_j) \cap V \neq \emptyset$, or equivalently, $B_j \cap T^n(V) \neq \emptyset$. Now every nonempty open subset of X contains a B_j, so we have just proved:

1.5. **Proposition.** *If T is a transitive map of a (necessarily separable) metric space X, then for every pair U, V of nonempty open subsets of X there is a non-negative integer n such that $T^{-n}(U) \cap V \neq \emptyset$, or equivalently, $U \cap T^n(V) \neq \emptyset$.*

Corollary. *If a map is transitive then the orbit of every nonvoid open set is dense.*

In case T is a *continuous* transitive map on X, the description (1) shows that *the set of transitive points of T is a dense G_δ.* Suppose, in addition, that X is complete. then Baire's Theorem asserts that the intersection of every countable collection of dense open sets is again dense, so every countable intersection of dense G_δ's is another dense G_δ. In summary:

1.6. **Proposition.** *Every transitive continuous mapping of a complete metric space X has a dense G_δ set of transitive points. Every countable collection of such maps on X has a dense G_δ set of common transitive points.*

If a complete metric space has no isolated points, then every dense G_δ set is uncountable (see [32, §5.13, page 103], for example), so continuous mappings of such spaces are either non-transitive, or have an uncountable dense set of transitive points.

It might appear that these remarks do not apply to the baker map, which is not continuous on I, but by the same trick used for irrational translation mod one, we also can think of B as acting on the unit circle. Indeed, letting $E : I \to \mathbb{T}$ denote the exponential map of §1.1 ($E(x) = e^{2\pi i x}$), we have $E \circ B = E^2$, which allows the transitivity of B to be transferred to the "squaring map" $z \to z^2$ of $\mathbb{T}$. Since the squaring map is continuous, it has a dense G_δ (and therefore uncountable) set of transitive points, and from this it's easy to see that this property transfers back to the (non-continuous) baker map.

We close this section with one more consequence of the characterization (1) of the set of transitive points, where now T is continuous and the separable metric space X is complete. Suppose that such a T obeys the conclusion of Proposition 1 above, i.e. that for each pair U, V of nonempty open subsets of U there is a non-negative integer n such that $T^{-n}(U) \cap V \neq \emptyset$. Then the union on the right-hand side of (1) is a G_δ that intersects every nonempty open set, i.e. it is a dense G_δ. The set of transitive points of T is therefore a countable intersection of dense G_δ's, so it is nonempty by Baire's Theorem [32, §5.6, page 97]. Thus for continuous maps of complete, separable metric spaces the converse of of Proposition 1 holds. These observations form the content of:

1.7. **Birkhoff's Transitivity Theorem.** *A continuous map T of a complete, separable metric space X is transitive if and only if for every pair U, V of nonempty open subsets of X there is a non-negative integer n such that $T^{-n}(U) \cap V \neq \emptyset$ (or equivalently: $U \cap T^n(V) \neq \emptyset$).*

Corollary. *A homeomorphism of a complete, separable metric space onto itself is transitive if and only if its inverse is transitive.*

So far our only invertible example is translation modulo one of the unit interval by an irrational number. In this case transitivity of the inverse is clear: it is just translation modulo one by the negative of that irrational number. Later we will encounter many other interesting examples of invertible transitive maps.

Transference of transitivity

We have used the exponential map $E : x \to e^{2\pi i x}$ to connect the behavior of some discontinuous maps T of the unit interval (irrational translation modulo one and the baker map) with continuous maps R of the unit circle (irrational rotation and the squaring map, respectively) via the "intertwining equation" $E \circ T = S \circ E$. This is a special case of something quite general:

1.8. **Definition.** *Suppose $T : X \to X$ and $S : Y \to Y$ are mappings of metric spaces, and $E : X \to Y$ is a continuous map of X onto Y for which $V \circ T = S \circ V$. In this case we call S a* factor *of T, and T an* extension *of S. If $V(X)$ is just dense in Y we'll say S is a* quasi-factor *of S, and T a* quasi-extension *of T. If V is a homeomorphism of X onto Y we say S and T are* conjugate.

Clearly conjugacy is an equivalence relation. Going back to our examples: Rotation of the circle through an irrational multiple of π is both a factor of irrational translation mod one, and a quasi extension of that map (via the "quasi-conjugacy" $V(e^{2\pi ix}) = x$, $0 \leq x < 1$ which maps $\mathbb{T}$ continuously onto the dense subset $[0,1)$ of the closed unit interval). Similarly, the squaring map on the circle is both a factor and a quasi-extension of the baker map of the unit interval.

1.9. **Proposition.** *If $T : X \to X$ is transitive then so is every quasi-factor of T.*

Proof. Suppose $V : Y \to Y$ is a continuous map of the metric space Y onto a dense subset, and $V \circ T = S \circ V$. We establish the desired result by showing that $V(\text{trans}\,(T)) \subset \text{trans}\,(S)$.

To this end suppose that x is a transitive point of T. An induction shows that $V \circ T^n = S^n \circ V$ for every non-negative integer n, hence $V(\text{orb}\,(T,x)) = \text{orb}\,(S, Vx)$. Since V is continuous with image dense in Y, the image of any dense subset of X is dense in Y, in particular this is true of the S-orbit of Vx, which was just revealed as the image under S of the dense T-orbit of x. □

1.10. **A shift map.** In proving transitivity for the baker map B, we used the fact that it acted as a shift on binary expansions of numbers in the unit interval. In fact, what we were really doing was proving transitivity for a certain extension of B. Let Σ be the space of all sequences of zeros and ones, and let β denote the *backward shift* on Σ, that is, if $x = (x(n) : 1 \leq n < \infty) \in \Sigma$, then

$$\beta x = (x(2), x(3), \dots).$$

The metric d defined on X by

$$d(x,y) = \sum_{n=1}^{\infty} 2^{-j}|x(j) - y(j)| \quad (x, y \in \Sigma)$$

is complete on σ, and d-convergence is the same as coordinatewise convergence. Thus the map $V : \Sigma \to I$ defined by

$$Vx = \sum_{n=1}^{\infty} 2^{-j} x(j) \quad (x \in \Sigma)$$

is continuous, with $V(\Sigma) = I$. Clearly $V \circ \beta = B \circ V$, which establishes the claim that the baker map B is a factor of β. In our proof of transitivity for β what we really did was construct a transitive vector for β, and interpret it, via V, as a point of I.

Here are two other well-known maps of the unit interval. I leave it mostly as exercises to show that they are conjugate to each other, and that both are factors of the squaring map.

1.11. **The quadratic map.** This is the mapping $Q : I \to I$ defined by $Q(x) = 4x(1-x)$.

Exercise. *Show that Q is a factor of the squaring map.*

Outline of proof. Define $V_1 : \mathbb{T} \to [-1,1]$ by $V_1(e^{ix}) = \mathrm{Re}\,(e^{ix}) = \cos x$. So V is the orthogonal projection of $\mathbb{T}$ onto $[-1,1]$; except for the points ± 1 it is two-to-one. Some common trigonometric identities show that if $z \in \mathbb{T}$ then $V_1(Qz) = 2V_1(z)^2 - 1$, so if σ denotes the squaring map of the circle and $Q_1(x) = 2x^2 - 1$ on $[-1,1]$, then $V_1 \circ \sigma = Q_1 \circ V_1$. Thus Q_1 is a factor of σ, and so inherits its transitivity.

Next, map $[-1,1]$ homeomorphically onto $I = [0,1]$ in the simplest way: $V_2(x) = (1+x)/2$. The map $Q_2 := V_2 \circ Q_1 \circ V_2^{-1}$ is therefore conjugate to Q_1, hence is a factor of σ. A little arithmetic shows that $Q_2(x) = (2x-1)^2$. Finally the homeomorphism $V_3(x) = 1 - x$ of I onto itself yields $V_3 \circ Q_2 \circ V_3^{-1} = Q$, hence Q is conjugate to Q_2, and so is a factor of σ, hence a transitive map.

Upon composing the maps V_j properly, you can see that the map $V : \mathbb{T} \to I$ that comes out of all this and exhibits the quadratic map as a factor of the squaring map ($V \circ \sigma = Q \circ V$) is: $V(e^{ix}) = \sin^2 x$.

1.12. **The Tent Map.** This is the map $T : I \to I$ defined by: $T(x) = 2x$ if $0 \le x \le 1/2$, and $T(x) = 2(1-x)$ if $1/2 \le x \le 1$.

The name comes from the shape of the graph, which is an inverted "V" based on I. The shape of the graph is qualitatively like that of the quadratic map, so one would conjecture that the two maps are conjugate. This would imply transitivity for the tent map, also a reasonable conjecture if one interprets the map's action in terms of mixing dough (it's just like the baker map, except that instead of cutting and translating the stretched out dough, you just double it over at its midpoint).

1.13. **Exercise.** *Show that if $V(x) = \sin^2(\frac{\pi}{2}x)$, so that V is a homeomorphism of I onto itself, then $V \circ T = Q \circ V$, hence the tent and quadratic maps are conjugate.*

Chaotic maps

The following discussion comes almost entirely from [3], a beautiful short paper which discusses the question: *What does it mean for a map to be "chaotic"?* There seems to be no consensus on this issue, but most authors seem to agree that chaotic maps should at least be transitive. However transitivity by itself does not seem to capture the essence of chaos. Consider for example the map of irrational translation mod one, acting on the unit interval. Although it is transitive, orbits that start close together stay close together, so at least in this respect the map is too regular to be considered truly chaotic.

To exclude such examples it seems desirable to also require a form of instability which asserts that each point $x \in X$ should have points arbitrarily close to it whose orbits, in some uniform sense, do not stay close to $\operatorname{orb}(T, x)$. More precisely:

1.14. **Definition.** *A mapping T of a metric space X* depends sensitively on initial conditions *(or:* has sensitive dependence*) if: There exists a number $\delta > 0$ such that for every $\varepsilon > 0$ and every $x \in X$ there is a point $y \in B(x, \varepsilon)$ such that $d(T^n x, T^n y) > \delta$ for some non-negative integer n.*

The number δ in this definition is is called a *sensitivity constant* for T.

Some authors take "chaotic" to mean "transitive plus sensitively dependent on initial conditions" (see, for example, [30]). However sensitive dependence has a flaw: *it is not, in general, preserved by conjugacy.* To see this, let T be the map of "multiplication by two" on the positive real axis, S the map of "translation by $\ln 2$" on $\mathbb{R}$, and $V : (0, \infty) \to \mathbb{R}$ the natural logarithm. Then $V \circ T = S \circ V$ so T and S are conjugate, but T does not have sensitive dependence, whereas S does.[2]

In his classic text [15], Devaney proposed a third property that chaotic maps should have: *dense sets of periodic points.* A point $x \in X$ is *periodic* for T if there is a *positive* integer n such that $T^n x = x$. The least such n is called the *period* of x; in particular, fixed points are precisely those of period one.

[2]However, as pointed out in [3], it is not difficult to show that if a map T of a *compact* metric space X has sensitive dependence, then so does any map conjugate to T. Indeed, suppose $V \circ T = S \circ V$, where $V : X \to Y$ is a homeomorphism onto Y. Let δ be a sensitivity constant for T, and let $D(\delta)$ denote the set of pairs $(x, x') \in X \times X$ with $d(x, x') \geq \delta$. This is a compact subset of $X \times X$, so its "V-image" is a compact subset of $Y \times Y$ that is disjoint from the diagonal of that space. Since this diagonal is also compact, these two compact sets lie some positive distance δ_Y apart, and one checks easily that δ_Y is a sensitivity constant for S.

Let per (T) denote the set of all the periodic points of T. The same argument that showed quasi-factors inherit transitivity now shows that they also inherit denseness of periodic points. Indeed, if $V \circ T = S \circ V$, then $V(\text{per}\,(T)) \subset \text{per}\,(S)$. Thus if V is continuous and has dense range, and per (T) is dense, then so is per (S). In particular, the property of having a dense set of periodic points is preserved by conjugacy.

Reasoning that chaotic maps should have pervasive elements of both randomness and predictability, Devaney defined a map to be "chaotic" if it has all three of these properties: transitivity, sensitive dependence, and a dense set of periodic points [15, Page 52]. It might appear that Devaney's definition, relying as it does on sensitive dependence, is not in general preserved by conjugacy, but in [3], Banks et. al. showed otherwise:

1.15. **Theorem.** *Suppose a continuous mapping T of a metric space X is transitive and also has a dense set of periodic points. Then T depends sensitively on initial conditions.*

Proof. The first step is to observe that each point of x lies uniformly far from some periodic orbit. More precisely:

> *There exists $\delta_0 > 0$ such that for each $x \in X$ there is a periodic point q with* $\text{dist}\,(x, \text{orb}\,(q)) > \delta$.

To see why this is so, note first that two periodic points for T either have the same orbit or have disjoint orbits. Since per (T) is assumed dense in X, there must be two periodic points with disjoint orbits. Let $2\delta_0$ be the distance between these orbits. Then any point of x has to lie at distance $\geq \delta_0$ away from at least one of these orbits.

I claim now that $\delta = \delta_0/4$ is a sensitivity constant for T.

To prove this, fix $x \in X$ and $0 < \varepsilon < \delta$. Our goal is to show that there is a point $y \in B(x, \varepsilon)$ and a positive integer m such that $d(T^m x, T^m y) > \delta$.

Since T is transitive, it has a dense set of transitive points, so one of these, call it t lies in $B(x, \varepsilon)$. If $d(T^m t, T^m x) > \delta$ for some m, we are done, with $y = t$. So assume that each point of orb (t) lies within δ of the corresponding point of orb (x), i.e. that the orbit of the transitive point t is "dragging the orbit of x through the space." In this case we'll see that for any periodic point $p \in B(x, \varepsilon)$, some point of orb (x) is more than δ away from the corresponding point of orb (p), and this will complete the proof (with $y = p$).

So fix such a periodic point p and denote its period by n. Note that the distance from p to orb (q) is $\geq 4\delta - \delta = 3\delta$.

Because T is continuous, there is a ball U centered at q with $T^j(U) \subset B(T^j q, \delta)$ for $j = 0, 1, 2, \ldots n-1$. Because t is a transitive point for T there is a positive integer k such that $T^k t \in U$, hence for $0 \leq j < n$ we have $T^{k+j}t$ within δ of orb (q), and so $T^{k+j}x$ lies within 2δ of orb (q). Now there's a unique integer j between 0 and $n-1$ such that $k+j$ is a multiple of n. Fix this j and set $m = k+j$. We've already observed that $T^m p = p$ lies more than 3δ distant from orb (q), and have just seen that $T^m x$ is at most 2δ distant from that same orbit. Thus $d(T^m p, T^m x) \geq \operatorname{dist}(T^m p, \operatorname{orb}(q)) - \operatorname{dist}(T^m x, \operatorname{orb}(q)) > 3\delta - 2\delta = \delta$, as desired. $\square$

1.16. **Definition.** *We say a mapping T of a metric space X is* chaotic *if it is transitive and has a dense set of periodic points.*

Thus chaotic maps have sensitive dependence, and chaotic-ness is preserved by conjugacy; indeed, it is inherited by quasi-factors.

Let's try this concept out on our examples. We've already seen that irrational translation modulo one (and also its factor, irrational rotation of the circle), although transitive, is not sensitively dependent, hence not chaotic. Of the remaining maps (all of which are transitive), recall that the tent map is conjugate to the quadratic map, which is a factor of the squaring map, which is a factor of the baker map, which is a factor of the backward shift β on the metric space Σ of all sequences of zeros and ones. So if we can prove β is chaotic, we'll know all of these factors are chaotic, too.

We already know β is transitive, so it's enough to find a dense set of periodic points. But this is obvious: the periodic points of β are just the periodic sequences of zeros and ones (these correspond to binary expansions of binary rationals in the unit interval), and any sequence $x \in \Sigma$ is approximated to within 2^{-n} by the periodic sequence you get by repeating *ad infinitum* the first n coordinates of x. Thus we have proved:

1.17. **Theorem.** *The backward shift, the squaring map, the tent map, and the quadratic map are all chaotic.*

2. Hypercyclicity: Basic examples

In this section we consider complete metric spaces that are also vector spaces over the complex numbers, and for which the vector operations are jointly continuous. That is, vector addition, viewed as a map $X \times X \to X$, and scalar multiplication, viewed as a map $\mathbb{C} \times X \to X$, are both continuous. Such spaces are called *F-spaces.* The most common examples of F-spaces are Hilbert spaces, and more generally Banach spaces. However there are others, for example:

(a) The Lebesgue spaces $L^p(\mu)$ where μ is a measure and $0 < p < 1$ are F-spaces, with the metric defined by

$$d(f,g) = \int |f-g|^p \, d\mu \qquad (f,g \in L^p(\mu)).$$

(b) If G is an open subset of the plane, then the space $C(G)$ of continuous, complex-valued functions on G can be metrized so that a sequence convergence in this metric if and only if it converges uniformly on compact subsets of G. The resulting space is thus an F-space. One way to obtain such a metric is to exhaust G by an increasing sequence of open subsets $\{G_n\}$, where the closure of each G_n is compact, and contained in G_{n+1}. Then let $d_n(f,g)$ be the supremum of $|f(z) - g(z)|$ as z ranges over G_n (finite because the closure of G_n is a compact subset of G), and set

$$d(f,g) = \sum_n 2^{-n} \frac{d_n(f,g)}{1 + d_n(f,g)} \qquad (f,g \in C(G)).$$

(c) The collection $H(G)$ of functions that are holomorphic on G is a closed subspace of $C(G)$, and therefore an F-space in its own right. In fact it was in just such a space that hypercyclicity (a.k.a. transitivity) was first observed for a linear operator. The result, due to G. D. Birkhoff, dates back to 1929, and asserts that for each $a \in \mathbb{C}$ the operator $T_a : H(\mathbb{C}) \to H(\mathbb{C})$ of "translation by a," defined by:

$$T_a f(z) = f(z+a) \qquad (z \in \mathbb{C},\ f \in H(\mathbb{C})) \tag{2}$$

is hypercyclic. (Note that T_a is invertible with inverse T_{-a}, which makes good on an earlier promise to provide more examples of transitive, invertible maps.)

The first example of hypercyclicity for an operator on a Banach space was exhibited by Rolewicz in 1969. The setting is the sequence space ℓ^p for $1 \le p < \infty$, and the operator is constructed from the backward shift B defined on ℓ^p in exactly the same way it was defined on the sequence space Σ in the last section:

$$Bx = (x(2), x(3), \dots) \qquad \text{where } x = (x(1), x(2), \dots) \in \ell^p. \tag{3}$$

B itself is a contraction on ℓ^p ($\|Bx\| \leq \|x\|$ for each $x \in \ell^p$), so there's no hope of its being hypercyclic, however Rolewicz proved that if you multiply B by any scalar of modulus > 1 the resulting operator *is* hypercyclic.

