

Classification of Linear Flows

by

Anthony Wynne

A thesis submitted in partial fulfilment of the requirements for the degree of

Master of Science

in

Mathematics

Department of Mathematical and Statistical Sciences

University of Alberta

©Anthony Wynne, 2016

Abstract

Flows on normed spaces can be classified using flow equivalences — maps on the space with the property that the structure of one flow is converted into the structure of another flow. Of particular interest are classifications that arise from flow equivalences that are either homeomorphisms or diffeomorphisms. It is possible to completely characterize such classifications based solely on a few simple properties of flows, at least in the case of linear flows on finite-dimensional normed spaces. Results concerning diffeomorphic classification are well known and can be found in many textbooks that discuss continuous dynamical systems. The situation is similar when it comes to homeomorphic classification of hyperbolic flows, but for arbitrary (possibly nonhyperbolic) flows results concerning homeomorphic classification are fairly obscure. This thesis aims to provide a complete discussion of the homeomorphic and diffeomorphic classification of linear flows on finite-dimensional normed spaces.

First and foremost I would like to acknowledge Dr. Arno Berger, my supervisor, for his invaluable support over the last few years. It was on his suggestion that I began investigating this topic, and he has pushed me to make this thesis as strong as possible. It is absolutely the case that without him this thesis would not exist.

I would also like to acknowledge my supervisory committee for taking the time to consider my thesis.

I would like to acknowledge all my teachers throughout my education and especially all those who taught me math. In particular, I would like to acknowledge Dr. James Muldowney — studying honours calculus under him in my first year of university was the principal motivator for my switch into the math program.

Finally, I would like to thank my parents. Without their support it would not have been possible for me to continue my studies as a graduate student.

Contents

- 1 Introduction 1
- Notation 2
- 2 Basic Properties of Linear Flows 5
 - 2.1 Linear Flows and Flow Equivalence 5
 - 2.2 Linear Flows and Bounded Linear Operators 9
 - 2.3 Flow Invariance 14
 - 2.4 Reducible Flows 18
 - 2.5 Irreducible Flows 24
- 3 Classification Theorems 33
 - 3.1 Diffeomorphic Equivalence 33
 - 3.2 Homeomorphic Equivalence of Hyperbolic Flows 35
 - 3.3 Homeomorphic Equivalence of General Linear Flows 43
- 4 Nonhyperbolic Flows 56
 - 4.1 Basic Constructions 56
 - 4.2 Strong and Weak Centres 63
 - 4.3 Block Counting 69
- 5 Examples and Other Considerations 74
 - 5.1 Examples 74
 - 5.2 Other Considerations 77
- A Operator Exponential 83
- B Operator Integral 88
- C Complexification and Realification 92

1 Introduction

The word flow in common parlance elicits associations with streams, rivers, currents in a lake, and so forth. One might imagine a collection of water molecules, each travelling along its own path over time. The individual molecules may ultimately end up being completely scattered, but one might expect similar short term behaviour. This conceptualization serves as a reasonable illustration of mathematical flows as well.

In mathematics, a flow is a map on a space with certain special properties. These properties have, at least from a geometric perspective, the effect of partitioning the space into a collection of disjoint paths. Flows have a real ‘time’ variable, and one can use this variable to travel forward and backward along the various paths. Although the basic flow properties do not limit the long term relative behaviour of two nearby paths, such paths are forced to behave similarly in the short term.

The concept of flows arose from the study of differential equations [6]. Initially the topic of differential equations was approached from a fairly quantitative direction — the primary concern was finding explicit solutions to various differential equations (often coming from the study of physics) given some initial condition. This approach persisted until the time of Poincaré (late nineteenth century) at which point there was a shift toward a more qualitative study, and questions concerning the general behaviour of solutions to differential equations gained prominence. Flows then arise as collections of solutions to differential equations, where each path represents a specific solution for some initial condition. As flows capture the behaviour of all solutions simultaneously, they are valuable considerations when it comes to assessing the qualitative behaviour of solutions to differential equations.

