Douglas E. Norton

The fundamental theorem of dynamical systems

Commentationes Mathematicae Universitatis Carolinae, Vol. 36 (1995), No. 3, 585--597

Persistent URL: http://dml.cz/dmlcz/118787

The fundamental theorem of dynamical systems

Douglas E. Norton

Abstract. We propose the title of The Fundamental Theorem of Dynamical Systems for a theorem of Charles Conley concerning the decomposition of spaces on which dynamical systems are defined. First, we briefly set the context and state the theorem. After some definitions and preliminary results, based both on Conley's work and modifications to it, we present a sketch of a proof of the result in the setting of the iteration of continuous functions on compact metric spaces. Finally, we claim that this theorem should be called The Fundamental Theorem of Dynamical Systems.

Keywords: chain recurrent set, attractor, decomposition
Classification: 58F12, 58F25, 26A18

1. Introduction

The story of the study of Dynamical Systems has a little bit of something for everybody. As a field of study, it is both as old as Newton's invention of Calculus and as current as a variety of brand new journals appearing now in the offerings of your local purveyor of mathematical literature. As an area of mathematical exploration, it is as abstract as cohomology and as applied as computer experimentation. Its tools range from traditional techniques of classical analysis to various branches of topology born in the twentieth century at least partially in response to some Dynamical Systems questions. It crosses interdisciplinary boundaries all over the map, from ecology to psychology to meteorology, from chemical kinetics to population genetics, from economics to mechanics Newtonian, Lagrangian, and Hamiltonian. It ranges from elegant proof of abstract theorems to colorful computer graphics plastered on T-shirts. Given such a broad and far-reaching definition of the context, can any result be appropriately dubbed the Fundamental Theorem of Dynamical Systems? This paper presents a fundamental result about some basic definitions of Dynamical Systems. This result is not new by any means, but it may not be widely known outside the circle of practitioners of the study of Dynamical Systems, and it may not be fully appreciated within that circle.

2. Context

The consideration of a result as "fundamental" to the study of Dynamical Systems seems premature without at least some discussion of what we mean by "Dynamical Systems".

A dynamical system consists of three ingredients: a setting in which the dynamical behavior takes place, such as the real line, a torus, or a locally compact metric space; a mathematical rule or description which specifies the evolution of the system; and an initial condition or state from which the system starts. The basic questions of Dynamical Systems are qualitative ones, dating back to the introduction of geometrical and topological tools into the study of differential equations by Poincaré. For a particular initial state, what happens to the system in the long run? How does "what happens" depend on the initial condition? How does it depend on the form of the mathematical description of the evolution of the system? On parameters within the description? On the properties of the space on which the system is defined? That is a fairly simplified description of what Dynamical Systems is all about, but it is pretty accurate. Given a particular setup and a mathematical description of how things are going to proceed, what happens?

The mathematical descriptions of how things are to proceed fall generally into two categories: discrete and continuous. A discrete dynamical system may be represented by a recursively defined sequence of numbers or points. The Fibonacci sequence and cellular automata are examples in which the recursive definition involves more than evaluating a single-variable function. Iterating a function is the standard discrete dynamical system, in which the function is applied to each term of the sequence to give the next term.

The other broad category of dynamical systems, continuous ones, is perhaps best represented by differential equations, certainly in light of the historical development. Associated closely with differential equations are *flows*, whose very name suggests the continuous nature of their definition. The dynamical systems in these two broad categories are linked by similar dynamic behaviors and mathematical analyses. This had led to earlier and earlier introduction of Dynamical Systems topics into the curriculum, in which students can discover much of the interesting dynamical behavior exhibited by more mathematically sophisticated systems without the prerequisites of differential topology, algebraic topology, or differential geometry. As we present some definitions and the theorem, we shall keep this continuous-discrete dichotomy in mind.