We begin with a sufficient condition that provides a unified proof of hypercyclicity for both these operators, and for many others as well. This result was discovered by Carol Kitai in her Toronto thesis [22], but she never published it, and it was rediscovered later by Gethner and Shapiro [17]. Don't be fooled by the seemingly complicated statement; as we'll see shortly, its proof is easy, and the result is often quite easy to use!

2.1. **Sufficient condition for hypercyclicity.** *Suppose a continuous linear transformation T on an F-space X satisfies these conditions:*

(a) *There exists a dense subset Y of X on which $\{T^n\}$ converges to zero pointwise.*

(b) *There exists a dense subset Z of X and a mapping $S : Z \to X$ (not necessarily either continuous or linear) such that:*

   (i) *TS is the identity map on Z*

   (ii) *$\{S^n\}$ converges to zero pointwise on Z.*

Then T is hypercyclic on X.

Proof. Fix two nonempty open subsets U and V of X. Using the density of Y and Z, choose $y \in U \cap Y$ and $z \in V \cap Z$. Then $T^n y \to 0$ and $S^n z \to 0$. Thus $x_n := y + S^n z \to y$, hence $x_n \in U$ for all sufficiently large n.

Now even though S and T need not commute, the fact that $TS = I$ on Z means that $T^n S^n = I$ on Z also. Thus by the linearity of T (used for the first and only time here), $T^n x_n = T^n y_n + z \to z$, hence $T^n x \in V$ for all sufficiently large n. Thus $T^n(U) \cap V \neq \emptyset$ for all sufficiently large n, which by Birkhoff's Transitivity Theorem (Theorem 1.7) is more than enough to imply that T is transitive. □ □

For our first application, recall the backward shift B defined on ℓ^p by (3).

2.2. **Rolewicz's Theorem** [31]**.** *For every scalar λ of modulus > 1 the operator λB is hypercyclic on ℓ^p for each $0 < p < \infty$.*

Proof. Fix a scalar λ with $|\lambda| = 1$. We will apply Theorem 2.1 to $T = \lambda B$. To this end let U denote the forward shift on ℓ^p:

$$Ux = (0, x(1), x(2), \dots) \qquad \text{where } x = (x(1), x(2), \dots) \in \ell^p.$$

and set $S = \lambda^{-1}U$. Since $BU = I$ on ℓ^p, we also have $TS = I$. Let $Z = \ell^p$ and note that $\|S^n x\| = |\lambda|^{-n}\|x\| \to 0$ as $n \to \infty$. let Y be the collection of finitely nonzero sequences in ℓ^p; a dense subspace of ℓ^p because $p < \infty$. Then for each $x \in Y$ we have $B^n x$ eventually zero, so the same is true of $S^n x$, hence the hypothesis of Theorem 2.1 are fulfilled. That's all there is to it! $\square$

We can think of λB as a special *weighted* backward shift. If $w = \{w_n\}$ is a bounded sequence of non-zero complex numbers, define the operator B_w on ℓ^p by:

$$B_w(x) = (x(1)w_1,\ x(2)w_2,\ \dots) \qquad (x = (x(1),\ x(2),\ \dots) \in \ell^p). \tag{4}$$

Then B_w is a continuous linear operator on ℓ^p, which converges pointwise to zero on the set Y of finitely non-zero sequences employed in the previous proof.

Suppose further that $\limsup_n |w_n| > 1$. Since none of the weights w_n are zero, we can define the right inverse operator S_w on $Z = \ell^p$ by: $S_w x = (0,\ x(1)/w(1),\ x(2)/w(2),\ \dots)$. Once again, all hypotheses of Theorem 2.1 are satisfied, so B_w is hypercyclic without any help from scalar multiplication. To summarize:

2.3. **Theorem** [17, §4]**.** *If w is a bounded sequence of complex numbers with no element zero, and $\limsup_n |w_n| > 1$, then the weighted backward shift B_w is hypercyclic on ℓ^p for every $0 < p < \infty$.*

Let's turn now to some examples from analytic function theory. We begin with a natural complement to Birkhoff's translation theorem that was, surprisingly, not proved until much later:

2.4. **MacLane's Differentiation Theorem** [25, 1952]**.** *The operator of differentiation is hypercyclic on $H(\mathbb{C})$.*

Proof. Let $Y = Z$ denote the dense subspace of all polynomials in $H(\mathbb{C})$. Let D denote the differentiation operator and S the operator of integration from 0 to z;

$$Sp(z) = \sum_{n=0}^{N} \frac{a_n}{n+1} z^{n+1} \qquad (\text{where } p(z) = \textstyle\sum_{n=0}^{N} a_n z^n).$$

Then for each polynomial p: $DSp = p$, $D^n p$ is eventually zero, and $S^n p \to 0$ uniformly on compact subsets of $\mathbb{C}$. Thus the hypotheses of Theorem 2.1 are satisfied, so D is hypercyclic on $\mathbb{C}$. □

In fact, this argument works just as well for $H(G)$ where G is any simply connected plane domain. All that is needed is that the polynomials be dense in $H(G)$, and this is provided by Runge's Theorem. This suggests the question of whether or not simple connectivity of a plane domain G is characterized by hypercyclicity of the differentiation operator on $H(G)$. This is in fact true—see [35] for the details, and for further variations on this theme.

For our final application we prove Birkhoff's result on the translation operator mentioned at the top of this section.

2.5. **Birkhoff's Translation Theorem** [5]. *For each complex number a, the operator T_a given by (2) is hypercyclic on $H(\mathbb{C})$.*

Proof. Instead of working with polynomials we consider exponentials E_λ defined by:

$$E_\lambda(z) = e^{\lambda z} \qquad (\lambda, z \in \mathbb{C}).$$

The key is that E_λ is an eigenvector for T_a with eigenvalue $e^{a\lambda}$, so in particular if $\mathrm{Re}\,(a\lambda) < 0$ then:

$$T_a^n E_\lambda = \exp\{n\mathrm{Re}\,(a\lambda)\} E_\lambda \to 0 \text{ in } H(\mathbb{C}) \qquad (n \to \infty).$$

Thus if Y is the linear span of the exponentials E_λ for $\mathrm{Re}\,(a\lambda) < 0$ (a half-plane H of points λ, whose boundary is the line through the origin orthogonal to a), we see that $T_a^n = T_{na} \to 0$ pointwise on Y. Similarly, let Z denote the linear span of the E_λ's for $\mathrm{Re}\,(a\lambda) > 0$. Then $T_a^n E_\lambda \to 0$, so if we can show that Y and Z are dense in $H(\mathbb{C})$ then we'll have the hypotheses of Theorem 2.1 satisfied with $S = T_a$.

In fact much more is true!

2.6. **Density Lemma.** *Suppose A is any subset of $\mathbb{C}$ with a limit point in $\mathbb{C}$, and let $E(A)$ be the linear span of the functions E_λ with $\lambda \in A$. Then $E(A)$ is dense in $H(\mathbb{C})$.*

Proof. Suppose Λ is a continuous linear functional on $H(\mathbb{C})$ that takes the value zero on each exponential function E_λ for $\lambda \in G$. By the Hahn-Banach Theorem it is enough to prove that $\Lambda \equiv 0$ on $H(\mathbb{C})$. For each $R > 0$ and each $f \in H(\mathbb{C})$ let $\|f\|_R := \max\{|f(z)| : |z| \leq R\}$. $\|\cdot\|_R$ is a norm on $H(\mathbb{C})$, and the open balls for each of these norms forms a basis for the topology of $H(\mathbb{C})$. Thus the inverse image of the unit disc under λ contains an $\|\cdot\|_R$-ball centered a the origin for some $R > 0$. In other words, Λ is a bounded linear functional relative to the norm $\|\cdot\|_R$, so by the Hahn-Banach theorem it extends to a bounded linear functional on $C(\{|z| \leq R\})$. By the Riesz Representation Theorem there is a finite Borel measure μ on the closed disc $\{|z| \leq R\}$ such that Λ is represented by integration against μ; in particular

$$\Lambda(f) = \int f\,d\mu \qquad (f \in H(\mathbb{C})).$$

Since the support of μ is compact, the function F defined on $\mathbb{C}$ by

$$F(\lambda) = \int E_\lambda\,d\mu = \int e^{\lambda z}\,d\mu(z)$$

is entire, and

$$D^n F(0) = \int z^n\,d\mu(z) \qquad (n = 0, 1, 2, \ldots).$$

But our hypothesis is that F vanishes on A, and since A has a finite limit point, the identity theorem for holomorphic functions insures that F vanishes on $\mathbb{C}$, hence the same is true of each of its derivatives. Thus $\int z^n\,d\mu(z) = 0$ for every non-negative integer, and therefore $\int f\,d\mu = 0$ for every holomorphic polynomial f, and so for every entire function f (every entire function is the limit, in $H(\mathbb{C})$ of the partial sums of its MacLaurin series). Thus Λ vanishes on $H(\mathbb{C})$, so by the Hahn-Banach Theorem, $E(A)$ is dense in $H(\mathbb{C})$.

This completes the proof of the Lemma, and with it, the proof of Birkhoff's Translation Theorem. □

Exercise. *Use the same idea to show that for any nonconstant polynomial p, the operator $p(D)$ is hypercyclic on $H(\mathbb{C})$, and even on $H(G)$ for any simply connected plane domain G.*

In fact the result of this exercise extends to any continuous linear operator L on $H(G)$ that commutes with D and is not a constant multiple of the identity; see [18, §5] for the details.

Examples of non-hypercyclicity

Why do most people find it surprising that linear operators can be hypercyclic, or mixing, or chaotic? Perhaps it's because the most common ones—the finite dimensional operators, are not. This is most easily seen by first considering a more general situation. Given an F-space X, denote the *dual space* of X (the space of continuous linear functionals on X) by X^*. Now if X is not locally convex we may have $X^* = \{0\}$ (e.g., $X = L^p([0,1])$ with $0 < p < 1$), but this does not affect our arguments. If T is a (continuous) linear operator on X, define its *adjoint* $T^* : X^* \to X^*$ by $T^*\Lambda := \Lambda \circ T$ for $\Lambda \in X^*$, so that T^* is a linear transformation of X^*.

2.7. **Theorem.** *Suppose T is a continuous linear operator on an F-space X. If the adjoint operator T^* has an eigenvalue, then T is not hypercyclic.*

Proof. We are saying that there is a continuous linear functional Λ on X that is not identically zero, and a complex number α such that $T^*\Lambda = \alpha\Lambda$, so for every positive integer n we have $T^{*n}\Lambda = \alpha^n\Lambda$. Thus if x is any vector in X then

$$\Lambda(T^n x) = T^{*n}\Lambda(x) = \alpha^n\Lambda(x) \qquad (n = 0, 1, 2, \dots),$$

hence $\Lambda(\mathrm{orb}\,(T,x)) = \{\alpha^n\Lambda(x)\}_0^\infty$. Now the set on the right-hand side of this identity is never dense in $\mathbb{C}$, whereas if $\mathrm{orb}\,(T,x)$ were dense in X, then its Λ-image would be dense in $\mathbb{C}$, by the continuity of Λ (the image of a dense set under a continuous map is dense in the image of the whole space). This contradiction proves the theorem. $\square$

Corollary. *There are no hypercyclic operators on finite dimensional F-spaces.*

Proof. Each finite dimensional F-space is isomorphic to $\mathbb{C}^n$ for some positive integer n. Therefore for each linear operator on such a space, the adjoint can also be viewed as an operator on $\mathbb{C}^n$, i.e., as an $n \times n$ complex matrix. So T^* has an eigenvalue, and therefore T is not hypercyclic. $\square$

Exercise. *No* compact *operator on a Banach space can be hypercyclic.*

3. Mixing transformations

In this section we return to the general setting of continuous maps of complete metric spaces. This is motivated by the fact that the proof of our sufficient condition for hypercyclicity (Theorem 2.1 of the last section) actually yields much more than was promised. In the first place, the conclusion itself is formally stronger than what's needed to be able to apply Birkhoff's Transitivity Theorem; the conclusion shows that for every pair U, V of nonempty open subsets of X there is a non-negative integer N such that $T^n(U) \cap V \neq \emptyset$ for all $n \geq N$ (Birkhoff's Theorem only demands *one* n for which the intersection is empty). Whenever this stronger property is true of a mapping of a metric space, we call that mapping *(topologically) mixing*. Thus, for example, all the linear operators we have proved to be transitive are actually mixing.

Next, the proof of Theorem 2.1 never made full use of linearity; it works just as well if X is merely a complete, metrizable topological group (with the group identity taking the place of the vector space zero element), and T is a continuous homomorphism of X into itself. The group need not even be abelian! For example the space Σ of $(0,1)$-sequences introduced in §1.10 is a group where addition is done coordinatewise modulo two, and the backward shift β on that space, which we proved to be transitive by constructing a point in Σ with dense orbit, actually obeys the hypotheses of our refined sufficient condition. Thus β isn't just transitive, it too is mixing!

3.1. **Exercise.** *Every quasi-factor of a mixing map is mixing.*

Clearly the maps of irrational translations mod one, though transitive, are not mixing. However all the other maps we discussed in §1 are factors of the backward shift on Σ, so they *are* mixing.

3.2. **Proposition.** *If X is a complete, separable metric space and T a continuous mapping on X that is mixing, then for every strictly increasing sequence $\{n(k)\}$ of positive integers there is a dense G_δ subset of points $x \in X$ for which $\{T^{n(k)}x : k \geq 0\}$ is dense in X.*

Proof. Fix a countable basis for the topology of X, and enumerate the *pairs* of basis sets, obtaining a sequence $\{U_i, V_i\}_1^\infty$ (with each of the basis sets occurring in each coordinate infinitely often). Because T is mixing:

$$\exists\ \nu_1 \in \mathbb{N} \text{ such that } n \geq \nu_1 \implies T^n(U_1) \cap V_1 \neq \emptyset.$$

In particular,

$$\exists\ k_1 \in \mathbb{N} \text{ with } T^{n(k_1)}(U_1) \cap V_1 \neq \emptyset.$$

Similarly

$$\exists\ \nu_2 \in \mathbb{N} \text{ such that } n \geq \nu_2 \implies T^n(U_2) \cap V_2 \neq \emptyset.$$

from which it follows that

$$\exists\ k_2 \in \mathbb{N} \text{ with } k_2 > k_2 \text{ and } T^{n(k_1)}(U_1) \cap V_1 \neq \emptyset.$$

Continuing in this manner you get a strictly increasing sequence $\{n(k_j)\}_{j=1}^{\infty}$ such that $T^{n(k_j)}U_j \cap V_j \neq \emptyset$. In particular, for any nonvoid open subsets U and V of X there exists an index k such that $T^{n(k)}U \cap V \neq \emptyset$.

The proof of Birkhoff's Transitivity Theorem, repeated almost word-for-word, now shows that there is a dense G_δ set of points $x \in X$ such that $\{T^n(k)x\}$ is dense in X. □

As a special case, for each fixed positive integer n we may take $n_k = kn$ in the result above, thus obtaining:

3.3. **Corollary.** *If T is a continuous, mixing transformation of a complete metric space X, then T^n is transitive for every positive integer n.*

By contrast:

3.4. **Proposition.** *Not every transitive map has transitive powers.*

Proof. To make examples of this phenomenon, let T be any continuous mixing map of a complete metric space X, and form a new $(\widetilde{X}, \widetilde{T})$ as follows:

$$\widetilde{X} = X \times \{1,2\} \quad \text{and } \widetilde{T}(x,1) = (Tx,2),\ \widetilde{T}(x,2) = (Tx,1) \text{ for } x \in X.$$

So $\widetilde{X}$ is the disjoint union of subsets $\widetilde{X_j} = X \times \{j\}$ where $j = 1,2$, and $\widetilde{T}$ maps $\widetilde{X_1}$ onto $\widetilde{X_2}$ in the same way T would map X onto itself, and it similarly interchanges $\widetilde{X_2}$ and $\widetilde{X_1}$. Perhaps one should think of each of these subsets as a copy of X, with the first one colored red and the second blue. Then $\widetilde{T}$ has the action of T, except it takes each red point to the "same" point, but now colored blue, and vice versa.

Now $\widetilde{X}$ is the product of two metric spaces, X, with its metric d and the two-point space $\{1,2\}$, with the discrete metric δ. Therefore the metric on $\widetilde{X}$ defined by:

$$\widetilde{d}((x,j),(y,k)) := d(x,y) + \delta(j,k)$$

puts the product topology on $\widetilde{X}$, thereby making it into a complete metric space on which $\widetilde{T}$ acts continuously.

Clearly $\widetilde{T}^2$ is not transitive since it takes $\widetilde{X_j}$ into itself for each $j = 1, 2$. On the other hand $\widetilde{T}$ itself transitive. Indeed, because T is assumed to be mixing, there is a dense G_δ subset of points $x \in X$ for which $\{T^{2n}x\}$ is dense in X, and a similar set of x's for which $\{T^{2n+1}x\}$ is dense. It follows that there are points x in X for which both sequences are dense. Thus $\{\widetilde{T}^{2n}(x,1)\}$ is dense in $\widetilde{X_1}$, while $\{\widetilde{T}^{2n+1}(x,1)\}$ is dense in $\widetilde{X_2}$, hence $\operatorname{orb}(\widetilde{T},(x,1))$ is dense in $\widetilde{X}$, i.e., $\widetilde{T}$ is transitive. $\square$

By refining this construction just slightly we can create examples where the underlying space $\widetilde{X}$ is connected. For this, suppose X is connected and $T : X \to X$ is continuous and mixing, but with a fixed point $p \in X$. For example, the tent and quadratic maps on $[0,1]$ have these properties (both fix the points 0 and 1).

Let $\hat{X}$ be $\widetilde{X}$ with $(p,1)$ and $(p,2)$ identified to a point we'll call $\hat{p}$. Define $\hat{T}$ on $\hat{X}$ in the obvious way: it's just $\widetilde{T}$ on $\hat{X}\backslash\{p\}$, and it fixes the point $\hat{p}$. The resulting $\hat{T}$ is continuous on $\hat{X}$, and if $V : \widetilde{X} \to \hat{X}$ is the identity on $\widetilde{X}\backslash\{(p,1),(p,2)\}$ and takes (p,j) to $\hat{p}$, then V is also continuous and $V \circ \widetilde{T} = \hat{T} \circ V$, i.e., $\hat{T}$ is a factor of $\widetilde{T}$, and therefore is also transitive. But just as before, $\hat{T}^2$ is not transitive, since for $j = 1, 2$ it still takes each of the disjoint sets $\widetilde{X_j}$ into itself. $\square$

For a more concrete example, it might be instructive to show that the map you get by this construction from the tent map on the unit interval is conjugate to the "double tent map" defined on $[-1,1]$ by: $\hat{T}(x) = T(x+1)$ for $-1 \le x \le 0$ and $= -T(x)$ for $0 \le x \le 1$.