When considering qualitative questions, it is often useful to introduce a notion of equivalence. For example, two simple flows on the Euclidean plane might both consist of spirals toward the origin, and from a qualitative perspective these flows may be viewed as essentially the same. As such, it is worthwhile to consider some notion of flow equivalence. At the most basic level, a flow equivalence between two flows is an automorphism of the space that converts the behaviour of the first flow into the behaviour of the second, so paths of the first flow are mapped into paths of the second, for instance. One can then add additional structure to the notion of flow equivalence, which in turn preserves additional structure

between flows. For example, one may consider a flow equivalence that is also a homeomorphism. Such an equivalence not only preserves basic flow properties but also very loosely preserves the general shape of paths.

This thesis is focused on fully characterizing diffeomorphic and homeomorphic flow equivalence for linear flows, a class of particularly well-behaved flows. The thesis consists of three main parts. In Chapter 2 the notions of linear flows and flow equivalence are formalized and a number of basic properties of linear flows are introduced. Then classification theorems for diffeomorphic and homeomorphic flow equivalence are discussed and proved in Chapter 3, with the caveat of a crucial assumption made regarding the existence of invariant subspaces with certain properties in the proof of the general homeomorphic classification theorem. Chapter 4, the final main chapter, is concerned with justifying this assumption.

Much of chapters 2 and 3 of this thesis is common knowledge in the theory of (continuous-time) linear dynamical systems. Details regarding the linear and diffeomorphic classifications can be found in many standard textbooks (for example, [1], [2], and [11]) and the same is true with regard to the homeomorphic classification of hyperbolic systems. On the other hand, the general (i.e., non-hyperbolic) classification is less well-known. The two pertinent articles, [8] and [9], are not as detailed as they could be and may contain minor inaccuracies. The goal of this thesis is to present a complete, self-contained and detailed account of the classification problem for linear flows. While standard terminology and tools are used throughout, and while the ideas of [9] in particular are followed wherever possible, some aspects of the proofs of the main results (Theorem 3.15 and its corollaries) have not appeared in the literature before. The three appendices provide a short introduction to the tools from linear analysis and algebra required for this thesis.

Notation

As usual \mathbb{N} , \mathbb{Z} , \mathbb{Q} , \mathbb{R} , and \mathbb{C} denote the set of all natural, integer, rational, real, and complex numbers. The natural numbers are the positive integers for the purposes of this thesis — when necessary \mathbb{N}_0 is used to denote the nonnegative integers. There will be several instances when only the positive real numbers are considered, so for simplicity they are represented by \mathbb{R}^+ . Similarly, \mathbb{R}^- represents the negative real numbers.

This thesis is only concerned with flows on finite-dimensional normed spaces, and such spaces are denoted primarily by X but also by Y and Z as necessary. The scalar field of these vector spaces is always either \mathbb{R} or \mathbb{C} . Generally both scalar fields will work. It will be made clear in the text when a specific scalar field is being considered; otherwise, \mathbb{K} will be used to denote the underlying scalar field. The choice of norm for a given space will generally be irrelevant (as long as one is consistent in their choice throughout), so $\|\cdot\|$ will be used to denote any arbitrary norm. It will be made clear in the text when a specific choice of norm is required, and such norms will be differentiated from $\|\cdot\|$ with a subscript, as in $\|\cdot\|_D$ for example. The only exception to this is the standard Euclidean norm on \mathbb{K}^d , which is denoted by $|\cdot|$ as usual. The unit sphere — that is, the set $\{x \in X : \|x\| = 1\}$ — appears occasionally and is denoted by \mathbb{S} . If a specific choice of norm is used, then the unit sphere with respect to that norm will be denoted using the same subscript, as in \mathbb{S}_D .

The dimension of a given normed space is denoted by $d \in \mathbb{N}_0$. Flows on a given vector space are denoted primarily by φ but also by ψ and γ as necessary. Similarly, bounded linear operators on a given normed space are denoted primarily by L but also by M and N as necessary. In particular I and O (as compared to 0 denoting either the number zero or the origin of a given normed space) are always used to represent the identity and zero operators respectively, as the vector space on which they are acting will always be clear from the context. It turns out that the set of all linear flows on a given normed space X can be identified with $\mathcal{L}(X)$, the set of all bounded linear operators on X . A subscript will sometimes be used to denote the bounded linear operator associated with a given flow (as in L_φ for example) but this subscript will be suppressed when it is clear from the context. It will occasionally be necessary to consider a fixed basis of X , and in those situations a given operator L and its matrix form with respect to that basis will be identified. In particular, I and O will also represent appropriately sized identity and zero matrices as necessary. If a matrix is block diagonal, say