3. The theorem

Charles Conley presents in his CBMS monograph *Isolated Invariants Sets and the Morse Index* [6] some significant results about invariant sets, attractor-repeller pairs, chain recurrence, Morse decompositions, and index theory (now called Conley Index Theory) in the setting of flows on compact metric spaces. One of the principal ideas in this work and a fundamental result in the study of Dynamical Systems is the Conley Decomposition Theorem. That result can be stated very simply:

Theorem 1. *Any flow on a compact metric space decomposes into a chain recurrent part and a gradient-like part.*

On the other side of that discrete-continuous Dynamical Systems dichotomy, the result reads as follows:

Theorem 2. *The iteration of a continuous function on a compact metric space decomposes the space into a chain recurrent part and a gradient-like part.*

A point exhibits some sort of recurrent behavior when the dynamical system returns the point to itself, or to a neighborhood of itself, in a particular way. Chain recurrence is one type of recurrence with "errors" allowed along the orbit. The term gradient-like suggests a one-way behavior, a flowing downhill, mathematically quantifiable. In either setting, the theorem is clearly "fundamental" in the sense of being basic, elemental, even easy to understand. We need specific definitions of various terms, such as chain recurrent, gradient-like, and Lyapunov function, but the *idea* is simple, simply the following. A particular type of recurrent behavior which is worth consideration on its own merits as a descriptive tool actually allows us to break down the very essence of the dynamic nature of a dynamical system into two parts: points which exhibit that carefully chosen type of recurrent behavior, and points which instead travel strictly one-way in a fashion that can be nicely quantified by an appropriate mathematical description.

Although Conley's results are in the setting of flows, we acknowledge that the iteration of continuous functions on compact metric spaces is the increasingly more common setting for the introduction of Dynamical Systems to the novice, and we restrict our attention to that setting in the discussion to follow.

4. Some dynamical preliminaries

Except for changing some variables, the following definitions and corresponding notations are from [11]. Throughout the remainder of this exposition, X will represent a compact metric space, with $f : X \to X$ representing a continuous function whose iteration provides the dynamics on X.

The *forward orbit* of x_0 is the set of all forward iterates of x_0. The point x_0 is a *fixed point* for f if $f(x_0) = x_0$. The point x_0 is a *periodic point* for f if $f^n(x_0) = x_0$ for some $n > 0$, and its forward orbit is called a *periodic forward orbit*. If $x_n = x_0$ but $x_k \neq x_0$ for $0 < k < n$, then n is called the *period* of the forward orbit. Finally, an *orbit* is a bi-infinite sequence of points $(\dots, x_{-2}, x_{-1}, x_0, x_1, \dots)$ such that $f(x_n) = x_{n+1}$ for every integer n.

Fixed points are certainly the simplest examples of recurrent dynamical behavior: under iteration, a fixed point not only comes back, it never goes away. When looking at subsets of the space rather than one point at a time, the relevant concept is that of an *invariant set*. A subset $Y \subseteq X$ is called an *invariant set* if $f(Y) = Y$. This definition does not suggest that any individual points remain fixed; they may all do quite a bit of traveling under the iteration of the function, but their orbits remain within the subset in question. Also, no points are left out, in the sense that an invariant set differs from a *positively invariant set*, which satisfies $f(Y) \subset Y$, in which the inclusion may be as a proper subset. Notice an important distinction between flows and iterating functions. Since a flow is defined for all $\mathbf{R}$, negative time is built in from the start, so reversing the direction of the flow and talking about the "backwards flow" is natural. On the other hand, without some extra hypotheses on the function to be iterated (invertible, homeomorphism, diffeomorphism, etc.), reversing the iteration, or inverting the function, may be multi-valued or undefined. Many treatments of these ideas include those extra hypotheses because of the natural parallels with the case of flows. For some results with continuity the only assumption on f, see [14], [15].

One basic question in the study of dynamical systems is that of the ultimate fate of points, and one representation of that idea is the "omega limit set". Just as ω is at the end of the Greek alphabet, the omega limit set of a point or a set is where it ends up under the action of the dynamical system. For a subset Y of the space X, $\omega(Y)$ is defined to be the maximal invariant set in the closure of the destination of the set Y under the dynamics of the system.