3.5. **Product maps.** The construction given above for $\widetilde{T}$ is actually a special case a much more far-reaching idea. Let $\{X_i\}$ be a finite or countable collection of metric spaces. On each X_i we can assume without loss of generality that the metric d_i is bounded by one (otherwise, replace it by $d_i/(1+d_i)$). Then the cartesian product X of the spaces X_i is a metric space in the metric $d = \sum_i 2^{-i} d_i$, and a sequence of points converges in this product metric if and only if it converges in each coordinate. Thus (X, d) has the product topology; it is compact if each (X_i, d_i) is compact, and complete if each (X_i, d_i) is complete.

Suppose T_i a mapping of X_i for each i. Then the "product map" T is defined on the product space X by letting T_i act in each coordinate. If each T_i is continuous X_i, then T is continuous on X (and conversely). In particular, the map $\widetilde{T}$ constructed above is an example of just such a product construction: Both set-theoretically and topologically $\widetilde{X} = X \times \{1, 2\}$, and $\widetilde{T} = T \times S$, where now $T : X \to X$ is the original map on X and $S : \{1, 2\} \to \{1, 2\}$ is the permutation that interchanges 1 and 2.

3.6. **Proposition.** *If X is the product of a finite or countable collection $\{X_i\}$ of metric spaces, and T_i is a mixing transformation of X_i for each i, then the product map T is mixing on the product space X.*

Proof. Suppose U and V are basic open sets in X, i.e. that $U = \prod_i U_i$ with U_i open in X_i and $U_i = X_i$ for all $i \geq$ some n_1. Similarly $V = \prod_i V_i$ with V_i open in X_i and equal to X_i for all $i \geq$ some, possibly different, index n_2. Let n be the larger of n_1 and n_2. Because each T_i is mixing we may choose a positive integer N such that

$$k \geq N \Rightarrow T_i^{-k}(U_i) \cap V_i \neq \emptyset \text{ for each } 1 \leq i \leq n,$$

hence for each such k:

$$T^{-k}(U) \cap V = \prod_{i=1}^{n} \left(T_i^{-k}(U_i) \cap V_i\right) \times \prod_{i \geq n+1} X_i \neq \emptyset,$$

which shows that T is mixing. □

3.7. **Definition.** *A mapping T of a metric space X is called* weakly mixing *if $T \times T$ is transitive on $X \times X$.*

We have just seen that every mixing transformation is weakly mixing, and it is easy to check that weakly mixing transformations are transitive (more generally, if a product mapping is transitive then the same is true of each coordinate mapping). However the converse is not true: there are transitive maps T that are not weakly mixing. The next result implies that the the examples constructed in the proof of Proposition 3.4 (with T is transitive but T^2 not) all have this property.

3.8. **Proposition.** *Suppose T is a continuous map of a complete metric space X. If T is weakly mixing then T^2 is transitive.*

Proof. We use Birkhoff's characterization of transitivity for this setting. Fix U and V open in X and nonempty. It is enough to find an even integer m such that $T^{-m}(U) \cap V \neq \emptyset$. Use the transitivity of $T \times T$ to choose $n > 1$ such that

$$(T \times T)^{-n}\left((U \times U) \cap (V \times T^{-1}(V))\right) \neq \emptyset$$

But the left-hand side of this expression is just $(T^{-n}(U) \cap V) \times (T^{-n}(U) \cap T^{-1}(V))$, hence $T^{-n}(U) \cap V$ and $T^{-(n-1)}(U) \cap V$ are both nonempty. Since either n or $n-1$ is even, it follows that T^2 is transitive. □

4. Hypercyclic composition operators

Historically the notion of transitivity was not foremost on the minds of operator theorists, an understandable oversight since when dealing with linear operators one thinks about invariant *subspaces*, but not so much about invariant *sets*. Suppose X is a vector space, T a linear transformation on X, and x a vector in X. Corresponding to the fact that $\operatorname{orb}(T, x)$ is the smallest T-invariant set containing x, the *linear span* of this orbit is the smallest T-invariant subspace containing x, and the closure of this linear span is the smallest *closed* T-invariant subspace of X containing x. If $\operatorname{span}\operatorname{orb}(T, x)$ is dense in X we say T is *cyclic* and call x a *cyclic vector* for T.

Thus noncyclic vectors generate proper, closed, invariant subspaces, and if there are no cyclic vectors then there are no closed invariant subspaces except for $\{0\}$ and X. The first example of a Banach space operator having only trivial closed invariant subspaces was constructed by Enflo [16] in the 1980's, and his work was later simplified by Read, who eventually showed that it is even possible for an operator on a Banach space to have no hypercyclic (i.e., transitive) vector [29], and hence no closed invariant subsets other than the zero vector and the whole space. However none of the examples produced so far is set in a Hilbert space, and for this special case it's a famous open question to decide if every bounded operator has a closed invariant subspace (or a closed invariant subset) $\neq \{0\}$ or the whole Hilbert space.

So it makes sense to think of transitivity of a linear operator as a very strong form of cyclicity. This is why operator theorists use the term "hypercyclic" instead of "transitive." In this section we'll find more examples to illustrate the point that hypercyclicity occurs surprisingly often. Then we'll explore a few of the ways in which hypercyclicity is a more "robust" concept than cyclicity.

Composition operators

Recall that if G is an open subset of the plane, then the space $H(G)$ of all complex-valued functions holomorphic on G can be made into an F-space by a complete metric for which a sequence $\{f_n\}$ in $H(G)$ converges to $f \in H(G)$ if and only if $f_n \to f$ uniformly on every compact subset of G. If G_1 and G_2 are open subsets of $\mathbb{C}$ and $\varphi : G_1 \to G_2$ a holomorphic map (not necessarily one to one or onto), then φ induces a *composition operator* $C_\varphi : H(G_2) \to H(G_1)$ defined by:

$$C_\varphi f = f \circ \varphi \qquad (f \in H(G_2)).$$

If G_1, G_2, and G_3 are all open subsets of $\mathbb{C}$ with $\varphi : G_1 \to G_2$ and $\psi : G_2 \to G_3$ holomorphic maps, then $C_{\psi\circ\varphi} = C_\varphi \circ C_\psi$. We will focus primarily on holomorphic self-maps φ of plane domains G, for which the composition formula above yields the iteration formula:

$$C_\varphi^n = C_{\varphi_n} \qquad (n = 0, 1, 2, \dots),$$

where, to avoid confusion with the n-fold pointwise product, φ_n denotes the composition of φ with itself n times. This simple observation suggests that there should be intriguing connections between the dynamical behavior of a composition operator C_φ with that of its inducing map φ. In particular, *which composition operators are hypercyclic on $H(G)$?* In order the operators foremost, we restrict attention to the simplest setting: $G = \mathbb{U}$, the open unit disc. Note, however, that the Riemann Mapping Theorem will allow us to transfer dynamical results about composition operators on $H(\mathbb{U})$ to $H(G)$ where G is any simply connected plane domain $\neq \mathbb{C}$. Indeed, Riemann's theorem guarantees that there is a univalent holomorphic map σ taking $\mathbb{U}$ onto G, hence the corresponding composition operator C_σ is an isomorphism (one-to-one, onto, linear, bi-continuous) of $H(G)$ onto $H(\mathbb{U})$. If $\varphi : G \to G$ is a holomorphic self-map of G, then $\psi := \sigma \circ \varphi \circ \sigma^{-1}$ is a holomorphic self-map of $\mathbb{U}$ that is (holomorphically) conjugate to φ. As for the corresponding operators,

$C_\psi = (C_\sigma)^{-1} C_\varphi C_\sigma$, so $C_\psi : H(G) \to H(G)$ is similar (i.e. linearly conjugate) to $C_\varphi : H(\mathbb{U}) \to H(\mathbb{U})$.

Henceforth φ will always denote a holomorphic self-map of $\mathbb{U}$, and we will abbreviate "holomorphic and one-to-one" by "univalent". Our first result severely limits the kinds of maps that can produce hypercyclic behavior.

4.1. **Proposition.** *If C_φ is hypercyclic on $H(\mathbb{U})$ then φ is univalent, and has no fixed point in $\mathbb{U}$.*

Proof. Suppose φ has a fixed point $p \in \mathbb{U}$. Then for $f \in H(\mathbb{U})$, any function in $\operatorname{orb}(C_\varphi, f)$ must have value $f(p)$ at p, hence the same must be true of any function in the closure of this orbit. Thus no C_φ-orbit is dense, hence C_φ is not hypercyclic.

Suppose φ is not univalent, so there exist distinct points $p, q \in \mathbb{U}$ with $\varphi(p) = \varphi(q)$. Then if $f \in H(\mathbb{U})$, each function in $\operatorname{orb}(C_\varphi, f)$ takes the same value at p as at q, and again this property gets passed on to functions in the closure of the orbit. Once again, no orbit can be dense. $\square$

Note that this proof did not make any use of the special properties of the unit disc, so the result is valid for any open set G.

In a sense to be made precise later, univalent self-maps of the unit disc are modelled by linear fractional transformations, so we consider this class of maps first. We need to know how these maps are classified in terms of their fixed points; then next few paragraphs review this matter.

4.2. **Linear fractional transformations.** Recall that a *linear fractional transformation* (henceforth, "a LFT") is a mapping of the form

$$\varphi(z) = \frac{az+b}{cz+d} \quad \text{where} \quad \Delta := ad - bc \neq 0. \tag{5}$$

The condition $\Delta \neq 0$ is necessary and sufficient for φ to be nonconstant, as you see when $c \neq 0$ by checking the formula:

$$\varphi(z) - \frac{a}{c} = -\frac{\Delta}{c}\,\frac{1}{cz+d}.$$

When $c = 0$ then the condition $\Delta \neq 0$ implies that neither d nor a is zero, hence φ is a nonconstant affine mapping of $\mathbb{C}$. We extend φ to a mapping of the extended complex plane

$\hat{\mathbb{C}} := \mathbb{C} \cup \{\infty\}$ onto itself by defining $\varphi(\infty) = a/c$ and $\varphi(-d/c) = \infty$ if $c \neq 0$, and $\varphi(\infty) = \infty$ if $c = 0$ (the affine case). If $\hat{\mathbb{C}}$ is then identified with the Riemann Sphere Σ via stereographic projection, then φ becomes a homeomorphism of Σ onto itself.

We employ the classification of LFT's in terms of their fixed points. Each LFT has one or two fixed points in $\hat{\mathbb{C}}$. If there's just one fixed point the map is called *parabolic.* If a parabolic map φ has its fixed point at ∞, then it's easy to check that φ is a translation: $\varphi(z) = z + \tau$ for some complex number τ. If, however, its fixed point is $p \in \mathbb{C}$, then the LFT $\alpha(z) = 1/(z - p)$ takes p to ∞, and so $\alpha \circ \varphi \circ \alpha^{-1}$ is an LFT that fixes ∞, and so is a translation. Thus: *the parabolic LFT's are precisely the ones that are (linear-fractionally) conjugate to translations.*

Note also that ∞ is an attractive fixed point for any translation, in the sense that the sequence of iterates converges to ∞ uniformly on compact subsets of $\mathbb{C}$. Thus if φ is any parabolic LFT with fixed point $p \in \hat{\mathbb{C}}$, then the sequence of iterates $\{\varphi_n\}$ converges to p uniformly on compact subsets of $\hat{\mathbb{C}} \backslash \{p\}$.

The other possibility is that φ has two distinct fixed points, say p and q. If these are 0 and ∞ then φ is a complex dilation: $\varphi(z) = \kappa z$ for some $\kappa \in \mathbb{C}$. In the general case, one fixes an LFT α that maps p to ∞ and q to zero, (e.g., $\alpha(z) = (z - p)/(z - q)$), in which case $\alpha^{-1} \circ \varphi \circ \alpha$ fixes 0 and ∞, hence is a "κ-dilation," as above. Now κ is not uniquely determined: the maps $z \to \kappa z$ and $z \to (1/\kappa)z$ are linear-fractionally conjugate to each other. But this is as bad is things can get—if one conjugation of φ to an LFT with fixed points 0 and ∞ gives you a multiplier of κ, then any other one will give either κ or $1/\kappa$ (see [34, Chapter 0] for more details).

If κ is positive we say φ is *hyperbolic*, otherwise φ is *loxodromic.* Note that in the hyperbolic or loxodromic cases one of the fixed points is attracting and the other repelling. The only other possibility is $|\kappa| = 1$ in which case φ is called *elliptic.* Here the fixed points are neither attracting nor repelling.

4.3. **Linear fractional maps of the unit disc.** How does our classification of linear fractional maps fare if we additionally require that the unit disc be taken into itself? Let LFT($\mathbb{U}$) denote this class of maps, and Aut ($\mathbb{U}$) denote the subclass of "conformal automorphisms of

$\mathbb{U}$, i.e., linear fractional maps that take $\mathbb{U}$ *onto* itself. Since we aim to study hypercyclic composition operators, we will focus on maps in LFT($\mathbb{U}$) that fix no point of $\mathbb{U}$.

If $\varphi \in$ LFT($\mathbb{U}$) is parabolic, then, because the fixed point is attractive, it must lie on the unit circle. Upon conjugating by an appropriate rotation we may assume this fixed point is 1. Now the map $z \to (1+z)/(1-z)$ takes the unit disc onto the open right half-plane $\mathbb{P}$, and takes 1 to ∞, so this map conjugates φ to a translation of $\mathbb{P}$ into itself, i.e., a map of the form $w \to w + \tau$ where $\operatorname{Re}\tau \geq 0$. Note that φ maps $\mathbb{U}$ *onto* itself if and only if $\operatorname{Re}\tau = 0$. So *parabolic maps are LFT-conjugate to translations of the right half-plane into itself,* with the automorphisms corresponding to the pure imaginary translations.

Similarly, if $\varphi \in$ LFT($\mathbb{U}$) is hyperbolic and an automorphism, then both its fixed points must lie on $\partial\mathbb{U}$ (as before, this is clear for the attractive one, and the repelling one is attractive for φ^{-1}, which is also an automorphism of $\mathbb{U}$). In this case any LFT that sends one of these fixed points to the origin and the other to ∞ conjugates φ to a positive dilation (as we have seen), and takes the unit circle to a straight line, with $\mathbb{U}$ going to one of the half-planes bounded by this line. Conjugation by an appropriate rotation takes this half-plane to $\mathbb{P}$, while leaving the dilation unchanged. Thus: *every hyperbolic automorphism of* $\mathbb{U}$ *is conjugate to a positive dilation of* $\mathbb{P}$.

Finally, if $\varphi \in$ LFT($\mathbb{U}$) is hyperbolic, but not an automorphism, then its attractive fixed point must still lie on $\partial\mathbb{U}$, while the repulsive one is either in $\mathbb{U}$ (the case we're not considering) or outside the closure of $\mathbb{U}$. By conjugating with a rotation we may assume the attractive fixed point is 1. Let q denote the repulsive one. Then the reflection q^* of q in the unit circle lies in $\mathbb{U}$, so for an appropriate unimodular constant ω the $\mathbb{U}$-automorphism $z \to \omega(p^* - z)/(1 - \overline{p^*}z)$ takes p^* to zero while fixing 1. This map therefore takes p to ∞, and therefore conjugates φ to a map $\Phi \in$ LFT($\mathbb{U}$) that fixes both 1 and ∞. Thus $\Phi(z) = az + b$ where $|a| + |b| \leq 1$ (since $\Phi(\mathbb{U}) \subset \mathbb{U}$) and $a + b = 1$ (since $\Phi(1) = 1$). It follows that $0 < a, b < 1$, so $\Phi(z) = az + 1 - a$ for some $0 < a < 1$. Note that Φ also belongs to LFT($\mathbb{U}$), and that φ is LFT($\mathbb{U}$)-conjugate to Φ. Note further that Φ is a hyperbolic automorphism of the half-plane $H := \{\operatorname{Re} z < 1\}$ and that its attractive fixed point is 1, the same as for φ.

With these results in hand we can classify those $\varphi \in$ LFT($\mathbb{U}$) that induce hypercyclic composition operators on $H(\mathbb{U})$. The result is surprisingly unsubtle:

4.4. **Theorem.** *Every $\varphi \in LFT(\mathbb{U})$ with no fixed point in $\mathbb{U}$ induces a hypercyclic composition operator on $H(\mathbb{U})$.*

Proof. The idea is to reduce everything, case by case, to Birkhoff's translation theorem (2.5). Fix $\varphi \in \mathrm{LFT}(\mathbb{U})$ with no fixed point in $\mathbb{U}$. Then φ is either parabolic or hyperbolic.

(a) If φ is parabolic then we know that there is a LFT σ mapping $\mathbb{U}$ onto $\mathbb{P}$ such that $\varphi = \sigma^{-1} \circ T_\tau \circ \sigma$, where $T_\tau(w) = w + \tau$ for $w \in \mathbb{P}$, and τ is a complex number with non-negative real part (so τ maps $\mathbb{P}$ into itself). As we saw at the beginning of this discussion, the composition operator $C_\sigma : H(\mathbb{P}) \to H(\mathbb{U})$ establishes a similarity (in the language of previous sections: a *linear conjugacy*) between C_φ on $\mathbb{U}$ and C_τ on $\mathbb{P}$. So it is enough to prove that C_τ is hypercyclic on $H(\mathbb{P})$.

For this, note that the restriction map $R : H(\mathbb{C}) \to H(\mathbb{P})$, defined by $Rf := f|_{\mathbb{P}}$ for $f \in H(\mathbb{C})$ is a 1-1 continuous linear map from $H(\mathbb{C})$ into $H(\mathbb{P})$ (it can be viewed as the composition operator induced by the identity map $\mathbb{P} \to \mathbb{C}$). Moreover $RC_\tau = C_\tau R$, where in the first instance C_τ is acting on $H(\mathbb{C})$ and in the second on $H(\mathbb{P})$. Finally, the range of R, namely $H(\mathbb{P})$ is dense in $H(\mathbb{C})$—this is a consequence of Runge's theorem which insures that the polynomials are dense in both spaces (their density in $H(\mathbb{C})$ is also obvious from the convergence properties of power series). Thus C_τ on $H(\mathbb{P})$ is semi-conjugate to C_τ on $H(\mathbb{C})$, and therefore the former operator inherits the hypercyclicity that Birkhoff's Translation Theorem provides for the latter one.