$$\begin{bmatrix} A_1 & O & O \\ O & A_2 & O \\ & & \ddots \\ O & O & A_m \end{bmatrix}$$

where the A_k are any combination of square matrices and scalars, then it will be denoted by $\text{diag}(A_1, A_2, \dots, A_m)$. Diagonal matrices will frequently arise when a basis is chosen so that the matrix form of an operator is in either real or complex Jordan canonical form. Complex Jordan blocks of size d with eigenvalue λ will be denoted by $J_d(\lambda)$. Real Jordan blocks of size d with nonreal eigenvalue pair $\{\lambda, \bar{\lambda}\}$ will be denoted by $\tilde{J}_d(\lambda)$. It is often necessary to consider the real and imaginary parts of a complex eigenvalue λ , denoted $\Re\lambda$ and $\Im\lambda$ respectively. Finally, the notation $\sigma(L)$ will be used to denote the set of all eigenvalues of the bounded linear operator L .

It will often prove useful to consider $\bigoplus_{k=1}^m X_k$, the direct sum of an appropriate collection of subspaces X_k of a given normed space X . When X is isomorphic to $\bigoplus_{k=1}^m X_k$, X will be directly identified with $\bigoplus_{k=1}^m X_k$, and $X = \bigoplus_{k=1}^m X_k$ will be written.

As this thesis can occasionally become notationally intense, with various different variables and indices all occurring simultaneously, it is worthwhile to attempt consistency in using these various variables and indices. Generally r and c denote fixed constants in \mathbb{R} and \mathbb{C} respectively. Elements of a space X or Y are denoted by x and y , while t and s are used to represent variables in \mathbb{R} . A subscript n is used to indicate a specific element of a sequence and also as an index for countably infinite sets. Subscripts t and s are used for indices associated with \mathbb{R} , \mathbb{R}^+ , or \mathbb{R}^- . A subscript k is used to indicate a specific component while m denotes the total number of components as in $\bigoplus_{k=1}^m X_k$. Other subscripts and variables will be used as needed.

As final notes, given a space $X = \bigoplus_{k=1}^m X_k$, the notation (x_1, \dots, x_m) is often used to represent a vector $x \in X$ in place of the standard vertical vector notation. It will sometimes be necessary to consider the binary representation of a number. In such a situation $\langle \cdot \rangle$ will be used, so for example $\langle 1010 \rangle$ is the number ten represented in binary. It will also sometimes be necessary to consider the floor and ceiling of some $r \in \mathbb{R}$, denoted $\lfloor r \rfloor$ and $\lceil r \rceil$ respectively, where the floor of r is the largest integer not larger than r while the ceiling of r is the smallest integer not smaller than r .

2 Basic Properties of Linear Flows

The first step toward classifying flows is to make precise the notion of linear flows. Moreover, it is also necessary to develop a notion (several notions in this case) of flow equivalence. One can then build upon these definitions to develop the basic properties of flows. In this chapter most of the properties of linear flows required for the next chapter's discussion of flow classification are developed.

2.1 Linear Flows and Flow Equivalence

Definition 2.1 A *flow* on X is a continuous map $\varphi : \mathbb{R} \times X \rightarrow X$ that satisfies the following:

- (i) $\varphi(0, x) = x$ for all $x \in X$ and
- (ii) $\varphi(s + t, x) = \varphi(s, \varphi(t, x))$ for all $s, t \in \mathbb{R}$ and $x \in X$.

A flow can be viewed geometrically as a collection of disjoint paths, where the set $\varphi(\mathbb{R}, x)$ is the path through x . Points $\varphi(t, x)$ for positive t can then be viewed as points encountered along the path travelling from x , while $\varphi(t, x)$ for negative t can be viewed as points encountered along the path travelling to x . It is a simple consequence of the preceding definition that any two non-disjoint paths are identical. If two paths $\varphi(\mathbb{R}, x)$ and $\varphi(\mathbb{R}, y)$ are not disjoint, then $\varphi(s_1, x) = \varphi(s_2, y)$ for some $s_1, s_2 \in \mathbb{R}$. It follows that $\varphi(t, x) = \varphi(t - s_1 + s_2, y) \in \varphi(\mathbb{R}, y)$ and $\varphi(t, y) = \varphi(t - s_2 + s_1, x) \in \varphi(\mathbb{R}, x)$ for all $t \in \mathbb{R}$.