Definition. For $Y \subset X$, let $\omega_n(Y) \equiv \bigcup_{k\geq n} f^k(Y)$ and $\omega(Y) \equiv \bigcap_{n\geq 0} \overline{\omega_n(Y)}$. Then $\omega(Y)$ is called the *omega limit set* of Y.

Since X is compact, if Y is nonempty, then $\omega(Y)$ is a nonempty compact set.

Consider the dynamical systems represented by the following diagrams, in which dots represent fixed points and arrows indicate the direction of the orbit of a point under iteration.

Example 1.

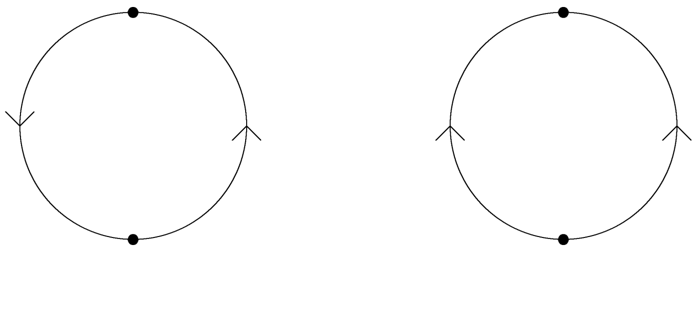

Figure 1

The phase portraits in (a) and (b) can be considered representations of the iteration of functions on the circle:

$$\text{(a)} \quad f(\theta) = \theta + \alpha\cos^2\theta, \qquad \text{(b)} \quad f(\theta) = \theta + \alpha\cos\theta,$$

where $\alpha < 1/2$. In both cases, $\theta = \pm\pi/2$ are fixed points, and each fixed point is its own omega limit set. Also in both cases, the omega limit set of any point in the right open half-circle is the fixed point $\theta = \pi/2$, so the union of the omega limit sets of all points in the right closed half-circle is the two point set $\theta = \pm\pi/2$.

However, the omega limit set of the right closed half-circle is the entire right closed half-circle. That is,

$$\bigcup_{\theta \in Y} \omega(\theta) \neq \omega(Y).$$

This is but one illustration of the utility of the definition of omega limit set for sets; the right closed half-circle is an invariant set in this example which is not reflected by a point-by-point omega limit set definition.

Omega limit sets will play an important role in the characterization of dynamics provided by the Fundamental Theorem.

5. Tools for the theorem

As motivation for the type of recurrence used in the theorem, consider this selection from the preface to [6]:

> The time evolution of some processes, such as the movement of the planets, is very accurately modeled by differential equations. Such accurate modeling requires the identification of a small number of measurable quantities, like the positions and velocities of the planets, whose behavior over the time span considered is almost independent of the neglected factors. There are other important processes, such as the fluctuation of animal populations, for which the identification of such quantities is not possible but where a rough relation between the more obvious variables and their "rates of change" may be evident . . . . If such rough questions are to be of use it is necessary to study them in rough terms, and that is the aim of these notes.

Chain recurrence is an incorporation of "roughness" or imprecision throughout the dynamic process, consideration of points which are periodic up to a slightly inexact following of the dynamics.

The modification of the idea of the orbit of a point which leads to just type of recurrence we need is called a *pseudo-orbit*. A pseudo-orbit is roughly an orbit with "errors" or deviations from the true orbits allowed at certain intervals of the dynamic process. For the iteration of a function, this deviation is possible at each stage of the iteration. For flows, the errors must be allowed only at intervals of time bounded away from zero; otherwise, the errors could accumulate in finite time to give behavior nothing at all like the original flow. The history of this approach is as simple as ABC, as in Anosov, Bowen, and Conley. The ideas were first developed by Conley in [5] and [6] and by Bowen in [3] and [4], based on related ideas of Anosov in [1]. We continue to restrict our attention to the iterating function case.