(b) Suppose φ is a hyperbolic automorphism. Then φ is LFT-conjugate to a dilation $\Delta_r : \mathbb{P} \to \mathbb{P}$, defined by $\Delta_R(w) = rw$ for some fixed $0 < r < 1$. The principal branch Λ of the logarithm takes $\mathbb{P}$ univalently onto the horizontal strip $S = \{W : |\mathrm{Im}\, W| < \pi/2\}$, and the corresponding composition operator effects a similarity between $C_{\Delta_\rho} : H(\mathbb{P}) \to H(\mathbb{P})$ and the translation operator $T_{\Lambda(\rho)}$, now acting on $H(S)$. Just as we saw in the parabolic case, the map that restricts an entire function to S is a continuous linear embedding of $H(\mathbb{C})$ into $H(S)$ with dense range, and it reveals $T_{\Lambda(\rho)}$, acting on $H(S)$ as a quasi-factor of the same operator acting on $H(\mathbb{C})$. Hence, as before, the translation operator on $H(S)$ is hypercyclic, and therefore so is its conjugate, C_φ on $H(\mathbb{U})$.

(c) Finally, suppose $\varphi \in \mathrm{LFT}(\mathbb{U})$ is hyperbolic, but *not* an automorphism. Then we have seen that φ is $\mathrm{LFT}(\mathbb{U})$-conjugate to the restriction to $\mathbb{U}$ of a hyperbolic automorphism Φ

of the half-plane $P = \{\operatorname{Re} z < 1\}$. So it is enough to show C_Φ is hypercyclic on $H(\mathbb{U})$. By part (b) above, C_Φ on $H(P)$ is conjugate to a composition operator induced by a hyperbolic automorphism of $\mathbb{U}$, and so is hypercyclic. The same observations we employed to prove part (b) show that C_Φ acting on $H(\mathbb{U})$ is a quasi-factor of C_φ acting on $H(P)$, so it too is hypercyclic. This completes the proof of our theorem. □

4.5. **Chaos.** In §1.16 we defined a map to be *chaotic* if it is transitive and has a dense set of periodic points. Which of the linear maps we've proved transitive are chaotic?

Answer: *All of them!*

To see this, recall that the operators of differentiation and translation have the exponential functions $E_\lambda(z) = e^{\lambda z}$ as eigenvectors: $DE_\lambda = \lambda E_\lambda$, and $T_a E_\lambda = e^{a\lambda} E_\lambda$. Thus if λ is an n-th root of unity, E_λ is a periodic point of D with period n. Any linear combination of such periodic points is again periodic, with period equal to the least common multiple of the original periods. Since the set of roots of unity is dense in the unit circle, the Density Lemma shows that D has a dense set of periodic points. The same argument works for T_a, except now you start with exponentials E_λ such that $e^{a\lambda}$ is a root of unity.

As for composition operators on $H(\mathbb{U})$ induced by linear fractional maps, our proof that they are all hypercyclic established that they are all quasi-factors of translation operators acting on $H(\mathbb{C})$, and therefore they inherit the chaotic behavior just proved for those operators. □

Remark. You can think of conformal automorphisms with no fixed point in $\mathbb{U}$ as "non-Euclidean translations" of the unit disc, with the attractive fixed point of φ (which necessarily lies on $\partial\mathbb{U}$) playing the role that ∞ plays in the Euclidean case. From this point of view the fact that such maps induce hypercyclic composition operators on $H(\mathbb{U})$ can be viewed as the non-Euclidean analogue of Birkhoff's Translation Theorem. Thus it is only fitting that the non-Euclidean result, which was proved about sixty years ago by Seidel and Walsh [37], actually follows from Birkhoff's.

Beyond linear-fractional

Having disposed of the linear-fractional case, it's time to ask whether or not a composition operator induced by an arbitrary univalent self-map of $\mathbb{U}$ is hypercyclic or chaotic on $H(\mathbb{U})$.

Given the special nature of the proofs used in the linear-fractional setting, one might be tempted to dismiss the question as too general to admit a definitive solution. However one would be wrong; one of the landmark theorems of classical analytic function theory renders the general problem an almost trivial consequence of our analysis of the linear-fractional situation. This is:

4.6. **The Linear Fractional Model Theorem.** *If φ is a univalent, holomorphic self-map of $\mathbb{U}$, then there exists a linear-fractional map ψ and a univalent map $\sigma : \mathbb{U} \to \mathbb{C}$ such that $\sigma \circ \varphi = \psi \circ \sigma$. If φ has no fixed point in $\mathbb{U}$ then there are two possibilities:*

(a) *ψ can be taken to be dilation: $\psi(z) = rz$ for some $0 < r < 1$, and $\sigma(\mathbb{U}) \subset \mathbb{P}$, or*

(b) *ψ can be taken to be a translation: $\psi(z) = z + \tau$ for some τ in $\mathbb{C}$.*

There is actually much more to this remarkable theorem, and we'll talk about this in a moment. Note that the two cases distinguished above correspond precisely to what happens for hyperbolic and parabolic maps in $\mathrm{LFT}(\mathbb{U})$, so in some sense these maps are "models" for univalent self-maps of $\mathbb{U}$ that fix no point of $\mathbb{U}$. Just how one distinguishes hyperbolic from parabolic behavior for such general self-maps of $\mathbb{U}$ is a fascinating question, which we will discuss shortly. But first note that theorem above makes short work of the hypercyclicity problem for composition operators on $H(\mathbb{U})$.

4.7. **Theorem.** *Suppose φ is a univalent holomorphic self-map of $\mathbb{U}$ that has no fixed point in $\mathbb{U}$. Then C_φ is chaotic on $H(\mathbb{U})$.*

Proof. The proof is the same one we used in the linear fractional case. In case ψ is a dilation, so that $\sigma(\mathbb{U}) \subset \mathbb{P}$, a further mapping by the principal branch Λ of the logarithm replaces it by a translation that takes $G := \Lambda(\sigma(\mathbb{U}))$ into itself. Since G is simply connected, Runge's Theorem insures that the polynomials are dense in $H(G)$, hence the entire functions are dense in $H(G)$. The rest follows just as before; our translation, and therefore C_φ itself, is exhibited as a quasi-factor of the same translation acting on $H(\mathbb{C})$. In the parabolic case the same proof works, except that now there is no need to call upon the logarithm. □

Remarks. Recall that in the proof of Birkhoff's Translation Theorem (Theorem 2.5) it was the sufficient condition 2.1 that proved hypercyclicity for translation operators on $H(\mathbb{C})$. Thus by the first paragraph of §3: *Every translation operator on $H(\mathbb{C})$ is mixing.* Since all the composition operators on $H(\mathbb{U})$ that we studied above turned out to be quasi-factors of such translations: *Every composition induced on $H(\mathbb{U})$ by a univalent, fixed-point-free holomorphic self-map of $\mathbb{U}$ is mixing.* So thus far, all our hypercyclic examples turn out to be both chaotic and mixing. This is not the case in general. For example, in [8] examples are given of hypercyclic translation operators T on some Hilbert spaces of entire functions for which the spectrum of T is the single point $\{1\}$. Since periodic points of linear operators are eigenvectors whose eigenvalue is a root of unity, these translation operators have no hope of being chaotic. We will see in the next section that, thanks to an example of Salas, not every hypercyclic operator is mixing.

Iteration and linear fractional models.

The Denjoy-Wolff Theorem. In a certain sense, every holomorphic self-map of $\mathbb{U}$ has an attractive fixed point: if there is not one in $\mathbb{U}$, then there is a unique boundary point that serves the purpose. This is the content of the famous *Denjoy-Wolff Theorem,* which figures importantly in many aspects of the study of composition operators. To simplify its statement let's adopt some terminology.

- A point $p \in \partial\mathbb{U}$ is a *boundary fixed point* of φ if φ has non-tangential limit p at p.
- The notation $\stackrel{\kappa}{\to}$ indicates uniform convergence on compact subsets of $\mathbb{U}$,
- If the derivative of φ has a nontangential limit at a boundary point p of $\mathbb{U}$, and the non-tangential limit of φ at p (whose existence follows easily from that of the derivative) has modulus one, we say φ has an *angular derivative* at p, and denote the limit by $\varphi'(p)$.

It may seem harsh to require that φ have non-tangential limit of modulus one at any boundary point at which its angular derivative is to exist, but when dealing with composition operators, this is exactly what makes the concept meaningful (see [34, Chapter 4], for example).

4.8. **The Denjoy-Wolff Iteration Theorem.** *Suppose φ is an analytic self-map of $\mathbb{U}$ that is not an elliptic automorphism.*

(a) *If φ has a fixed point $p \in \mathbb{U}$, then $\varphi_n \stackrel{\kappa}{\to} p$ and $|\varphi'(p)| < 1$.*

(b) *If φ has no fixed point in $\mathbb{U}$, then there is a point $p \in \partial\mathbb{U}$ such that $\varphi_n \stackrel{\kappa}{\to} p$. Furthermore:*

  ∗ *p is a boundary fixed point of φ; and*

  ∗ *the angular derivative of φ exists at p, with $0 < \varphi'(p) \le 1$.*

(c) *Conversely, if φ has a boundary fixed point p at which $\varphi'(p) \leq 1$ then φ has no fixed points in $\mathbb{U}$, and $\varphi_n \xrightarrow{\kappa} p$.*

The fixed point p to which the iterates of φ converge is called the *Denjoy-Wolff* point of φ. Part (a), which is an exercise based on the Schwarz Lemma, is not really part of the original theorem; it is included here only for completeness. For a proof of Theorem 4.8, and for further connections with the theory of composition operators, see [34, Chapter 5] or [13, Section 2.4].

Classification of linear-fractional maps

The Denjoy-Wolff Theorem suggests a "linear-fractional-like" classification of arbitrary holomorphic self-maps of $\mathbb{U}$. For motivation, let's review how the *linear-fractional* self-maps of $\mathbb{U}$ fall into distinct classes determined by their fixed-point properties (cf. [34]: Chapter 0). These are:

- *Maps with interior fixed point.* We didn't concentrate much on this case previously, but by an argument based on the Schwarz Lemma, the interior fixed point is either attractive, or the map is an elliptic automorphism. In both cases the map is conjugate to a dilation $z \to \lambda z$ for some complex number λ with $0 < |\lambda| \leq 1$.
- *Hyperbolic maps with attractive fixed-point on $\partial\mathbb{U}$.* Upon chasing through the classification of these maps as conjugate to dilations of the right half-plane, you see that they are the self-maps of $\mathbb{U}$ having no fixed point in $\mathbb{U}$, with derivative < 1 at the attractive boundary fixed point.
- *Parabolic maps.* These have exactly one fixed point on the Riemann Sphere, necessarily lying on $\partial\mathbb{U}$. These maps are characterized by the fact that they have derivative $= 1$ at the fixed point.

The parabolic self-maps of $\mathbb{U}$ fall into two subclasses, which one distinguishes by examining their the action of the corresponding maps on the right half-plane:

- *The automorphisms.* These are distinguished by the property that each orbit is separated in the hyperbolic metric (meaning that, for each $z \in \mathbb{U}$, the hyperbolic distance between successive points of the orbit $(\varphi_n(z))$ stays bounded away from zero).

- *The nonautomorphisms.* For these, the orbits are not hyperbolically separated, i.e., the hyperbolic distance between successive orbit points tends to zero.

An elementary argument establishes these last two statements. The first just reflects the fact that automorphisms are hyperbolic isometries. The second is best viewed in the context of the right half-plane $\mathbb{P}$. Suppose ψ is a parabolic self-map of $\mathbb{U}$ with fixed point at 1, and let

$$T(w) = \frac{w+1}{w-1}, \quad \text{and} \quad \Psi := T \circ \psi \circ T^{-1}.$$

Thus T is a linear-fractional mapping of $\mathbb{U}$ onto $\mathbb{P}$ that takes 1 to ∞, and one easily checks that $\Psi(w) = w + \psi''(1)$. It follows that $\psi''(1)$ has non-negative real part (otherwise Ψ could not map $\mathbb{P}$ into itself), and since ψ is not an automorphism of $\mathbb{U}$, $\psi''(1)$ cannot be pure imaginary. Now hyperbolic discs in $\mathbb{P}$ of fixed radius have this property: their Euclidean size is proportional to the real part of their hyperbolic center (see section 4, or [35, Chapter 4] for the details). Our hypothesis on the translation distance $\psi''(1)$ insures that for each $w \in \mathbb{P}$ the Ψ-orbit $(\Psi_n(w))$ has unbounded real part, but fixed Euclidean distance $|\psi''(1)|$ between successive points. Thus for all sufficiently large n, the hyperbolic disc of radius ε about $\Psi_n(w)$ contains $\Psi_{n+1}(w)$, hence the orbit of w is not separated.

Motivated by the classification of linear-fractional self-maps of $\mathbb{U}$, and encouraged by the restrictions the Denjoy-Wolff Theorem places on the values the derivative of an arbitrary self-map can take at the Denjoy-Wolff point, we introduce the following general classification scheme.

4.9. **Classification of arbitrary self-maps.** A holomorphic self-map φ of $\mathbb{U}$ is of:

- *dilation type* if it has a fixed point in $\mathbb{U}$;
- *hyperbolic type* if it has no fixed point in $\mathbb{U}$ and has derivative < 1 at its Denjoy-Wolff point;
- *parabolic type* if it has no fixed point in $\mathbb{U}$ and has derivative $= 1$ at its Denjoy-Wolff point.

As in the linear-fractional case, the maps of parabolic type fall into two subclasses:

- *Automorphic type:* Those with an orbit that's separated in the hyperbolic metric of $\mathbb{U}$.
- *Non-automorphic type:* Those for which *no* orbit is hyperbolically separated.

It can be shown that either all orbits are separated or none are separated (for a special case of this, see [7, §4]). With these ideas in hand we can state the full-strength version of the:

4.10. **Linear-Fractional Model Theorem.** *Suppose φ is a univalent holomorphic self-map of $\mathbb{U}$. Then there exists a holomorphic univalent map $\sigma : \mathbb{U} \to \mathbb{C}$ and a linear-fractional map ψ such that $\psi(\mathbb{U}) \subset \mathbb{U}$, $\psi(\sigma(\mathbb{U})) \subset \sigma(\mathbb{U})$), and*

$$\sigma \circ \varphi = \psi \circ \sigma. \tag{6}$$

Furthermore:

(a) *ψ, viewed as a self-map of $\mathbb{U}$, has the same type as φ.*
(b) *If φ is of hyperbolic type then ψ may be taken to be a conformal automorphism of $\mathbb{U}$.*
(c) *If φ is of either hyperbolic or parabolic-automorphic type, then σ may be taken to be a self-map of $\mathbb{U}$.*

We call the pair (ψ, G) (or, equivalently, (ψ, σ)) a *linear-fractional model* for φ).

The fact that ψ maps the simply-connected domain $G = \sigma(\mathbb{U})$ into itself follows immediately from the functional equation (6). This equation establishes a conjugacy between the original map φ acting on the unit disc and the linear-fractional map ψ acting on G. Since the action of ψ is known, the subtleties of φ lie encoded in the geometry of G.

History of the LFM Theorem. The Linear-Fractional Model Theorem is the work of a number of authors, whose efforts stretch over nearly a century. The dilation case is due to Koenigs ([23]: 1884). In this case equation (6) is *Schröder's equation:* $\sigma \circ \varphi = \lambda\sigma$, where (necessarily) $\lambda = \varphi'(0)$ (see [35, Chapter 6] for more details). The hyperbolic case is due to Valiron. If one replaces the unit disc by the right half-plane, sending the Denjoy-Wolff point to ∞, then the resulting functional equation is again Schröder's equation, but this time λ is the reciprocal of the angular derivative of the original disc map at the Denjoy-Wolff point ([38]: 1931). Finally the parabolic cases were established by Baker and Pommerenke ([28, 2]: 1979), and independently by Carl Cowen ([10]: 1981). Once again the situation is best viewed in the right half-plane, rather than the unit disc, with the Denjoy-Wolff point placed at ∞. Then equation (6) is just $\sigma \circ \varphi = \sigma + i$ in the automorphic case [28], and $\sigma \circ \varphi = \sigma + 1$ in the nonautomorphic case [2]. In [10] Cowen unified the proof of the Linear-Fractional Model Theorem by means of a Riemann-surface construction that disposes of all the cases in one stroke (see also [13, Theorem 2.53]). He later introduced linear-fractional models into the study of composition operators, using them to investigate spectra [11]. These models have also figured prominently in previously-mentioned work on subnormality [12] and compactness [36].

Distinguishing the parabolic models. The problem of distinguishing the two parabolic cases of the Linear-Fractional Model Theorem is, in general, quite delicate. It's shown in [7, §4] that if φ has enough differentiability at the Denjoy-Wolff point, then cases are distinguished by the second derivative of φ at that point. There is, however, some subtlety here; it's shown in [7, §6] that, for example, C^2-differentiability at the Denjoy-Wolff point is not enough to allow the second derivative to distinguish the two cases.

Necessity of Univalence. Although we have stated the Linear-Fractional Model Theorem only for univalent maps φ, the result is true even if φ is not univalent, provided we are willing to give up the conclusion of univalence for the intertwining map σ. (In case φ is of dilation type, with fixed point $p \in \mathbb{U}$, we must also assume that $\varphi'(p) \neq 0$.)

5. Why hypercyclicity is interesting

We've observed hypercyclic phenomena in some interesting classes of operators—weighted shifts and composition operators. There is much more to say about composition operators (see §6), but right now I'd like to shift gears and discuss some of the functional analytic aspects of hypercyclicity.

Suppose, for example, that you are an operator theorist interested in invariant subspaces. Then you are interested in *cyclicity,* so why bother with hypercyclicity, except that it is a formally stronger concept than the one you want to study? One answer is that if an operator is hypercyclic then it will, in general, have a far greater proliferation of cyclic vectors than one that is merely cyclic. Indeed, we have already noted that for a hypercyclic operator the collection of hypercyclic vectors is a dense G_δ set. By contrast, *the cyclic vectors for a cyclic operator need not be dense.*

Here is an example. On $H(\mathbb{U})$ let M_z denote the operator of "multiplication by the independent variable z", i.e.,

$$M_z f(z) = z f(z) \qquad (f \in H(\mathbb{U}) \text{ and } z \in \mathbb{U})$$

(admittedly, there is some abuse of notation here).