A trivial example of a flow is the map satisfying $(t, x) \mapsto x$ for all $(t, x) \in \mathbb{R} \times X$. It is not difficult to construct other simple but less trivial flows as well.

Example 2.2 Consider the map φ on $\mathbb{R} \times \mathbb{R}^2$ given by $\varphi(t, (x_1, x_2)) = (e^t(tx_2 + x_1), e^tx_2)$. This flow can be written in the form $\varphi(t, (x_1, x_2)) = e^tM(t)(x_1, x_2)$ where $M(t)$ is the 2×2 matrix

$$\begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix}$$

for all $t \in \mathbb{R}$. Continuity of φ is now clear as $t \mapsto e^tM(t)$ is clearly continuous. Moreover, it is also clear that $e^tM(t) = I$ when $t = 0$ so φ satisfies (i). Finally, it is easily verified that

$e^s M(s)e^t M(t) = e^{s+t} M(s+t)$, so $\varphi(s, \varphi(t, (x_1, x_2))) = \varphi(s+t, (x_1, x_2))$ for all $s, t \in \mathbb{R}$ and $(x_1, x_2) \in \mathbb{R}^2$, and thus φ satisfies (ii) as well. \diamond

Example 2.3 It is also easily verified that the map φ on $\mathbb{R} \times \mathbb{R}^2$ satisfying the equation $\varphi(t, (x_1, x_2)) = (\sin(t)x_2 + \cos(t)x_1, \cos(t)x_2 - \sin(t)x_1)$ is a flow. This is due to the fact that φ is really just $R(t)(x_1, x_2)$ where $R(t)$ is the standard rotation matrix in \mathbb{R}^2 with t in place of θ . The flow properties of φ are immediate consequences of the properties of rotations of the Euclidean plane. \diamond

Flows can also be viewed as families of automorphisms. Linear flows are then defined simply by requiring that these automorphisms are all linear.

Definition 2.4 Let φ be a flow on X , and fix $t \in \mathbb{R}$. The map $\varphi_t : X \rightarrow X$ given by $\varphi_t(x) := \varphi(t, x)$ is the *time- t map* of φ .

Definition 2.5 A flow on X is *linear* if its time- t maps are linear for all $t \in \mathbb{R}$.

Examples 2.2 and 2.3 are clearly both linear, as for each fixed $t \in \mathbb{R}$ the maps φ_t are both of the form Mx for some 2×2 matrix M . Linear flows have many useful properties that do not hold for flows in general, several of which will be demonstrated throughout this chapter. Of particular note is that the set of all linear flows can be identified with the set of bounded linear operators in a canonical fashion. This will be proved in the next section. For now though, it is a straightforward application of the basic flow properties to show that the time- t maps of a flow are all homeomorphisms.

Proposition 2.6 *Let φ be a flow on X . Then φ_t is a homeomorphism for all $t \in \mathbb{R}$.*

Proof. φ_t and φ_{-t} inherit the continuity of φ for each $t \in \mathbb{R}$. Flow property (ii) requires that $\varphi_t \circ \varphi_{-t} = \varphi_0 = \varphi_{-t} \circ \varphi_t$ for each $t \in \mathbb{R}$, and since flow property (i) requires that $\varphi_0 = I$, the invertibility of φ_t for each $t \in \mathbb{R}$ follows. \square

Since X is assumed to be finite-dimensional, it follows that every time- t map of a linear flow is a bounded linear operator.

Another essential notion required to discuss the classification of linear flows is that of flow equivalence. Two flows are considered equivalent if there exists an invertible morphism that

maps the paths of one flow into the paths of the other flow and conversely, in some structured fashion. The exact properties of the morphism may lead to different classifications.

Definition 2.7 Let φ and ψ be flows on X and Y respectively. Then φ and ψ are *flow equivalent* if there exists a bijection (a *flow equivalence*) $h : X \rightarrow Y$ such that

$$h(\varphi(t, x)) = \psi(t, h(x))$$

for all $t \in \mathbb{R}$ and $x \in X$. In particular, φ and ψ are *homeomorphically equivalent* (and h is a *homeomorphic flow equivalence*) if h is a homeomorphism. Similarly, φ and ψ are *diffeomorphically equivalent* (and h is a *diffeomorphic flow equivalence*) if h is a diffeomorphism. Finally, φ and ψ are *linearly equivalent* (and h is a *linear flow equivalence*) if h and h^{-1} are linear.