Measuring deviations from true orbits makes working in a metric space a natural choice. The definitions that follow, and even the main results ahead, can be stated in the language of topologies and open sets, avoiding explicit use of the metric; however, in this introduction to the topic, we stay with the metric, with computation and visualization our motivating concerns. Throughout the discussion that follows, let (X, d) denote a compact metric space.

Definition. Given $x, y \in X$ and $\varepsilon > 0$, an *ε-pseudo-orbit* from x to y means a sequence of points $(x = x_0, x_1, \ldots, x_n = y)$ with $n > 0$ such that for $k = 0, 1, \ldots, n-1$, we have $d(f(x_k), x_{k+1}) < \varepsilon$.

Another way to describe a pseudo-orbit is as a sequence of points that would be considered an orbit if the positions of the points were only known up to a given finite accuracy ε. See Figure 2 below for a simple representation of an *ε-pseudo-orbit* from x_0 to x_4.

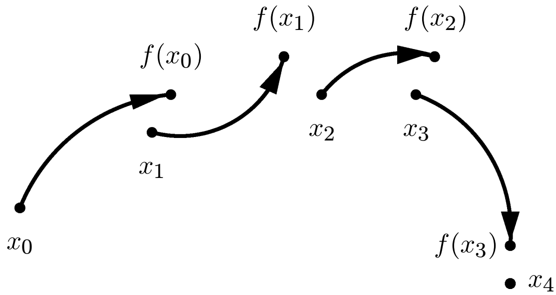

Figure 2

Some authors, including Conley in his development of some of these ideas, use the term "ε-chain" for ε-pseudo-orbit. Since we think of ε-pseudo-orbits as derived from orbits, we will use the pseudo-orbit terminology, although the word "chain" will be used extensively as a modifier for terms describing related structures and behaviors.

The idea of ε-pseudo-orbits can be used to describe a partial order on the set X.

Notation. For x and y in X, we will write $x > y$ to mean that for every $\varepsilon > 0$, there is an ε-pseudo-orbit from x to y; that is, for every $\varepsilon > 0$ there is an $n > 0$ and there are values for x_k so that $(x_0, x_1, \ldots, x_n)$ is an ε-pseudo-orbit from x to y.

Notice that $x > y$ is a statement *for all* $\varepsilon > 0$. If for an arbitrary choice of ε there is an ε-pseudo-orbit from x to y, then the relation $x > y$ holds. So $x > y$ is a rather strong relation between x and y. They are "almost" connected by an orbit, because the deviations from a true orbit can be chosen as small as desired. There remains a difference, however, between "$x > y$" and "y can be reached by an orbit from x". For example, consider Example 1 in the previous section. On the (b) side, there is no orbit from $\theta = -\pi/2$ to $\theta = +\pi/2$ because $\theta = -\pi/2$ is a fixed point of the map. No matter how small an ε-jump to some point a near

$\theta = -\pi/2$ is taken, the true forward orbit of a will end up within another ε-jump of $\theta = +\pi/2$, so we have $-\pi/2 > +\pi/2$ (in this partial order sense).

To describe recurrence in this context, we invoke the alternative terminology of chains and call the behavior chain recurrence. A point $x \in X$ is called *chain recurrent* if $x > x$. That is, x is chain recurrent if there is an ε-pseudo-orbit from x back to itself for any choice of $\varepsilon > 0$.

Figure 3 below is a simple representation of an ε-pseudo-orbit from x_0 back to itself.