5.1. **Proposition.** *$f \in H(\mathbb{U})$ is a cyclic vector for M_z if and only if f has no zero in $\mathbb{U}$.*

Before proving this result, lets note that by Hurwitz's Theorem (an immediate consequence of the argument principle) any limit in $H(\mathbb{U})$ of a sequence of never-vanishing holomorphic functions is either identically zero or never-vanishing itself. Thus Proposition 5.1 has this consequence:

5.2. **Corollary.** *The operator M_z on $H(\mathbb{U})$ is cyclic, but its collection of cyclic vectors is not dense.*

Proof of Proposition 5.1. If f vanishes at some point of $\mathbb{U}$ then so does everything in $\text{orb}\,(M_z, f)$, hence so does anything in the closure of this orbit. So the orbit-closure can't be dense, i.e., f can't be cyclic (note that this argument works for any plane domain).

Conversely, if f vanishes nowhere on $\mathbb{U}$ then $1/f \in H(\mathbb{U})$, and so by power-series convergence, there is a sequence of polynomials $\{p_n\}$ that converges in $H(\mathbb{U})$ (i.e., uniformly on compact subsets of $\mathbb{U}$) to $1/f$. Thus the sequence $\{p_n f\}$, which is contained in $\text{span}\,\{\text{orb}\,(M_z, f)\}$, converges in $H(\mathbb{U})$ to 1, and so 1 belongs to the closure of the linear span of $\text{orb}\,(M_z, f)$. Now the aforementioned linear span is M_z-invariant, hence so is its closure, and so every polynomial belongs to this closure. In other words, the linear span of $\text{orb}\,(M_z, f)$ is dense in $H(\mathbb{U})$, i.e., f is cyclic for M_z. $\square$

Note that the argument of the last paragraph works for any simply connected plane domain—all that's needed is the polynomials approximating $1/f$, and in this generality the power series convergence argument gives way to Runge's theorem. In fact this characterization of cyclic vectors holds in *any* plane domain, but now one needs to use the full strength of the fact that $H(G)$ is a topological algebra with identity. I leave this as an exercise for the interested reader, noting only that the key to success is the fact that a proper closed ideal of $H(G)$ is maximal if and only if it consists of all functions that vanish at a pre-assigned point of G [24, Proposition 13.9, page 110].

Recall that all of our hypercyclic examples so far have been chaotic, and therefore, by Theorem 1.15 these maps have sensitive dependence on initial conditions. In fact, for linear operators, hypercyclicity itself implies sensitive dependence.

5.3. **Proposition.** *Every hypercyclic operator on an F-space has sensitive dependence on initial conditions.*

Proof. Suppose T is hypercyclic on X and let $HC(T)$ denote the collection of hypercyclic vectors for T. Fix $x \in X$ and note that $HC(T)+x$ is a dense G_δ subset of X (because $HC(T)$ has this property, and translation by x is a homeomorphism of X). Thus any neighborhood of x contains a point y of $HC(T)+x$. Since $x-y$ is a hypercyclic vector for T, the orbit of x exhibits the appropriate divergence from the orbit of y. □

Exercise. *If a continuous linear operator on an F-space X is hypercyclic, then every vector in X is a sum of two hypercyclic vectors.*

So far it has been the sufficient condition 2.1 that has yielded hypercyclicity in all our examples. In the first two paragraphs of §3 we noted two ways in which this condition is more powerful than first advertised: (a) When properly rephrased it works for continuous homomorphisms of complete, metrizable, topological groups; (b) It provides, not just hypercyclicity, but a stronger property: *mixing.*

Question. *Does every mixing operator on an F-space obey the hypotheses of Theorem 2.1?*

I don't know the answer to this one, but by the end of this section I hope you'll agree that the question is a reasonable one. There is a corresponding question for hypercyclicity whose formulation depends on the observation that the proof of Theorem 2.1 (our sufficient condition for hypercyclicity) works just as well under much weakened hypotheses. Here is the theorem that this proof actually gives:

5.4. **The Hypercyclicity Criterion.** *Suppose T is a continuous linear transformation an F-space X, and that for some subsequence of positive integers $n(k) \nearrow \infty$:*

(a) *There exists a dense subset Y of X on which $\{T^{n(k)}\}$ converges to zero pointwise.*

(b) *There exists a dense subset Z of X and a sequence $\{S_k\}$ of mappings $S : Z \to X$ (not necessarily either continuous or linear) such that:*

  (i) *$\{T^{n(k)}S_k\}$ converges pointwise on Z to the identity map on Z,*

  (ii) *$\{S_k\}$ converges to zero pointwise on Z.*

Then T is hypercyclic on X.

Note that if T satisfies the hypercyclicity criterion then so does $T \oplus T$, i.e., in the terminology of §3.7, T is *weakly mixing*. It is an open question whether or not every hypercyclic operator on an F-space satisfies the hypotheses of the hypercyclicity criterion. However in this direction we have a striking recent result of Bès and Peris [4]:

5.5. **Theorem.** *Suppose T is a continuous linear transformation of a separable F-space X. Then T is weakly mixing on X (i.e., $T \oplus T$ is hypercyclic on $X \oplus X$) if and only if T satisfies the hypotheses of the hypercyclicity criterion.*

Proof. It's easy to see that if T satisfies the hypotheses of the hypercyclicity criterion on X, then the same is true of $T \oplus T$ on $X \oplus X$, hence T is weakly mixing. It's the *converse* that requires some work!

Suppose, then, that $T \oplus T$ is hypercyclic on $X \oplus X$. Fix a vector $(x, y) \in X$ that is hypercyclic for $T \oplus T$. We will verify the hypotheses of the hypercyclicity criterion for T with $Y = Z = \operatorname{orb}(T, x)$ (dense in X because x is a hypercyclic vector for T); the trick is to find the subsequence $\{n(k)\}$ and the approximate right-inverses S_k.

Since T is hypercyclic, its range is dense in X, hence the range of T^N is also dense for every positive integer N. From this it is easy to check that for each N the vector $(x, T^N y)$ is hypercyclic for $T \oplus T$. In particular, for each zero-neighborhood U in X there is a vector $u \in U$ such that (x, u) is hypercyclic for $T \oplus T$.

Let's denote the distance from a vector $v \in X$ to the origin by $\|v\|$ (this is just for notational convenience: unless X is a Banach space, $\|\cdot\|$ will generally *not* be a norm). By the work of the last paragraph, we can inductively choose a strictly increasing sequence $\{n(k)\}$ of positive integers and a sequence $\{u_k\}$ of vectors in X so that:

$$\|u_k\|,\ \|T^{n(k)}x\|,\ \text{and } \|T^{n(k)}u_k - x\| \text{ are all } < \frac{1}{k} \quad \forall\, k \in \mathbb{N}. \tag{7}$$

From the second of these conditions, $T^{n(k)} \to 0$ pointwise on Y. We define the approximate right inverses from $Z = Y \to X$ by setting

$$S_k(T^n x) = T^n u_k \qquad (k = 0, 1, 2, \dots).$$

The definition is well-made because x is a hypercyclic vector for T, hence the points in its orbit must be distinct (if two points of an orbit coincide, the entire orbit is eventually periodic, hence finite). Because $\|u_k\| \to 0$ it follows that $S_k \to 0$ pointwise on Y. Finally, for each $k, n \in \mathbb{N}$:

$$T^{n(k)} S_k(T^n x) = T^{n(k)+n} u_k = T^n T^{n(k)} u_k \to T^n x \text{ as } k \to \infty,$$

where the convergence on the right is provided by the last inequality of (7) above. Thus $T^{n(k)} S_k \to I$ pointwise on Y, as desired. $\square$

This proof has a curious consequence: *If an operator on an F-space satisfies the hypotheses of the hypercyclicity criterion, then it satisfies those hypotheses with $Y = Z$.*

5.6. **Hypercyclicity vs. mixing.** At this point it's appropriate to mention a striking example, due to Hector Salas [33] of *a bilateral weighted shift T on $\ell^2(\mathbb{Z})$ which, along with its adjoint, is hypercyclic.* By way of contrast, note that in the Rolewicz example λB where B is the backward shift on ℓ^2 and λ a scalar of modulus > 1, the (Banach space) adjoint can be identified with $\lambda^{-1} S$, where S is the forward shift, so T^* is, in this case, a strict contraction, hence not hypercyclic. This example, and others like it, made the Salas example seem quite surprising.

Furthermore, in Salas's example, the weighting coefficients of the shift T are real, so with respect to the standard orthonormal basis for $\ell^2(\mathbb{Z})$ the matrices for both T and T^* have only real entries. It follows from this and an intriguing unpublished result of James Deddens (see below) that $T \oplus T^*$ is not even cyclic! Now we saw in Proposition 3.6 that any direct sum of mixing transformations is again mixing, so in the Salas example, either T or T^* is not mixing. In particular:

5.7. **Theorem.** *There exist hypercyclic operators on Hilbert space that are not mixing, hence do not satisfy the hypotheses of our first sufficient condition, Theorem 2.1.*

Because the result of Deddens is striking, useful, easily proved, and unpublished, I'd like to end this section by giving it a proper statement and proof.

5.8. **Deddens's Theorem (1982).** *Suppose T is a bounded linear operator on a separable Hilbert space $\mathcal{H}$ whose matrix, with respect to some orthonormal basis of $\mathcal{H}$, consists entirely of real entries. Then $T \oplus T^*$ is not cyclic.*

Proof. Suppose $f, g \in \mathcal{H}$. We'll prove the theorem by writing down a vector in $\mathcal{H} \oplus \mathcal{H}$ that is orthogonal to the $T \oplus T^*$-orbit of (f, g). Let $\{e_n\}_0^\infty$ be the promised orthonormal basis for $\mathcal{H}$ relative to which all the matrix entries $\langle Te_n,\, e_m\rangle$ are real. Now f and g have representations $f = \sum a_n e_n$ and $g = \sum b_n e_n$ relative to this basis, with square-summable coefficient sequences $\{a_n\}$ and $\{b_n\}$ respectively. Therefore there exist vectors $\overline{f}$ and $\overline{g}$ in $\mathcal{H}$ defined by $\overline{f} = \sum \overline{a_n} e_n$, and $\overline{g} = \sum \overline{b_n} e_n$.

I claim that $(-\overline{g}, \overline{f})$ is orthogonal in $\mathcal{H} \oplus \mathcal{H}$ to the orbit of (f, g). Indeed, for each non-negative integer n:

$$\langle (T \oplus T^*)^n (f, g),\, (-\overline{g}, \overline{f})\rangle = \langle (T^n f, T^{*n} g),\, (-\overline{g}, \overline{f})\rangle$$

$$:= -\langle T^n f,\, \overline{g}\rangle + \langle T^{*n},\, \overline{f}\rangle = \langle -T^n f,\, \overline{g}\rangle + \langle g,\, T^n \overline{f}\rangle$$

$$= \langle -T^n f,\, \overline{g}\rangle + \overline{\langle T^n \overline{f},\, g\rangle} = \langle -T^n f,\, \overline{g}\rangle + \langle T^n f,\, \overline{g}\rangle$$

$$= 0,$$

where the next-to-last equality we finally use the fact that all the matrix coefficients $\langle Te_n,\, e_m\rangle$ are real. □

6. Composition operators on H^2

So far we've considered composition operators only on the full space $H(\mathbb{U})$ of functions holomorphic on the unit disc. Now I'd like to shift the scene to a more subtle setting: the *Hardy space* H^2, which is a subspace of $H(\mathbb{U})$ that, in its natural norm, is a Hilbert space. H^2 is arguably the best place to study the interaction between the theories of linear operators and analytic functions. The purpose of this section is to prepare the way for the next one, in which we'll study hypercyclicity for linear-fractionally induced composition operators on H^2, discovering in the process some interesting contrasts with the $H(\mathbb{U})$ case.

In this section I'll develop some basic properties of H^2 and prove that that every composition operator restricts to a continuous mapping of H^2 into itself. If you've already had an introduction to Hardy spaces and composition operators on them, skip this section. What's here comes almost *verbatim* from [34, Chapter 1].

6.1. **The Hardy space H^2.** For $f \in H(\mathbb{U})$ and every non-negative integer n, let $\hat{f}(n) = f^{(n)}(0)/n!$. Then the series $\sum_{n=0}^\infty \hat{f}(n) z^n$ is the Taylor series of f with center at the origin:

it converges uniformly on compact subsets of $\mathbb{U}$ to f. The *Hardy space* H^2 is the collection of functions $f \in H(\mathbb{U})$ with $\sum_{n=0}^{\infty} |\hat{f}(n)|^2 < \infty$.

We equip H^2 with the norm that is naturally associated with its definition:

$$\|f\| = \left(\sum_{n=0}^{\infty} |\hat{f}(n)|^2 \right)^{1/2}, \tag{8}$$

and note that this norm arises from the natural inner product

$$\langle f, g\rangle := \sum_{n=0}^{\infty} \hat{f}(n)\overline{\hat{g}(n)} \qquad (f, g \in H^2). \tag{9}$$

Let T be the "Taylor transformation" from H^2 into the sequence space ℓ^2 defined by $Tf = \hat{f}$. The mapping T is clearly linear, and from the definition of the H^2 norm, it is an isometry: $\|Tf\| = \|f\|$ for every $f \in H^2$.

6.2. **Proposition.** *T maps H^2 onto ℓ^2. In particular, H^2 is a Hilbert space in the inner product (9).*

Proof. Because square-summable sequences are bounded, a simple geometric series estimate shows that if the complex sequence $\vec{a} = \{a_n\}_0^\infty$ lies in ℓ^2, then the associated power series $\sum_{n=0}^{\infty} a_n z^n$ converges uniformly on compact subsets of $\mathbb{U}$ to an analytic function f. By the uniqueness of power series representations, $a_n = \hat{f}(n)$ for every n, hence $Tf = \vec{a}$, so $T(H^2) = \ell^2$. □

Thus H^2 is the sequence space ℓ^2, disguised as a space of analytic functions. Note in particular that:

6.3. **Corollary.** *The sequence of monomials $\{z^n : n = 0, 1, 2, \ldots\}$ is an orthonormal basis for H^2.*

Some properties of the functions in H^2 can be easily discerned from the definition of the space. Here is one.

6.4. **Growth Estimate.** *For every $f \in H^2$ and $z \in \mathbb{U}$, $|f(z)| \le \|f\|(1 - |z|^2)^{-1/2}$.*

Proof. Use successively the triangle inequality and the Cauchy-Schwarz Inequality on the power series representation for f:

$$\begin{aligned} |f(z)| &= |\sum_{n=0}^{\infty} \hat{f}(n) z^n| \leq \sum_{n=0}^{\infty} |\hat{f}(n)|\,|z|^n \\ &\leq \left(\sum_{n=0}^{\infty} |\hat{f}(n)|^2\right)^{1/2} \left(\sum_{n=0}^{\infty} |z|^{2n}\right)^{1/2} \\ &= \|f\| \frac{1}{(1-|z|^2)^{1/2}}. \end{aligned}$$

□

The exercise below shows that the exponent in the Growth Estimate is best possible.

Exercise. *For α real let $f_\alpha(z) = (1-z)^{-\alpha}$. Show that $f_\alpha \in H^2$ if and only if $\alpha < 1/2$.*

Suggestion: Use the Binomial theorem and Stirling's formula to show that $\hat{f_\alpha}(n) \approx n^{\alpha-1}$.

6.5. **Corollary.** *Convergence in H^2 implies uniform convergence on compact subsets of $\mathbb{U}$.*

Proof. Suppose $\{f_n\}$ is a sequence of functions in H^2, f is a function in H^2, and $\|f_n - f\| \to 0$. Our goal is to show that $f_n \to f$ uniformly on compact subsets of $\mathbb{U}$.

For this, suppose K is a compact subset of $\mathbb{U}$. Let $r = \max\{|z| : z \in K\}$. Then for $z \in K$, the Growth Estimate yields:

$$|f_n(z) - f(z)| \leq \frac{\|f_n - f\|}{(1-|z|^2)^{1/2}} \leq \frac{\|f_n - f\|}{(1-|r|^2)^{1/2}},$$

which shows that as $n \to \infty$,

$$\max_{z \in K} |f_n(z) - f(z)| \leq \frac{\|f_n - f\|}{(1-|r|^2)^{1/2}} \to 0,$$

i.e. that $f_n \to f$ uniformly on K. □

However some properties of H^2 do not follow easily from the definition. For example, is every bounded analytic function in H^2? In order to answer this question reasonably, we need a different description of the norm.

6.6. **Proposition.** *A function $f \in H(\mathbb{U})$ belongs to H^2 if and only if*

$$\lim_{r\to 1-} \frac{1}{2\pi}\int_{-\pi}^{\pi} |f(re^{it})|^2\, d\,\theta < \infty.$$

When this happens, the limit of integrals on the left is $\|f\|^2$.

Proof. The functions $e^{in\theta}$ form an orthonormal set in the space $L^2([0,2\pi])$, hence for each $0 \le r < 1$ the integral on the right is $\sum_{n=0}^{\infty} |\hat{f}(n)| r^{2n}$. The result now follows from the monotone convergence theorem. $\square$

It is now an easy matter to show that every bounded function in $H(\mathbb{U})$ belongs to H^2. In fact, we can do better. Let H^∞ denote the collection of bounded analytic functions on $\mathbb{U}$, and for $b \in H^\infty$ let $\|b\|_\infty = \sup\{|b(z)| : z \in \mathbb{U}\}$. The integral representation given above for the H^2 norm shows immediately:

6.7. **Proposition.** *If $b \in H^\infty$ and $f \in H^2$ then $bf \in H^2$ and $\|bf\| \le \|b\|_\infty \|f\|$.*

In particular, upon taking $f \equiv 1$ we obtain:

Corollary. *If $b \in H^\infty$ then $b \in H^2$ with $\|b\| \le \|b\|_\infty$.*

6.8. **Multiplication operators act on H^2.** Proposition 6.7 reveals an interesting class of linear transformations on H^2. For $b \in H^\infty$ let M_b denote the operator of (pointwise) multiplication by b. That is, $M_b f = bf$. Clearly M_b, when viewed as a mapping on all of $H(\mathbb{U})$, is linear (note that for this we don't need b to be bounded). According to Proposition 6.7, M_b maps H^2 into itself, with $\|M_b f\| \le \|b\|_\infty \|f\|$ for each $f \in H^2$, hence M_b is a bounded linear operator on H^2 with norm $\le \|b\|_\infty$. We call M_b the *multiplication operator* induced by b. The most famous of these is the one induced by the identity map $b(z) \equiv z$. If we identify H^2 with the sequence space ℓ^2 this mapping of "multiplication by z" gets revealed as the forward shift on ℓ^2, which appeared in previous sections as the right inverse of the backward shift.