The ‘flow’ of flow equivalence will often be dropped, as flow equivalence is the only notion of equivalence that appears in this thesis. The requirement that $h(\varphi(t, x)) = \psi(t, h(x))$ is related to the idea that h maps paths into paths since $h(\varphi(\mathbb{R}, x)) = \psi(\mathbb{R}, h(x))$. As t is fixed here, this notion of equivalence does not admit much variation when it comes to how the paths of one flow are mapped into the paths of the other flow. One could weaken this definition to allow for more flexibility; for example, a common alternate definition of equivalence is as above except that $h(\varphi(t, x)) = \psi(rt, h(x))$ for some fixed $r \in \mathbb{R}^+$. While it is still true that $h(\varphi(\mathbb{R}, x)) = \psi(\mathbb{R}, h(x))$ under this weakened definition, the additional factor r allows for a broader notion of equivalence.

Although Definition 2.7 allows for equivalences between flows on different spaces, this thesis is primarily interested in equivalences between flows on the same space. With that said, there are instances where equivalences between flows on different spaces lead to equivalences between flows on the same space. These instances occur as equivalence is transitive. More generally, flow equivalence is an equivalence relation, exactly as one would hope.

Theorem 2.8 *Flow equivalence is an equivalence relation, as are homeomorphic, diffeomorphic, and linear flow equivalence.*

Proof. For reflexivity, note that the identity map is trivially a flow equivalence between any flow and itself. For symmetry, suppose that φ and ψ are two flow equivalent flows on X

and Y respectively, so that there exists a flow equivalence $h : X \rightarrow Y$ between φ and ψ . It follows that h^{-1} is a flow equivalence between ψ and φ since

$$h^{-1}(\psi(t, y)) = h^{-1}(\psi(t, h(h^{-1}(y)))) = h^{-1}(h(\varphi(t, h^{-1}(y)))) = \varphi(t, h^{-1}(y))$$

for all $t \in \mathbb{R}$ and $y \in Y$, so ψ and φ are flow equivalent. Finally, for transitivity, suppose φ , ψ , and γ are three flows on X , Y , and Z respectively such that φ is flow equivalent to ψ , and such that ψ is flow equivalent to γ . Then there exist two flow equivalences $g : X \rightarrow Y$ and $h : Y \rightarrow Z$ between φ and ψ and between ψ and γ respectively. It follows that $h \circ g$ is a flow equivalence between φ and γ since

$$(h \circ g)(\varphi(t, x)) = h(\psi(t, g(x))) = \gamma(t, (h \circ g)(x))$$

for all $t \in \mathbb{R}$ and $x \in X$, so φ and γ are flow equivalent. The proofs for the homeomorphic, diffeomorphic, and linear cases are similar. \square

To simplify the process of examining the various types of equivalence, it is worthwhile to seek out relationships between these notions. For example, two diffeomorphically equivalent flows are necessarily also homeomorphically equivalent, as differentiability implies continuity. The following proposition summarizes a number of similar relationships between equivalences.

Proposition 2.9 *Let φ and ψ be two flows on X and Y , respectively. If φ and ψ are linearly equivalent, then they are also diffeomorphically equivalent. If φ and ψ are diffeomorphically equivalent, then they are also homeomorphically equivalent.*

Proof. The first statement follows since finite-dimensional linear maps are trivially differentiable, as they are already the best linear approximations of themselves. The second statement is a direct consequence of the fact that differentiability implies continuity. \square

Note that Definition 2.7 does not require the flows to be linear. Note also that h need not be unique — it follows immediately from Proposition 2.6 and flow property (ii) that each φ_t is a flow equivalence between φ and itself. It simplifies matters if an equivalence fixes the origin, and, at least for equivalent linear flows, it is always possible to find such an equivalence.