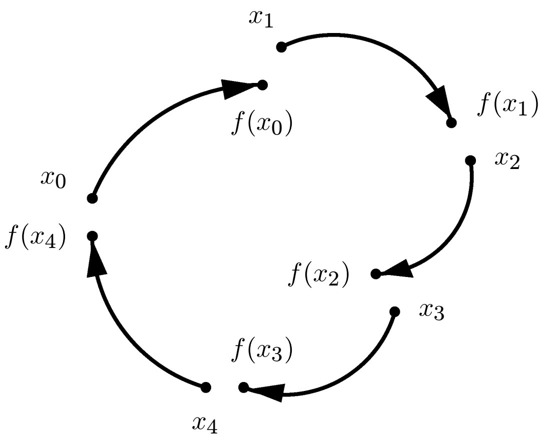

Figure 3

This does *not* mean that x_0 is necessarily chain recurrent; there is an ε-pseudo-orbit from x_0 back to itself for a particular fixed $\varepsilon > 0$, but there may not be one for other choices of $\varepsilon > 0$. Chain recurrence is a weaker form of recurrence than periodicity, but it is not so weak that just finding one ε-pseudo-orbit for one choice of ε will do the trick.

Notation. $R(X, f) \equiv$ *the chain recurrent set of f on X* $\equiv \{x \in X : x \text{ is chain recurrent}\}$.

When the context allows no ambiguity, both $R(X)$ and $R(f)$ are used as alternative notations for the chain recurrent set.

Since any true orbit is an ε-pseudo-orbit for any ε, any periodic point is chain recurrent. So $\mathrm{Per}(f) \subset R(f)$, where $\mathrm{Per}(f) \equiv \{x \in X : x \text{ is periodic}\}$.

Consider again Example 1 above. In part (a), $\theta = \pm\pi/2$ are fixed points. On the right half-circle, points move up toward the fixed point $\theta = \pi/2$, and on the left half-circle, points move down toward the fixed point $\theta = -\pi/2$. Then any point in the circle is chain recurrent: for any $\varepsilon > 0$, one can follow the orbit of a point to a distance less than ε from a fixed point, utilize an ε-jump to get past the fixed point, follow the orbit again, jump past the other fixed point, and continue to return to the starting value. So the chain recurrent set is the entire circle. In particular, $R(f) \neq \mathrm{Per}(f)$.

In part (b) of the example, $\theta = \pm\pi/2$ are again fixed points. On the right half-circle, points move up toward the fixed point $\theta = \pi/2$. This time, on the left half-circle, points also move up toward the fixed point $\theta = \pi/2$. Then only the fixed points are chain recurrent: for any other θ there is no ε-pseudo-orbit from θ back to itself for small enough ε. So in this case, $R(f) = \mathrm{Per}(f)$.

Consider briefly the structure of the chain recurrent set. The partial order $x > y$ generates an equivalence relation by defining $x \sim y$ to mean that $x > y$ and $y < x$. The relation "$\sim$" partitions $R(f)$ into equivalence classes, or *chain classes*. Each chain class is *chain transitive*: within each chain class, any two points have an ε-pseudo-orbit connecting them, for any $\varepsilon > 0$. Chain classes are by definition the maximal chain transitive sets in the sense that no chain class lies within a chain transitive set that is strictly larger than itself.

We note briefly a few properties of the chain recurrent set. It is closed. It is invariant. It contains all the periodic points, $\omega(x)$ for each $x \in X$, and points exhibiting some other forms of recurrence which we have not discussed here, such as the nonwandering set. (See [6].) Finally, the chain recurrent set satisfies a pair of properties which can be shown to be lacking in some other types of recurrent set. Any invariant set Y satisfies $R(R(Y)) = R(Y)$. Also, $R(\omega(x)) = \omega(x)$ for any $x \in Y$. For proofs of these and other properties of the chain recurrent set, in these and other settings, see, for example, [2], [6], [10], [14], [15].