Do Composition operators act on H^2?

This is not a trivial question. Suppose you have $f \in H^2$ and want to determine if $C_\varphi f \in H^2$. Using the definition of H^2 we would substitute $\varphi(z)$ for z in the power series expansion of f, expand the various powers of the power series of φ by the binomial theorem, and regroup the resulting double series to identify the new powers of z, which are now complicated numerical series involving the coefficients of f and those of the powers of φ. Done this way, there seems to be no reason why $C_\varphi f$ should be in H^2. A calculation using the alternate characterization of H^2 provided by Proposition 6.6 fares just as badly, since it raises the specter of an unpleasant, and possibly non-univalent, change of variable in an integral.

After these pessimistic observations, it is remarkable that composition operators *do* preserve the space H^2, and do so continuously. The key to this is the following result, proved by Littlewood and published in 1925.

6.9. **Littlewood's Subordination Theorem.** *Suppose φ is a holomorphic self-map of $\mathbb{U}$ and $\varphi(0) = 0$. Then C_φ is a contraction mapping on H^2.*

Proof. The proof is helped significantly by the *backward shift operator* B, defined on H^2 by

$$Bf(z) = \sum_{n=0}^{\infty} \hat{f}(n+1)z^n \qquad (f \in H^2).$$

The name comes from the fact that B shifts the power series coefficients of f one unit to the left, and drops off the constant term. Clearly, $\|Bf\| \le \|f\|$ for each $f \in H^2$, and one might expect this fact to play an important role in the proof, but surprisingly it does not! Only the following two identities are needed, and they hold for any $f \in H(\mathbb{U})$:

$$f(z) = f(0) + zBf(z) \qquad (z \in \mathbb{U}), \tag{10}$$

$$B^n f(0) = \hat{f}(n) \qquad (n = 0, 1, 2, \dots). \tag{11}$$

To begin the proof, suppose first that f is a (holomorphic) polynomial. Then $f \circ \varphi$ is bounded on $\mathbb{U}$, so by the work of the last section there is no doubt that it lies in H^2; the real issue is its *norm*.

We begin the norm estimate by substituting $\varphi(z)$ for z in (10) to obtain

$$f(\varphi(z)) = f(0) + \varphi(z)(Bf)(\varphi(z)) \qquad (z \in \mathbb{U}).$$

Let us rewrite this equation in the language of composition and multiplication operators:

$$C_\varphi f = f(0) + M_\varphi C_\varphi B f\,. \tag{12}$$

At this point, the assumption $\varphi(0) = 0$ makes its first (and only) appearance. It asserts that all the terms of the power series for φ have a common factor of z, hence the same is true for the second term on the right side of equation (12), rendering it orthogonal in H^2 to the constant function $f(0)$. Thus,

$$\|C_\varphi f\|^2 = |f(0)|^2 + \|M_\varphi C_\varphi Bf\|^2 \leq |f(0)|^2 + \|C_\varphi Bf)\|^2, \tag{13}$$

where the last inequality follows from Proposition 6.7 above (since $\|\varphi\|_\infty \leq 1$). Now successively substitute $Bf, B^2f, \cdots$ for f in (13) to obtain:

$$\begin{aligned}
\|C_\varphi Bf\|^2 &\leq |Bf(0)|^2 + \|C_\varphi B^2 f\|^2 \\
\|C_\varphi B^2 f\|^2 &\leq |B^2 f(0)|^2 + \|C_\varphi B^3 f\|^2 \\
&\vdots \qquad\qquad \vdots \\
\|C_\varphi B^n f\|^2 &\leq |B^n f(0)|^2 + \|C_\varphi B^{n+1} f\|^2.
\end{aligned}$$

Putting all these inequalities together, we get

$$\|C_\varphi f\|^2 \leq \sum_{k=0}^{n} |(B^k f)(0)|^2 + \|C_\varphi B^{n+1} f\|^2$$

for each non-negative integer n.

Now recall that f is a polynomial. If we choose n be the degree of f, then $B^{n+1}f = 0$, and this reduces the last inequality to

$$\|C_\varphi f\|^2 \leq \sum_{k=0}^{n} |(B^k f)(0)|^2 = \sum_{k=0}^{n} |\hat{f}(k)|^2 = \|f\|^2,$$

where the middle line comes from property (11) of the backward shift. This shows that C_φ is an H^2-norm contraction, at least on the vector space of holomorphic polynomials.

To finish the proof, suppose $f \in H^2$ is not a polynomial. Let $f_n(z) = \sum_{k=0}^{n} \hat{f}(k) z^k$, the n-th partial sum of the Taylor series of f. Then $f_n \to f$ in the norm of H^2, so by Corollary 6.5 $f_n \to f$ uniformly on compact subsets of $\mathbb{U}$, hence $f_n \circ \varphi \to f \circ \varphi$ in the same manner. It is clear that $\|f_n\| \leq \|f\|$, and we have just shown that $\|f_n \circ \varphi\| \leq \|f_n\|$. Thus for each fixed $0 < r < 1$ we have

$$\begin{aligned}
\frac{1}{2\pi} \int_{-\pi}^{\pi} |f_n(\varphi(re^{it}))|^2 \, d\theta &= \lim_{n\to\infty} \frac{1}{2\pi} \int_{-\pi}^{\pi} |f_n(\varphi(re^{it}))|^2 \, d\theta \\
&\leq \limsup_{n\to\infty} \|f_n \circ \varphi\| \leq \limsup_{n\to\infty} \|f_n\| \leq \|f\|.
\end{aligned}$$

To complete the proof, let r tend to 1, and appeal one last time to Proposition 6.6. $\square$

To prove that C_φ is bounded even when φ does not fix the origin, we need to study conformal automorphisms of $\mathbb{U}$ from a different point of view. For each point $p \in \mathbb{U}$, define the holomorphic function α_p on $\mathbb{U}$ by:

$$\alpha_p(z) := \frac{p - z}{1 - \overline{p}z}.$$

The map so defined belongs to $\mathrm{Aut}\,(\mathbb{U})$, interchanges p with the origin, and is its own inverse (see, for example, [32], §12.2–12.6, pp. 254–256). Write $p = \varphi(0)$. Then the holomorphic function $\psi = \alpha_p \circ \varphi$ takes $\mathbb{U}$ into itself and fixes the origin. By the self-inverse property of α_p we have $\varphi = \alpha_p \circ \psi$, and this translates into the operator equation $C_\varphi = C_\psi C_{\alpha_p}$. We have just seen that C_ψ maps H^2 into itself. Thus, the fact that C_φ does the same will follow from the first sentence of the next result.

6.10. **Lemma.** *For each $p \in \mathbb{U}$, C_{α_p} is a bounded linear operator on H^2, with*

$$\|C_{\alpha_p}\| \leq \left(\frac{1 + |p|}{1 - |p|}\right)^{\frac{1}{2}}.$$

Proof. Suppose first that f is holomorphic in a neighborhood of the closed unit disc, say in $R\mathbb{U} = \{|z| < R\}$ for some $R > 1$. Then the limit in formula (6.6) can be passed inside the integral sign, with the result that

$$\|f\|^2 = \frac{1}{2\pi}\int_{-\pi}^{\pi} |f(e^{it})|^2 d\,\theta. \tag{14}$$

This opens the door to a simple change of variable in which the self-inverse property of α_p figures prominently:

$$\begin{aligned}
\|f \circ \alpha_p\|^2 &= \frac{1}{2\pi}\int_{-\pi}^{\pi} |f(\alpha_p(e^{it}))|^2 d\,\theta \\
&= \frac{1}{2\pi}\int_{-\pi}^{\pi} |f(e^{it})|^2 |\alpha_p'(e^{it})| dt \\
&= \frac{1}{2\pi}\int_{-\pi}^{\pi} |f(e^{it})|^2 \; \frac{1 - |p|^2}{|1 - \bar{p}e^{it}|^2} dt \\
&\leq \frac{1 - |p|^2}{(1 - |p|)^2} \cdot \left(\frac{1}{2\pi}\int_{-\pi}^{\pi} |f(e^{it})|^2 dt\right) \\
&= \frac{1 + |p|}{1 - |p|} \cdot \|f\|^2.
\end{aligned}$$

Thus the desired inequality holds for all functions holomorphic in $R\mathbb{U}$; in particular it holds for polynomials. It remains only to transfer the result to the rest of H^2, and for this we simply repeat the argument used to finish the proof of Littlewood's Subordination Theorem. □

At this point we have assembled everything we need to show that composition operators map H^2 into itself.

6.11. **Theorem.** *Suppose φ is a holomorphic self-map of $\mathbb{U}$. Then C_φ is a bounded linear operator on H^2, and*

$$\|C_\varphi\| \leq \sqrt{\frac{1+|\varphi(0)|}{1-|\varphi(0)|}}.$$

Proof. As outlined earlier, we have $C_\varphi = C_\psi C_{\alpha_p}$, where $p = \varphi(0)$, and ψ fixes the origin. Since each of the operators on the right-hand side of this equation sends H^2 into itself, the same is true of C_φ.

As for the inequality, this follows from Lemmas 6.9 and 6.10. I leave the details to you.

□

7. Hypercyclic composition operators on H^2

Now it's time to consider hypercyclicity for composition operators on H^2. The argument that showed such operators can only be induced by fixed-point-free univalent (holomorphic) self-maps of $\mathbb{U}$ works again to give the same result for H^2. So we begin, as before, with linear fractional maps, once again looking initially at the automorphisms. In dealing with composition operators on $H(\mathbb{U})$ we could cavalierly map the unit disc to other domains, ultimately identifying our operators as quasi-factors of translation operators on $H(\mathbb{C})$. this doesn't work in H^2, where the situation is a lot more rigid, so we have to take our stand pretty much in the unit disc. The material of this section follows very closely that of [34, §7.1, §7.2].

7.1. **Theorem.** *Suppose $\varphi \in \mathrm{Aut}\,(\mathbb{U})$ fixes no point of $\mathbb{U}$. Then C_φ is hypercyclic on H^2.*

Proof. As discussed in §4.3, the automorphism φ, being non-elliptic, has a unique attractive fixed point $\alpha \in \partial\mathbb{U}$. If there is another fixed point β, then this too must lie on the unit circle since it is the attractive fixed point for the inverse of φ, which is again an automorphism of

$\mathbb{U}$. Suppose first that we are dealing with the case of two fixed points. We will produce the cast of characters required for the hypothesis of the Theorem 2.1.

Let Y denote the set of functions that are continuous on the closed unit disc, analytic on the interior, and which vanish at α. We claim that $C_\varphi^n \to 0$ on Y. For this, note that for every $\zeta \in \partial\mathbb{U}\backslash\{\beta\}$ we have $\varphi_n(\zeta) \to \alpha$, hence if $f \in Y$ then $f(\varphi_n(\zeta)) \to f(\alpha) = 0$. Since f and φ_n are continuous, we can use boundary integral representation (14) of the H^2 norm, the Lebesgue Bounded Convergence Theorem, yields the desired result:

$$\|C_\varphi^n f\|^2 = \frac{1}{2\pi}\int_{-\pi}^{\pi} |f(\varphi_n(e^{it}))|^2 d\,\theta \to 0 \qquad (n \to \infty).$$

There are several ways to see that Y is dense in H^2. Here is one based on elementary Hilbert space theory. Suppose $f \in H^2$ is orthogonal to Y. Then for every non-negative integer n, the polynomial $z^{n+1} - \alpha z^n$ belongs to Y, so it is orthogonal to f:

$$0 =< f, z^{n+1} - \alpha z^n >= \hat{f}(n+1) - \bar{\alpha}\hat{f}(n).$$

It follows upon iterating this identity that $\hat{f}(n) = \bar{\alpha}^n \hat{f}(0)$ for each n. Since α has modulus one, all the Taylor coefficients of f have the same modulus, and since $f \in H^2$, these coefficients must all be zero. Thus the only H^2 function orthogonal to Y is the zero function. Since Y is a linear subspace of H^2, it must therefore be dense.

We note for further reference that the only property required here of α is that it lie outside of $\mathbb{U}$; the argument actually shows:

If $\alpha \notin \mathbb{U}$ then the set of polynomials that vanish at α is dense in H^2.

To finish the proof let $S = C_\varphi^{-1} = C_{\varphi^{-1}}$. As noted above, φ^{-1} is also an automorphism of the disc, with attracting fixed point β (the repulsive fixed point of φ). So if we take Z to be the set of continuous functions on the disc that are holomorphic in the interior and vanish at β, then S maps Z into itself, and the previous arguments apply to show that Z is dense and $S^n \to 0$ on Z. The hypotheses of Theorem 2.1 are therefore satisfied, so C_φ is hypercyclic, indeed, even mixing, on H^2.

The case where φ has just one fixed point is even easier; take Y as before, and set $Z = Y$. I leave the details to you. $\square$

7.2. **The Linear Fractional Hypercyclicity Theorem.** *Suppose that $\varphi \in LFT(\mathbb{U})$ has no fixed point in $\mathbb{U}$. Then:*

(a) *C_φ is hypercyclic on H^2 unless φ is a parabolic non-automorphism.*

(b) *If φ is a parabolic non-automorphism, then C_φ fails to be hypercyclic in a very strong sense: Only constant functions can be limit points of C_φ orbits.*

Proof of (a). We have already proved the result for automorphisms, so it remains to do it for hyperbolic non-automorphisms. Let φ be such a map, and suppose α and β are its fixed points, with α the attractive one. As before we seek to find the dense sets Y and Z, and the map S that will satisfy the hypotheses of Theorem 2.1.

The space Y is exactly the one that worked in the automorphic hypercyclicity result, and it works again with no change in the argument. It is the space Z that requires some care.

Suppose first that the repulsive fixed point β lies on the line through the origin and α, but is on the other side of the origin from α. Let Δ be the disc whose boundary is the circle perpendicular to this line that passes through α and β, so now $\mathbb{U}$ is inside Δ and $\partial\mathbb{U}$ is tangent to $\partial\Delta$ at α.

Since φ fixes α and β, and preserves angles, it maps the boundary of Δ onto itself, and therefore takes Δ onto either itself or the exterior of $\partial\Delta$. But φ takes $\mathbb{U}$ into itself, so the latter possibility is ruled out. Therefore φ is a conformal automorphism of Δ.

Let Z be the collection of functions that are continuous on $\overline{\Delta}$, analytic on Δ, and which vanish at β. As we noted above, the fact that β lies outside $\mathbb{U}$ insures that the polynomials that vanish at β form a dense subset of H^2, thus Z is dense. Define the map $S : Z \to Z$ by

$$Sf(z) = f(\varphi^{-1}(z)) \qquad (z \in \mathbb{U}).$$

The fact that $\varphi^{-1}(z)$ is not always in $\mathbb{U}$ is of no importance here, nor is the fact that S is neither defined nor bounded on H^2. What is important is that $\varphi^{-1}(\Delta) \subset \Delta$, that S is defined on Z, and that $C_\varphi S$ is the identity on Z, all of which is obvious. In addition, the fact that $\varphi^{-n}(\zeta) \to \beta$ for each $\zeta \in \partial\mathbb{U}$ yields, precisely as in the last section, that $S^n \to 0$ on Y. Thus the hypotheses of the Theorem 2.1 are again satisfied, so C_φ is hypercyclic on H^2.

If the repulsive fixed point β is not in the required position (it could even be at ∞ for example), then there is a conformal automorphism γ_α of $\mathbb{U}$ that fixes α and takes β to the desired position. (This is a simple exercise: instead of the unit disc, work in the right half-plane, with ∞ in place of α. An appropriate affine map gives the desired automorphism of

$\mathbb{P}$). Then $\varphi = \gamma_\alpha \circ \psi \circ \gamma_\alpha^{-1}$, where ψ is a linear fractional self-map of $\mathbb{U}$ that has its fixed points arranged properly, so C_ψ is hypercyclic, and C_φ is similar to C_ψ, hence also hypercyclic. $\square$

Remarks. Once it has been observed that φ is an automorphism of the larger disc Δ, a more elegant line of argument suggests itself. Define the space $H^2(\Delta)$ in some obvious way, show that it is a dense subspace of H^2 and has a stronger topology. Then use the Automorphism Theorem to conclude that C_φ is hypercyclic on $H^2(\Delta)$, and transfer this hypercyclicity to H^2 by noting that C_φ on H^2 is a quasi-factor of C_φ on $H^2(\Delta)$.

It is also interesting to investigate why the proof given above for hyperbolic non-automorphisms does not work for parabolic ones. The point is that in the hyperbolic case the big disc Δ is precisely the union of the successive inverse images of $\mathbb{U}$ under φ:

$$\Delta = \bigcup_n \varphi^{-n}(\mathbb{U}).$$

If, on the other hand, φ is a *parabolic* non-automorphism, then this union turns out to be $\hat{\mathbb{C}}\backslash\{\alpha\}$ (as is easily seen by representing the map as a translation of the right half-plane strictly into itself), hence the set Z defined in the proof above contains only the zero function.

Proof of (b). Now we assume that $\varphi \in \mathrm{LFT}(\mathbb{U})$ is a parabolic non-automorphism, so it has only one fixed point in $\hat{\mathbb{C}}$, and this lies on $\partial\mathbb{U}$. Without loss of generality we may take this fixed point to be $+1$ (otherwise conjugate φ by an appropriate rotation to produce a similar composition operator induced by a parabolic automorphism with fixed point at $+1$).

We return to an earlier idea, seeking to understand parabolic self-maps of $\mathbb{U}$ by mapping the unit disc to the right half-plane, where the parabolic map becomes a translation. Since our particular map fixes the point 1, we use the transformation $w = (1+z)/(1-z)$, which takes our the original map φ to an LFT Φ that maps $\mathbb{P}$ into itself and fixes only the point at ∞. Thus Φ is translation of the right half-plane, $\Phi(w) = w + a$ where necessarily $\mathrm{Re}\, a > 0$ ($\mathrm{Re}\, a$ is ≥ 0 because $\Phi(\mathbb{P}) \subset \mathbb{P}$, and > 0 because Φ is not an automorphism of $\mathbb{P}$). Similarly, the n-th iterate φ_n of φ gets transformed into "translation by na:" $\Phi_n(w) = w + na$.