The idea of a gradient-like dynamical system is an extension from gradient flows of the idea of functions that decrease on solutions, called *Lyapunov functions*. A dynamical system is called *gradient-like* if there is some continuous real-valued function which is strictly decreasing on nonconstant solutions. A system is called *strongly gradient-like* "if the chain recurrent set is totally disconnected (and consequently equal to the rest point set)" [6]. A system is called *chain recurrent* "if the chain recurrent set is the whole space" [6]. Note that strongly gradient-like and chain recurrent are the extremes, with strongly gradient-like a stronger statement than gradient-like. In particular, a dynamical system can be gradient-like but also chain recurrent; for an example, see [6]. The connection between the gradient terminology and chain recurrence is the Fundamental Theorem: chain recurrence is the type of recurrence that, when factored out of the dynamics, leaves the remainder of the space flowing downhill with respect to a Lyapunov function.

6. Back to the theorem

Now we have the tools to understand the theorem.

Theorem 2 (reprise). *The iteration of a continuous function on a compact metric space decomposes the space into a chain recurrent part and a gradient-like part.*

We present here just a flavor of the proof by way of a couple of details: the structure of the chain recurrent set in terms of attractors, and a glance at the complete Lyapunov function which provides the gradient-like structure for the theorem.

A study of the qualitative behavior of a dynamical system inevitably involves the discussion of sets called attractors. An *attractor* is a set $A \subset X$ which attracts neighboring points, in some well-defined way, under the action of a flow or an iterated function f. Definitions abound; see, for example, [10], [12], [16], and references listed therein.

We will utilize the definition given by Conley ([5], [6]) for flows on compact metric spaces, and modified and used by McGehee [11] for the setting of continuous functions on compact metric spaces.

Definition. A set A is an *attractor* for an iterated function f if

(1) A is a nonempty compact invariant set and
(2) there exists a neighborhood U of A such that $\omega(U) = A$.

The *domain of attraction* of an attractor A is the set of all points attracted to A:

$$D(A) \equiv \{x \in X : \omega(x) \subset A\}.$$

The domain of attraction is sometimes called the *basin of attraction*. The set of all points in X outside the influence of the attractor A is called the *repeller complementary to* A and is represented by the following notation:

$$A^* \equiv X - D(A).$$

That is, $A^* \equiv \{x \in X : \omega(x) \not\subset A\}$. (It is true that for any attractor A and any $x \in X$, either $\omega(x) \subset A$ or $\omega(x) \cap A = \varnothing$.) In the case of flows, the dual repeller A^* is an attractor for the reverse flow, and much is made of the symmetry in the dynamical structure of attractor-repeller pairs. While the symmetry does not carry over to the iterating function setting, attractor-repeller pairs nonetheless provide a useful breakdown of X into three sets: points in the attractor A (and which stay there since an attractor is an invariant set), points not in A but which end up there in the sense of ω-limit set, and point not in A which do not approach A under iteration. For more on attractors in this context, see [6], [11], [14], [15].

We skip right to the punch line.

Theorem 3. *For a continuous iterated function on a compact metric space X, there are only countably many attractors A and complementary repellers A^*, and the chain recurrent set is the countable intersection of the pairwise unions of those attractor-repeller pairs:*

$$R(X) = \bigcap \{A \cup A^* : A \ \textit{ is an attractor on } X\}.$$

For proofs in various settings, see [6], [9], [10], [15]. Although the collection of attractors is not the same as the breakdown of the chain recurrent set into chain transitive components, the attractors and their complementary repellers provide both the structure to the chain transitive components and the pieces from which to construct a complete Lyapunov function for the space.

Definition. A *complete Lyapunov function* for the space X with respect to a continuous function f is a continuous, real-valued function g on X satisfying:

(a) g is strictly decreasing on orbits outside the chain recurrent set;
(b) $g(R(X))$ is a compact nowhere dense subset of $\mathbf{R}$;
(c) if $x, y \in R(X)$, then $g(x) = g(y)$ if and only if $x \sim y$; that is, for any $c \in (R(X))$, $g^{-1}(c)$ is a chain transitive component of $R(X)$.

On both sides of the flow/iterating function fence, the construction of the complete Lyapunov function for X is accomplished piecewise from the attractor-repeller pairs. Here is just a quick taste, in our setting of the iteration of a continuous function on a compact metric space, taken directly from [14]. Other constructions are similar.

Lemma. *For each attractor A_n, there is a continuous function g_n such that $g_n^{-1}(0) = A_n$, $g_n^{-1}(1) = A_n^*$, and g_n is strictly decreasing on orbits of points in $X - (A_n \cup A_n^*)$.*

Each function g_n is obtained from sums of suprema of the function

$$g_0(x) = \frac{d(x, A_n)}{d(x, A_n) + d(x, A_n^*)}$$

applied to iterates of the function f, using the attractor-repeller pair (A_n, A_n^*).

Theorem 4. *If f is a continuous function on a compact metric space X, then there is a complete Lyapunov function $g : X \to \mathbf{R}$ for f.*

The Lyapunov function in this case can take the form

$$\sum_{n=1}^{\infty} \frac{2g_n(x)}{3^n},$$

where the g's are from the Lemma, on the attractor-repeller pairs, one pair at a time.

The partial order ">" on the points of the space X generates a partial order on the chain classes of X which is then reflected by the complete Lyapunov function on X. In general, there are many orderings of the components of $R(X)$ by different complete Lyapunov functions, all of which respect the order imposed by the dynamics.

There are modifications and extensions of the Fundamental Theorem in many settings: for example, semiflows, homeomorphisms, diffeomorphisms, relations, on spaces compact, locally compact, noncompact, infinite dimensional, etc. A few references are: [7], [10], [11], [13]–[15].

7. The Fundamental Theorem

Now that we have both a statement and an understanding of the theorem, the question remains: why should this be considered the Fundamental Theorem of Dynamical Systems?

In some sense, all the questions of Dynamical Systems are either variations on or extensions of the following: what is the ultimate fate or dynamical behavior of each point in the system, and how do those behaviors fit together to give a description of the ultimate fate or behavior of the system as a whole? The Conley Decomposition Theorem is a concise description of all possible types of motion from a very basic point of view, a simple statement of "what happens": points are recurrent, or they run downhill on solutions. It is a one-sentence statement, in a nutshell, of all that can happen: parts of the space are either chain recurrent or gradient-like. It is more than a vacuous statement of the Law of the Excluded Middle: "points either exhibit this type of recurrence or they do not". Instead, it gives real alternatives of some mathematically quantifiable behavior. The theorem is fundamental in the sense that it deals with the basic question of the field. It is also fundamental in that it encompasses such big ideas in such a small, concise statement.

Consider briefly some other uses of the phrase "Fundamental Theorem". The Fundamental Theorem of Arithmetic states that every counting number can be expressed uniquely as a product of primes. That is, the elemental building block of Arithmetic, the counting number, can be decomposed uniquely into its basic parts: its prime factors. The Fundamental Theorem of Algebra states that every polynomial with real or complex coefficients can be factored into linear factors in the field of complex numbers. That is, the elemental building block of Algebra, the polynomial, can be decomposed uniquely into its basic parts: its linear factors. One description of the Conley Decomposition Theorem is that the elemental piece of a Dynamical System, the space on which the dynamics take place, can be decomposed uniquely into its basic dynamical parts: points whose dynamics can be described as exhibiting a particular type of recurrence, and points which proceed in a gradient-like fashion and provide the dynamical structure for fitting the recurrent parts together.