Our proof will depend on knowing how quickly the φ-orbits of points in $\mathbb{U}$ get close to each other, and to the attractive fixed point $+1$. Suppose z in $\mathbb{U}$ and w is the corresponding point in $\mathbb{P}$, so

$$w = \frac{z+1}{z-1} \quad \text{and} \quad z = \frac{w-1}{w+1}\ .$$

Then:

$$1 - z = \frac{2}{w+1} \quad \text{and} \quad 1 - |z|^2 = \frac{4\mathrm{Re}\, w}{|w+1|^2}$$

It follows, then, that if $z \in \mathbb{U}$, w is the point of $\mathbb{P}$ corresponding to $\varphi(z)$, and w_0 that of $\varphi(0)$, then $w + na$ corresponds to $\varphi_n(z)$ and $w_0 + na$ to $\varphi_n(0)$. Thus:

$$1 - |\varphi_n(z)|^2 = \frac{4(\mathrm{Re}\, w + n\mathrm{Re}\, a)}{|w+na+1|^2}$$

and

$$\varphi_n(z) - \varphi_n(0) = \frac{2(w_0 - w)}{(w_0 + na + 1)(w + na + 1)} \ ,$$

from which follows:

$$\lim_{n\to\infty} n(1 - |\varphi_n(z)|^2) = c_1 \quad \text{and} \quad \lim_{n\to\infty} n^2|\varphi_n(z) - \varphi_n(z_0)| = c_2, \tag{15}$$

where c_1 and c_2 are non-zero constants that depend on z and a.

Now fix $f \in H^2$. Our goal is to show that if the orbit of f under C_φ clusters at some $g \in H^2$, then g must be a constant function. For this we need a growth estimate on *differences* of functional values that is analogous the one obtained in Growth Estimate 6.4 for the values themselves. We begin with the derivative. For $z \in \mathbb{U}$,

$$\begin{aligned} |f'(z)|^2 &= \left|\sum_{n=1}^{\infty} n\hat{f}(n) z^{n-1}\right|^2 \\ &\leq \left(\sum_{n=1}^{\infty} |\hat{f}(n)|^2\right)\left(\sum_{n=1}^{\infty} n^2 |z|^{2(n-1)}\right) \\ &\leq \|f\|^2 \frac{2}{(1-|z|^2)^3} \ . \end{aligned}$$

Upon taking square roots on both sides of the last inequality, we get this growth estimate on the derivative of f:

$$|f'(z)| \leq \sqrt{2}\, \frac{\|f\|}{(1-|z|^2)^{3/2}} \qquad (z \in \mathbb{U}).$$

To get the desired estimate on differences, suppose $z, w \in \mathbb{U}$ and $|z| \leq |w|$. To estimate $f(z) - f(w)$ we integrate f' over the line segment joining z and w, and use the inequality above:

$$
\begin{aligned}
|f(z)-f(w)| &\le \int_z^w |f'(\zeta)||d\zeta| \\
&\le \sqrt{2}\,\|f\| \int_z^w \frac{|d\zeta|}{(1-|\zeta|^2)^{3/2}} \\
&\le \sqrt{2}\;\|f\|\,\frac{|w-z|}{(1-|w|^2)^{3/2}}.
\end{aligned}
$$

Thus for each pair of points $z, w \in \mathbb{U}$,

$$
|f(z)-f(w)| \le \sqrt{2}\,\|f\|\,\frac{|w-z|}{(\min\{1-|w|, 1-|z|\})^{3/2}}. \tag{16}
$$

In (16) above, substitute $\varphi_n(z)$ for z, and $\varphi_n(0)$ for w, and use the estimates of (15) above; the result is

$$
\begin{aligned}
|f(\varphi_n(z)) - f(\varphi_n(0))| &\le \text{const.}\, \frac{|\varphi_n(z)-\varphi_n(0)|}{(\min\{1-|\varphi_n(z)|, 1-|\varphi_n(0)|\})^{3/2}} \\
&\le \text{const.}\, \frac{n^{-2}}{n^{-3/2}} = \frac{\text{const.}}{\sqrt{n}},
\end{aligned}
$$

where the constant in each line depends on f, z, and φ, but *not* on n. Thus

$$
\lim_n [f(\varphi_n(z)) - f(\varphi_n(0))] = 0 \qquad (z \in \mathbb{U}). \tag{17}
$$

To finish the argument, suppose $g \in H^2$ is a cluster point of the C_φ - orbit of f. Then for some sequence $n_k \nearrow \infty$ we have $f \circ \varphi_{n_k} \to g$ in the norm of H^2, and therefore pointwise on $\mathbb{U}$. By (17) this implies

$$
g(z) - g(0) = \lim_k [f(\varphi_{n_k}(z)) - f(\varphi_{n_k}(0))] = 0,
$$

hence $g \equiv g(0)$. Thus only constant functions can be limit points of the C_φ-orbit of an H^2 function. □

Hypercyclicity for more general composition operators on H^2

Suppose φ is a holomorphic self-map of $\mathbb{U}$, with linear fractional model (ψ, G). As usual, denote by σ the Riemann map of $\mathbb{U}$ onto G. The basic principle behind operating here is:

> *If the polynomials in σ are dense in H^2, then cyclic behavior for C_ψ on H^2 can be transferred to cyclic behavior for C_φ.*

For example, suppose C_ψ is hypercyclic on H^2, so that ψ is a holomorphic self-map of $\mathbb{U}$ that's either hyperbolic or a parabolic automorphism. According to the Linear Fractional

Model Theorem, we may take $G \subset \mathbb{U}$, so if we let $V : H(\mathbb{U}) \to H(\mathbb{U})$ denote C_σ acting on the restrictions to G of functions in H^2, we can interpret the functional equation $\psi \circ \sigma = \sigma \circ \varphi$ as asserting that $VC_\psi = C_\varphi V$. Now our hypothesis on σ is that V has dense range in H^2, so C_φ *is a quasi-factor of* C_ψ and therefore C_φ inherits the hypercyclicity of C_φ. The same holds for other concepts such as cyclicity, chaos, mixing.

One way to insure that the polynomials in σ are dense in H^2 is to employ

7.3. **Walsh's Theorem.** *Suppose G is a simply connected domain whose boundary is a Jordan curve. Let the holomorphic function σ map $\mathbb{U}$ univalently onto G. Then the polynomials in σ are dense in H^2.*

The result that is usually called Walsh's Theorem actually asserts that the polynomials in z are uniformly dense in $A(G)$, the subalgebra of $C(\overline{G})$ consisting of functions holomorphic on G (see, for example, [26, Theorem 3.9, page 98]). A theorem of Carathéodory asserts that F extends continuously and univalently to $\overline{G}$, so Walsh's original result asserts, in our situation, that the polynomials in F are dense in $A(\mathbb{U})$, which is clearly dense in H^2, and this yields Theorem 7.3 (see [34, §8.1] for more details).

The main point of the monograph [7] is to find conditions on a univalent holomorphic self-map φ of $\mathbb{U}$ which guarantee that G is the interior of a Jordan curve. The main result shows that:

> *If the Denjoy-Wolff point ω of φ lies on $\partial\mathbb{U}$, and*
>
> * *the closure of $\varphi(\mathbb{U})$ touches $\partial\mathbb{U}$ only at ω, and*
>
> * *φ has "sufficient differentiability" at ω,*
>
> *then the composition operator induced on H^2 by φ has the same hypercyclic behavior as its linear fractional model, i.e., C_φ is hypercyclic (indeed chaotic and mixing) if φ is of hyperbolic or parabolic automorphic type, and not hypercyclic if φ is of parabolic non-automorphic type.*

Here "sufficient differentiability" varies from case to case, but C^4 works for all of them. Of course the negative result about parabolic non-automorphic type maps cannot be deduced from regularity of the model. Instead, enough differentiability is assumed at the Denjoy-Wolff point to allow the estimates that worked in the linear fractional case to be used for the more general one. For specific references to these results, see Table II on page 12 of

[7]. To give you a feeling for how the arguments go, I present below the hyperbolic case, taken almost word-for-word from [34, §8.3, pp. 134–137]. In order to simplify notation, let's agree that a *Jordan domain* is a simply connected plane domain whose boundary is a Jordan curve.

7.4. **A Valiron-type Theorem.** *Suppose φ is a univalent holomorphic self-map of $\mathbb{U}$ of hyperbolic type, with C^2-smoothness at its Denjoy-Wolff point ω. Suppose further that the closure of $\varphi(\mathbb{U})$ lies in $\mathbb{U}\cup\{\omega\}$. Then there exists a hyperbolic $\psi \in LFT(\mathbb{U})$ and a holomorphic univalent map σ of $\mathbb{U}$ onto a Jordan domain contained in $\mathbb{U}$, such that $\sigma \circ \varphi = \psi \circ \sigma$.*

Proof. We may without loss of generality suppose that the Denjoy-Wolff point is $+1$. Write $\lambda = \varphi'(1)$, let H denote the half-plane $\{\operatorname{Re} z < 1\}$, and set $\psi(z) = \lambda z + (1-\lambda)$, so ψ is a hyperbolic automorphism of H. We'll first find a mapping σ with all the required properties, except that its image will lie in H. The theorem as stated will follow upon mapping H conformally onto $\mathbb{U}$.

We'll be able to copy the original Koenigs argument almost word-for-word if we use the map $z \to 1-z$ to map $\mathbb{U}$ onto the open disc $\mathbb{U}_0$ of radius one and center $+1$. Let's still use the notation φ for the resulting self-map of $\mathbb{U}$, which now fixes the origin, at which it is assumed to be C^2-smooth. This means that φ has the "finite Taylor expansion"

$$\varphi(z) = \lambda z + z^2 B(z) \qquad (z \in \overline{\mathbb{U}}_0) \tag{18}$$

where $0 < \lambda = \varphi'(0) < 1$, and the function B is bounded on the closure of $\mathbb{U}_0$. The idea is to resurrect Koenigs's original proof for the interior fixed point case, but where Koenigs used the Schwarz Lemma, we will employ an estimate derived from (18). We will obtain a solution σ of Schröder's equation $\sigma \circ \varphi = \lambda\sigma$ on $\mathbb{U}_0$, with σ a *Jordan map* with *positive real part.*

We proceed, just as did Koenigs, obtaining σ as a limit of normalized iterates $\sigma_n = \lambda^{-n}\varphi_n$. Note first that since φ has positive real part, so does each of its iterates, and therefore so does each map σ_n. We claim that the sequence $\{\sigma_n\}$ converges uniformly on $\overline{\mathbb{U}}_0$.

For this, let $\beta = \max\{|B(z)| : z \in \overline{\mathbb{U}}_0\}$, and let Δ be the intersection of $\overline{\mathbb{U}}_0$ with the closed disc of radius $(1-\lambda)/2\beta$ centered at the origin. By (18) above we have

$$|\varphi(z)| \le \left(\frac{1+\lambda}{2}\right)|z| \qquad (z \in \Delta). \tag{19}$$

Since $(1+\lambda)/2 < 1$, this last inequality insures that $\varphi(\Delta) \subset \Delta$, so inequality (19) can be iterated for each $z \in \Delta$ to yield

$$|\varphi_n(z)| \leq \left(\frac{1+\lambda}{2}\right) |\varphi_{n-1}(z)| \leq \left(\frac{1+\lambda}{2}\right)^2 |\varphi_{n-2}(z)| \cdots \leq \left(\frac{1+\lambda}{2}\right)^n |z|.$$

By our definition of Δ, this last estimate shows that

$$|\varphi_n(z)| \leq \frac{1-\lambda}{2\beta} \left(\frac{1+\lambda}{2}\right)^n \qquad (z \in \Delta) \tag{20}$$

for each non-negative integer n. Note that the origin is now playing the role of the Denjoy-Wolff point for φ, so $\varphi_n \to 0$ uniformly on compact subsets of $\mathbb{U}_0$. We are assuming that φ takes $\overline{\mathbb{U}}_0$ into $\mathbb{U}_0 \cup \{0\}$, so the closure of $\varphi(\overline{\mathbb{U}}_0) \backslash \Delta$ in $\mathbb{U}_0$ is compact. Thus $\varphi_n \to 0$ *uniformly on* $\overline{\mathbb{U}}_0$. In particular, there is a positive integer N such that

$$\varphi_n(\overline{\mathbb{U}}_0) \subset \Delta \qquad (n \geq N).$$

Following Koenigs, we set

$$F(z) = \frac{\varphi(z)}{\lambda z} \qquad (z \in \overline{\mathbb{U}}_0),$$

and note that for each $z \in \overline{\mathbb{U}}_0$ the expansion (18) implies

$$|1 - F(z)| = \lambda^{-1}|z||B(z)| \leq \lambda^{-1}\beta|z|. \tag{21}$$

Now fix $z \in \overline{\mathbb{U}}_0$. If $j \geq N$ then $\varphi_j(z) \in \Delta$, so using respectively (21) and (20) above (with $j - N$ in place of n and $\varphi_N(z)$, which belongs to Δ, in place of z), we obtain

$$|1 - F(\varphi_j(z))| \leq \frac{\beta}{\lambda}|\varphi_{j-N}(\varphi_N(z))| \leq \frac{1-\lambda}{2\lambda}\left(\frac{1+\lambda}{2}\right)^{j-N}$$

for each $z \in \overline{\mathbb{U}}_0$. Since N is independent of the point $z \in \overline{\mathbb{U}}_0$, this last inequality shows that each term of the infinite series $\sum |1 - F(\varphi_j(z))|$ is dominated by the corresponding term of a convergent geometric series. Thus $\sum |1 - F(\varphi_j(z))|$ converges uniformly on $\overline{\mathbb{U}}_0$, by the Weierstrass M-test. Since this convergence is passed on to infinite product

$$\prod_{j=0}^{\infty} F(\varphi_j(z)) = z^{-1} \lim_{n\to\infty} \sigma_n(z),$$

the sequence $\{\sigma_n\}$ therefore converges uniformly on $\overline{\mathbb{U}}_0$ to a function σ that fixes the origin, is continuous on $\overline{\mathbb{U}}_0$, is holomorphic and univalent on $\mathbb{U}_0$, and obeys Schröder's equation on $\overline{\mathbb{U}}_0$. Furthermore, σ has positive real part on $\mathbb{U}_0$ since it is non-constant there and, as we noted above, each σ_n has positive real part.

Thus, in order to show that $\sigma(\overline{\mathbb{U}}_0)$ is a Jordan domain, it only remains to check that σ is univalent on $\partial\mathbb{U}$. But this follows from Schröder's equation and the univalence of φ on $\overline{\mathbb{U}}_0$. The argument is this: If $\sigma(z_1) = \sigma(z_2)$ for a pair of points $z_1, z_2 \in \partial\mathbb{U}_0$, then upon multiplying by λ and using Schröder's equation (which we have just proved holds on the *closed* disc) we see that $\sigma(\varphi(z_1)) = \sigma(\varphi(z_2))$. If neither z_1 nor z_2 is zero, then both φ-images belong to $\mathbb{U}_0$, on which we know σ is univalent. Thus $\varphi(z_1) = \varphi(z_2)$, so $z_1 = z_2$ since φ is assumed to be univalent on $\overline{\mathbb{U}}_0$.

Suppose on the other hand that one of the original points, say z_1 is zero. Then Schröder's equation and the fact that $\sigma(0) = 0$ yield $\sigma(\varphi(z_2)) = 0$. But if $z_2 \neq 0$, then $\varphi(z_2) \in \mathbb{U}_0$, contradicting the fact that $\operatorname{Re}\sigma > 0$ on $\mathbb{U}_0$. Thus $z_2 = 0$, so σ is univalent on $\overline{\mathbb{U}}_0$.

In summary, we have shown that there is a continuous, univalent map σ defined on $\overline{\mathbb{U}}_0$ that has non-negative real part, is holomorphic on $\mathbb{U}_0$, and satisfies Schröder's equation $\sigma\circ\varphi = \lambda\sigma$ on $\overline{\mathbb{U}}_0$.

Upon transferring this result back to the unit disc by means of the map $z \mapsto 1 - z$ our accomplishment looks like this:

> *If φ obeys the hypotheses of the Theorem, then it has a Jordan model (ψ, σ), where $\psi(z) = \lambda z + (1 - \lambda)$, and σ maps $\mathbb{U}$ into the half-plane $\{\operatorname{Re} z < 1\}$.*

The only problem remaining is that G need not lie in $\mathbb{U}$, but this is easily remedied. Let τ be a linear fractional transformation that takes the half-plane $\{\operatorname{Re} z < 1\}$ onto the unit disc, and fixes $+1$. Necessarily $\tau \in \mathrm{LFT}(\mathbb{U})$. Define:

- $\widetilde{\psi} = \tau\circ\psi\circ\tau^{-1}$, another member of $\mathrm{LFT}(\mathbb{U})$ with attractive fixed point at $+1$,
- $\widetilde{G} = \tau(G)$, a Jordan sub-domain of $\mathbb{U}$, and
- $\widetilde{\sigma} = \tau\circ\sigma$, a univalent mapping of $\mathbb{U}$ onto $\widetilde{G}$.

Since ψ is hyperbolic, so is $\widetilde{\psi}$, and $\widetilde{\sigma}\circ\varphi = \widetilde{\psi}\circ\widetilde{\sigma}$, so $(\widetilde{\psi}, \widetilde{G})$ is the desired linear fractional model for φ. □

8. Why hypercyclicity is VERY interesting

In this section we show that hypercyclic operators have special properties not possessed by general transitive mappings.

We observed in the first section of these notes that nontrivial transitive mappings might have non-transitive squares: in fact, the nontrivial permutation map of the discrete metric space $\{1, 2\}$ onto itself as this property, as does every product of this map with a transitive one. With this construction we see that even continuous group homomorphisms can be transitive without having transitive squares. Carol Kitai, in her 1982 dissertation ([22], Remark 2.13) asked if this could happen in the linear setting, and in 1995 Shamim Ansari gave a striking argument to show that that it cannot:

8.1. **Ansari's Theorem** [1]. *If T is a hypercyclic operator on a metrizable topological vector space X, then T^n is hypercyclic for any positive integer n.*

Here is a way to think about trying to prove Ansari's Theorem: Let x be a hypercyclic vector for T, and fix an integer $n > 1$. Then

$$\operatorname{orb}(T, x) = \operatorname{orb}(T^n, x) \cup \operatorname{orb}(T^n, Tx) \cdots \cup \operatorname{orb}(T^n, T^{n-1}x), \tag{22}$$

Now in any topological space, if a finite union of sets is dense, at least one of the sets must be somewhere dense (i.e., its closure must contain a nonempty open set). To see why this is so, we may assume the collection of sets with dense union is *minimal*, i.e., if we remove one them, then union of what remains is not dense. So remove one of the sets in this minimal collection. The closure of what remains misses a nonvoid open subset of the space, and this open subset must therefore belong to the closure of what was removed. So that removed set is somewhere dense. Returning to our hypercyclic situation, (22) therefore guarantees that least one of the sets $\operatorname{orb}(T^n, T^k x)$ is somewhere dense. Thus Ansari's Theorem will be proved if we can show that for linear operators, somewhere dense orbits are everywhere dense. Note that our examples of (nonlinear) transitive maps with nontransitive squares also show that, in the general situation, somewhere dense orbits need not be dense.