The Fundamental Theorem of Calculus is not a result about breaking down into parts but about bringing together the perhaps unexpectedly related derivative and integral. It connects the branches of the field in a fundamental way. In a somewhat similar way, the Conley Decomposition Theorem is about not only the behavior of a system but the study of the discipline:

> It has always seemed to me that this is the correct framework, at the coarsest level, for studying dynamical systems. One can present purely gradient-like systems such as Morse-Smale flows, then study chain recurrent phenomena like Anosov diffeomorphisms or the Smale horseshoe, and finally investigate the way these two ingredients fit together to form more general systems. [8]

This brief exposition does not claim more than it states clearly. It makes no claims about a Fundamental Personage in the History of Dynamical Systems or a Fundamental Paper in Dynamical Systems; the name of Poincaré and a certain paper by Smale, among others, come to mind. This paper is instead about a theorem, one which deals with Poincaré's basic question about the qualitative behavior of a dynamical system, which deals with Smale's classification of Dynamical Systems by qualitative dynamical behavior, and which opens the door to some more results about structural stability and other topics, all of which are at the heart of Dynamical Systems. This paper presents to you the Fundamental Theorem of Dynamical Systems.

Acknowledgements. The author gratefully acknowledges assistance at various stages from R. Easton and M. Hurley, and expresses particular appreciation for many general insights and specific suggestions from G. Hall and R. McGehee.

Dedication. This paper is dedicated to the memory of Charles Cameron Conley. There are many of us who owe Charlie a lot, for our view of Dynamical Systems, for much of what we know about the field, and for the insights he shared, not only with some of us individually, but with the entire mathematical community. I can hardly believe it has been a decade since he left us. When he did leave us, he left a lot of himself behind. Thanks, Charlie.

References

[1] Anosov D.V., *Geodesic Flows on Closed Riemannian Manifolds of Negative Curvature*, Proceedings of the Steklov Institute of Mathematics, Vol. 90, American Mathematical Society, Providence, R.I., 1969.

[2] Block L., Franke J.E., *The chain recurrent set, attractors, and explosions*, Ergodic Theory and Dynamical Systems **5** (1985), 321–327.

[3] Bowen R., *Equilibrium States and the Ergodic Theory of Axiom A Diffeomorphisms*, Lecture Notes in Mathematics, Vol. 470, Springer Verlag, New York, 1975.

[4] ———, *On Axiom A Diffeomorphisms*, CBMS Regional Conference Series in Mathematics, Vol. 35, American Mathematical Society, Providence, R.I., 1978.

[5] Conley C., *The Gradient Structure of a Flow, I*, IBM RC 3932, #17806, 1972; reprinted in Ergodic Theory and Dynamical Systems **8*** (1988), 11–26.

[6] ———, *Isolated Invariant Sets and the Morse Index*, CBMS Regional Conference Series in Mathematics, Vol. 38, American Mathematical Society, Providence, R.I., 1978.

[7] Easton R., *Isolating blocks and epsilon chains for maps*, Physica D **39** (1989), 95–110.

[8] Franks J., *Book review*, Ergodic Theory and Dynamical Systems **7** (1987), 313–315.

[9] ———, *A Variation on the Poincaré-Birkhoff Theorem*, in: Hamiltonian Dynamical Systems, K.R. Meyer and D.G. Saari, eds., American Mathematical Society, Providence, R.I., 1988, pp. 111–117.

[10] Hurley M., *Chain recurrence and attraction in non-compact spaces*, Ergodic Theory and Dynamical Systems **11** (1991), 709–729.

[11] McGehee R.P., *Some Metric Properties of Attractors with Applications to Computer Simulations of Dynamical Systems*, preprint, 1988.

[12] Milnor J., *On the concept of attractor*, Communications in Mathematical Physics **99** (1985), 177–195.

[13] Norton D.E., *Coarse-Grain Dynamics and the Conley Decomposition Theorem*, submitted, 1994.

[14] Norton D.E., *The Conley Decomposition Theorem for Maps: A Metric Approach*, submitted, 1994.
[15] ———, *A Metric Approach to the Conley Decomposition Theorem*, Thesis, University of Minnesota, 1989.
[16] Ruelle D., *Small random perturbations of dynamical systems and the definition of attractors*, Communications in Mathematical Physics **82** (1981), 137–151.

Department of Mathematical Sciences, Villanova University, 800 Lancaster Avenue, Villanova, Pennsylvania 19085, USA

(*Received* October 7, 1994)