Here is another open question, this one posed by Domingo Herrero in 1992 [20]: Suppose $\{x_1, x_2, \ldots x_N\}$ is a finite subset of an F-space X, and $\bigcup_{j=1}^{n} \operatorname{orb}(T, x_j)$ is dense in X; is T hypercyclic on X? Herrero's question was just recently answered in the affirmative, independently by George Costakis and Alfredo Peris. To get a nice statement, let's call an operator for which a finite union of orbits is dense "multihypercyclic."

8.2. **The Costakis-Peris Theorem** [9, 27]. *Every multi-hypercyclic operator is hypercyclic.*

Note that, like Ansari's theorem, this result would follow easily if we could prove that for linear operators, somewhere dense orbits are everywhere dense. Peris in [27] posed this problem explicitly, and within the last few months Paul Bourdon and Nathan Feldman solved it affirmatively, thus providing a unified proof of both Ansari's Theorem and the Costakis-Peris Theorem.

8.3. **The Bourdon-Feldman Theorem** [6]. *Suppose T is a continuous linear operator on a locally convex F-space X, and $x \in X$. If* $\mathrm{orb}\,(T, x)$ *is somewhere dense in X, then it is dense in X.*

Proof. Actually, Bourdon and Feldman prove their result for any locally convex (Hausdorff) topological vector space over the complex scalars. You'll see below that neither completeness nor metrizability ever plays an essential role in the argument. However the Hahn-Banach theorem enters at one point in the proof, so local convexity is required. I don't know if this result is true in topological vector spaces that are not locally convex.

For convenience we'll employ the following notation throughout the proof:

- $\mathrm{orb}\,(T, x)$ will henceforth be abbreviated just $\mathrm{orb}\,(x)$.
- The closure of $\mathrm{orb}\,(x)$ will be denoted $\mathrm{clorb}\,(x)$.
- The interior of the $\mathrm{clorb}\,(x)$ will be denoted $\mathrm{clorb}^\circ\,(x)$.
- $\mathcal{P}$ is the collection of (holomorphic) polynomials with complex coefficients. If $\mathcal{S} \subset \mathcal{P}$ and $y \in X$ then $\mathcal{S}(T) = \{p(T) : p \in \mathcal{S}\}$, and $\mathcal{S}(T)y = \{p(T)y : p \in \mathcal{S}\}$.

In this notation, to say that a vector $y \in X$ *cyclic* for an operator T on X is to assert that $\mathcal{P}(T)y$ is dense in X.

The proof will be broken up into five steps. Throughout, we assume $x \in X$ has a somewhere dense orbit, i.e., $\mathrm{clorb}^\circ\,(x) \neq \emptyset$ (although in Step I this will not be used).

Step I. *If* $y \in \mathrm{orb}\,(x)$ *then* $\mathrm{clorb}^\circ\,(y) = \mathrm{clorb}^\circ\,(x)$.

Proof. That $\mathrm{clorb}^\circ\,(y) \subset \mathrm{clorb}^\circ\,(x)$ follows from the corresponding set containment for orbits. Now $\mathrm{clorb}\,(x)$ differs from $\mathrm{clorb}\,(y)$ by just a finite set of isolated points, from which follows the reverse containment. □

STEP II. *Each element of* orb (x) *is a cyclic vector for* T.

Proof. Suppose $y \in \operatorname{orb}(x)$. By Step I, $\operatorname{orb}(y)$ is somewhere dense, and since $\operatorname{orb}(y) \subset \operatorname{orb}(x) \subset \mathcal{P}(T)x$ we see that $\mathcal{P}(T)x$ is somewhere dense. Now $\mathcal{P}(T)x$ is a vector space, hence so is its closure. Since this closure contains an open set, it is the whole space. Thus y is cyclic for T. □

The next step provides the crucial element of the argument. To put it in perspective, note that $\operatorname{clorb}(x)$ is T-invariant.

STEP III. *The complement in* X *of* $\operatorname{clorb}^\circ(x)$ *is* T*-invariant.*

Proof. By Step I we may replace x by any element of its orbit without disturbing $\operatorname{clorb}^\circ(x)$. Now $\operatorname{clorb}^\circ(x)$ is a nonempty open set, each point of which is a limit point of $\operatorname{orb}(x)$, so some point of $\operatorname{orb}(x)$ belongs to $\operatorname{clorb}^\circ(x)$. Replace x by this point. In other words:

Without loss of generality we may henceforth assume that $x \in \operatorname{clorb}^\circ(x)$.

Suppose, in order to reach a contradiction, that $X \backslash \operatorname{clorb}^\circ(x)$ is *not* T-invariant, i.e., that for some $y \notin \operatorname{clorb}^\circ(x)$ we have $Ty \in \operatorname{clorb}^\circ(x)$. Actually:

We may assume that $y \notin \operatorname{clorb}(x)$.

Indeed, if we're unlucky and y is in $\operatorname{clorb}(x)$ then (since it's not in $\operatorname{clorb}^\circ(x)$) it must be on the boundary of $\operatorname{clorb}(x)$. In particular, there is a point $y' \notin \operatorname{clorb}(x)$, but close enough to y that Ty' is close enough to Ty to keep it in the open set $\operatorname{clorb}^\circ(x)$ (we use the continuity of T here). Then rename y' as y.

We may also assume that $y = p(T)$ *for some* $p \in \mathcal{P} \backslash \{0\}$.

This is a similar argument: x is cyclic for T (Step I), i.e., $\mathcal{P}(T)x$ is dense in X, so we can find $p \in \mathcal{P}$ so that $p(T)x$ is close enough to y that it lies in the (open) complement of $\operatorname{clorb}(x)$, and so that its T-image stays in $\operatorname{clorb}^\circ(x)$ (continuity of T, and open-ness of $\operatorname{clorb}^\circ(x)$ again). Note that p is not the zero polynomial, since $p(T)x \neq Tp(T)x$.

Now because $\operatorname{clorb}(x)$ is T-invariant, and contains $Tp(T)x$,

$$\operatorname{clorb}(x) \supset T^n p(T)x = p(T)T^{n+1}x \qquad (n = 0, 1, 2, \dots),$$

hence $\operatorname{clorb}(x) \supset p(T)(\operatorname{orb}(Tx))$. Upon taking closures:

$$\operatorname{clorb}(x) \supset p(T)\operatorname{clorb}(Tx) \supset p(T)\operatorname{clorb}^\circ(Tx) = p(T)\operatorname{clorb}^\circ(x), \tag{23}$$

(the last equality following from Step I). But $x \in \operatorname{clorb}^\circ(x)$ (recall that we showed earlier that there was no loss of generality in assuming this), so by (23) above, $p(T)x \in \operatorname{clorb}(x)$. But this is a contradiction; we have chosen p so that $p(T)x \notin \operatorname{clorb}(x)$. □

Step IV. $\mathcal{P}(T)$ has dense range for every $p \in \mathcal{P}\backslash\{0\}$.

Proof. This is the only place where we use local convexity and the fact that the scalar field is $\mathbb{C}$. Fix $p \in \mathcal{P}\backslash\{0\}$ and factor it into linear factors: $p(z) = (z-\alpha_1)(z-\alpha_2)\dots(z-\alpha_n)$, where $\alpha_1, \dots, \alpha_n$ are complex numbers. Then $p(T)$ has a similar decomposition into linear factors $T - \alpha_j I$, so it is enough to show that each such factor has dense range.

For this, fix $\alpha \in \mathbb{C}$ and suppose the range of $T - \alpha I$ is not dense in X. Then, thanks to local convexity, the Hahn-Banach Theorem provides a continuous linear functional Λ on X that annihilates $\operatorname{ran} T - \alpha I$, but is not identically zero on X. Now $\Lambda \circ (T - \alpha I) = 0$, so $\Lambda \circ T = \alpha\Lambda$. We are saying, of course, that α is an eigenvalue of T^*, with eigenvector Λ. Now just as in the proof of Theorem 2.7 we have:

$$\Lambda(\operatorname{orb}(x)) = \{\alpha^n \Lambda(x) : n = 0, 1, 2, \dots\}. \tag{24}$$

In this case $\Lambda(\operatorname{clorb}^\circ(x))$ is a nonvoid open subset of $\mathbb{C}$ (continuous linear functionals are open maps), hence $\Lambda(\operatorname{orb}(x))$ is somewhere dense in $\mathbb{C}$. But it's a simple exercise to check that $\{\alpha^n \Lambda(x)\}$ is *nowhere dense* in $\mathbb{C}$, regardless of the value of $\Lambda(x)$, so thanks to (24) we have arrived at a contradiction. This all began with the assumption that $\operatorname{ran}(T - \alpha I)$ is not dense in X, so that couldn't have been correct. □

Step V. *Completion of the proof.* We are assuming that $\operatorname{clorb}^\circ(x) \neq \emptyset$ and want to show that $\operatorname{orb}(x)$ is dense, i.e., that $\operatorname{clorb}(x) = X$.

So suppose not. Recall that x is cyclic for T (Step II), so $\mathcal{P}(T)x$ is dense in X, and therefore one can find a subcollection $\mathcal{Q} \subset \mathcal{P}\backslash\{0\}$ of polynomials such that $\mathcal{Q}(T)x$ is a dense subset of the nonvoid open set $X\backslash\operatorname{clorb}(x)$, and therefore a dense subset of $X\backslash\operatorname{clorb}^\circ(x)$. We showed in Step III that this latter set is T-invariant, so $\mathcal{Q}(T)(\operatorname{orb}(x)) \subset X\backslash\operatorname{clorb}^\circ(x)$, hence by continuity of T:

$$X\backslash\operatorname{clorb}^\circ(x) \supset \overline{\mathcal{Q}(T)(\operatorname{orb}(x))} \supset \mathcal{Q}(T)(\operatorname{clorb}(x)) \tag{25}$$

For convenience, write $\mathcal{P}'$ for $\mathcal{P}\backslash\{0\}$.

Claim: $p \in \mathcal{P}' \Rightarrow p(T)x \notin \partial \mathrm{clorb}^{\circ}(x)$.

Suppose we have proved this Claim. The vector subspace $\mathcal{P}(T)x$, being dense in X, is infinite dimensional, so there is no doubt that $\mathcal{P}'(T)$ is connected.[3]

This connected set is the disjoint union of two subsets:

$$G = (\mathcal{P}'(T)x) \cap \mathrm{clorb}^{\circ}(x) \quad \text{and} \quad H = (\mathcal{P}'(T)x) \cap (X \backslash \mathrm{clorb}^{\circ}(x)).$$

Clearly G is relatively open in $\mathcal{P}'(T)x$, and thanks to the Claim, so is H! Furthermore, neither G nor H is empty: $x \in G$, and $\mathcal{Q}(T)x \subset H$. This contradicts the connectedness of $\mathcal{P}'(T)x$, and finishes the proof.

It remains to prove the Claim. Suppose for the sake of contradiction that $p(T)x \in \partial \mathrm{clorb}^{\circ}(x)$ for some $p \in \mathcal{P}'(T)$. Consider the set

$$D := \mathrm{clorb}^{\circ}(x) \cup \mathcal{Q}(T)x,$$

which is dense in X because the collection $\mathcal{Q}$ of polynomials has been chosen to make $\mathcal{Q}(T)$ dense in the complement of $\mathrm{clorb}^{\circ}(x)$. Now

$$p(T)D = p(T)\mathrm{clorb}^{\circ}(x) \cup p(T)\mathcal{Q}(T)x,$$

and by the T-invariance of $X \backslash \mathrm{clorb}^{\circ}(x)$ (Step III), the second term on the right lies in $X \backslash \mathrm{clorb}^{\circ}(x)$. But so does the first term on the right! Indeed, $p(T)x \notin \mathrm{clorb}^{\circ}(x)$ (because it's assumed to be in $\partial \mathrm{clorb}^{\circ}(x)$), so by a now-familiar argument, $p(T)\mathrm{clorb}^{\circ}(x) \subset X \backslash \mathrm{clorb}^{\circ}(x)$.

Thus $p(T)D$ lies entirely in $X \backslash \mathrm{clorb}^{\circ}(x)$, so is disjoint from the nonvoid open set $\mathrm{clorb}^{\circ}(x)$, thus contradicting the density of D. This completes the proof of the Claim, and with it, the proof of the Bourdon-Feldman Theorem. □

The Bourdon-Feldman Theorem is a beautiful piece of work, but don't let my introduction to it fool you into thinking it makes trivial the results of Ansari and Costakis-Peris. The proof of the Bourdon-Feldman Theorem has many elements that went into proving the previous two results. As we mentioned earlier, the original question it answers was asked by Peris in [27]; furthermore the phenomena of cyclicity and connectedness play an important role in Ansari's proof [1].

[3]Since we work with complex scalars here, all that's really needed is that the vector space have dimension $\neq 0$, however the argument above shows that this part of the argument works even for real scalars.

References

[1] S. I. Ansari, *Hypercyclic and cyclic vectors*, J. Functional Analysis 128 (1995), 374–383.
[2] I. N. Baker and Ch. Pommerenke, *On the iteration of analytic functions in a half-plane II*, J. London Math. Soc. (2) 20 (1979), 255–258.
[3] J. Banks, J. Brooks, G. Cairns, G. Davis, and P. Stacey, *On Devaney's definition of chaos*, American Math. Monthly 99 (1992), 332–334.
[4] J. Bès and A. Peris, *Hereditarily hypercyclic operators,* J. Functional Analysis 167 (1999), 94–112.
[5] G. D. Birkhoff, *Démonstration d'un théoreme elementaire sur les fonctions entieres,* C.R. Acad. Sci. Paris 189 (1929), 473–475.
[6] P. S. Bourdon and N. S. Feldman, *Somewhere dense orbits are everywhere dense,* preprint, Washington and Lee University, January 2001.
[7] P. S. Bourdon and J. H. Shapiro, *Cyclic phenomena for composition operators,* Memoirs Amer. Math. Soc. #596, American Mathematical Society, Providence, R.I., 1997.
[8] K. C. Chan and J. H. Shapiro, *The cyclic behavior of translation operators on Hilbert spaces of entire functions*, Indiana Univ. Math. J. 40 (1991) 1421–1449.
[9] G. Costakis, *On a conjecture of D. Herrero concerning hypercyclic operators*, C. R. Acad. Scil. Paris 330 (2000), 179–182.
[10] C. C. Cowen, *Iteration and the solution of functional equations for functions analytic in the unit disc*, Trans. Amer. Math. Soc. 265 (1981), 69–95.
[11] C. C. Cowen, *Composition operators on* H^2, J. Operator Th. 9 (1983),77–106.
[12] C. C. Cowen and T. L. Kriete III, *Subnormality and composition operators on* H^2, J. Funct. Anal. 81 (1988), 298-319.
[13] C. C. Cowen and B. D. MacCluer, *Composition Operators on Spaces of Analytic Functions*, CRC Press, Boca Raton, 1995.
[14] A. Denjoy, *Sur l'itération des fonctions analytiques*, C.R. Acad. Sci. Paris 182 (1926), 255–257.
[15] R.L. Devaney, *An Introduction to Chaotic Dynamical Systems*, second ed., Addison-Wesley, Reading, Mass., 1989.
[16] P. Enflo, *On the invariant subspace problem for Banach spaces,* Acta Math. 158 (1987), 213–313.
[17] R.M. Gethner and J.H. Shapiro, *Universal vectors for operators on spaces of holomorphic functions,* Proc. Amer. Math. Soc. 100 (1987), 281–288.
[18] G. Godefroy and J. H. Shapiro, *Operators with dense, invariant, cyclic vector manifolds*, J. Funct. Anal. 98 (1991), 229–269.
[19] C. G. Grosse-Erdmann, *Universal families and hypercyclic operators*, Bull . Amer. Math. Soc. 36 (1999) 345–381.
[20] D. A. Herrero, *Hypercyclic operators and chaos*, J. Operator Theory 28 (1992), 93–103.
[21] D. A. Herrero and Z. Wang, *Compact perturbations of hypercyclic and supercyclic operators*, Indiana Univ. Math. J. 39 (1990), 819-830.
[22] C. Kitai, *Invariant closed sets for linear operators*, Thesis, Univ. of Toronto, 1982.
[23] G. Koenigs, *Recherches sur les intégrales de certaines équationes functionelles*, Annales Ecole Normale Superior (3) 1 (1884), Supplément, 3–41.
[24] D. H. Luecking and L. A. Rubel, *Complex Analysis, A Functional Analysis Approach,* Springer-Verlag 1984.
[25] G. R. MacLane, *Sequences of derivatives and normal families,* J. D'Analyse Math. 2 (1952), 72–87.
[26] A. I. Markushevich, *Theory of Functions of a Complex Variable*, Vol III, Prentice Hall, 1967.
[27] A. Peris, *Multihypercyclic operators are hypercyclic,* Math. Z. 236 (2001), 779–786.
[28] Ch. Pommerenke, *On the iteration of analytic functions in a half-plane I*, J. London Math Soc. (2) 19 (1979), 439–447.
[29] C. Read, *On the invariant subspace problem for Banach spaces,* Israel J. Math. 63 (1988), 1–40.
[30] C. Robinson, *Dynamical Systems,* CRC Press, 1995.
[31] S. Rolewicz, *On orbits of elements,* Studia Math. 33 (1969), 17–22.
[32] W. Rudin, *Real and Complex Analysis, Third Edition,* McGraw-Hill, New York, 1987.
[33] H. Salas, *A hypercyclic operator whose adjoint is also hypercyclic,* Proc. Amer. Math. Soc. 112 (1991), 765–770.

[34] J. H. Shapiro, *Composition Operators and Classical Function Theory,* Springer-Verlag, 1993.

[35] J. H. Shapiro, *Simple Connectivity and linear chaos*, Rend. Circ. Mat. Palermo, Supplemento Serie II, Numero 56 (1998), 27–48.

[36] J. H. Shapiro, W. Smith, and D. A. Stegenga, *Geometric models and compactness of composition operators*, J. Functional Analysis 127 (1995), 21–62.

[37] W. Seidel and J. Walsh, *On approximation by Euclidean and non-Euclidean translates of an analytic function,* Bull. Amer. Math. Soc. 47 (1941), 916–920.

[38] G. Valiron, *Sur l'iteration des fonctions holomorphes dans un demi-plan*, Bull des Sci. Math. (2) 55 (1931), 105–128.

Michigan State University, East Lansing, MI 48824-1027, USA

E-mail address: shapiro@math.msu.edu