Proposition 2.10 *Let φ and ψ be linear flows on X and Y , respectively. If φ and ψ are homeomorphically (respectively, diffeomorphically or linearly) equivalent, then there exists a homeomorphic (respectively, diffeomorphic or linear) equivalence h between φ and ψ such that $h(0) = 0$.*

Proof. If in particular φ and ψ are linearly equivalent, then any linear equivalence h will work, as necessarily $h(0) = 0$. Otherwise, let \tilde{h} be any homeomorphic equivalence between φ and ψ , and let $h : X \rightarrow Y$ be given by $h(x) = \tilde{h}(x) - \tilde{h}(0)$ for all $x \in X$.

The map h is clearly invertible with $h^{-1} : Y \rightarrow X$ given by $h^{-1}(y) = \tilde{h}^{-1}(y + \tilde{h}(0))$ for all $y \in Y$. Also since φ and ψ are linear for each $t \in \mathbb{R}$

$$\begin{aligned} h(\varphi(t, x)) &= \tilde{h}(\varphi(t, x)) - \tilde{h}(0) \\ &= \tilde{h}(\varphi(t, x)) - \tilde{h}(\varphi(t, 0)) \\ &= \psi(t, \tilde{h}(x)) - \psi(t, \tilde{h}(0)) \\ &= \psi(t, \tilde{h}(x) - \tilde{h}(0)) \\ &= \psi(t, h(x)) \end{aligned}$$

for all $t \in \mathbb{R}$ and $x \in X$, so h is a flow equivalence; moreover, as a consequence of basic properties of limits h is homeomorphic whenever \tilde{h} is. Finally, $h(0) = \tilde{h}(0) - \tilde{h}(0) = 0$. The diffeomorphic case is similar. \square

As a consequence of the preceding proposition, from this point on it will be assumed that all equivalences fix the origin.

2.2 Linear Flows and Bounded Linear Operators

It is often non-obvious, based solely on Definitions 2.1 and 2.7, whether or not two flows are equivalent. As such, it is desirable to further develop the notion of flows; ideally, flows could then be classified based solely on some easily determined properties. In the case of linear flows, this can be achieved by identifying the set of flows with the set of bounded linear operators. Questions concerning equivalence can then be answered by examining the related operator. Toward this end it is first shown that linear operators induce linear flows via the operator exponential. The operator exponential is defined based on the Taylor

Series representation of the real exponential. All of the familiar properties of the real exponential carry over to the operator exponential with only minor adjustments (adjustments are needed to handle the fact that the algebra of linear operators is not commutative) and the exponential of a bounded linear operator is itself a bounded linear operator. A more detailed discussion of the operator exponential can be found in Appendix A.

Proposition 2.11 *Let $L \in \mathcal{L}(X)$. The map $\varphi : \mathbb{R} \times X \rightarrow X$ given by $\varphi(t, x) := e^{tL}x$ is a linear flow on X .*

Proof. First fix $(t_0, x_0) \in \mathbb{R} \times X$ and $\epsilon \in \mathbb{R}^+$, and consider that by the triangle inequality

$$\|e^{tL}x - e^{t_0L}x_0\| \leq \|e^{tL} - e^{t_0L}\|\|x - x_0\| + \|e^{tL} - e^{t_0L}\|\|x_0\| + \|e^{t_0L}\|\|x - x_0\|$$

for all $t \in \mathbb{R}$ and $x \in X$. The map $t \mapsto e^{tL}$ is continuous (in fact, differentiable) so there exists a $\tilde{\delta} \in \mathbb{R}^+$ such that $\|e^{tL} - e^{t_0L}\| < \min\{\frac{\epsilon}{3\|x_0\|+1}, \sqrt{\frac{\epsilon}{3}}\}$ for all $t \in \mathbb{R}$ satisfying $|t - t_0| < \tilde{\delta}$. Now set $\delta := \min\{\tilde{\delta}, \frac{\epsilon}{3\|e^{t_0L}\|}\}$. It follows, for all $(t, x) \in \mathbb{R} \times X$ satisfying $\|t - t_0\|, \|x - x_0\| < \delta$, that $\|e^{tL}x - e^{t_0L}x_0\| < \sqrt{\frac{\epsilon}{3}}\sqrt{\frac{\epsilon}{3}} + \frac{\epsilon}{3\|x_0\|+1}\|x_0\| + \|e^{t_0L}\|\frac{\epsilon}{3\|e^{t_0L}\|} < \epsilon$. Since (t_0, x_0) and ϵ were arbitrary, φ is continuous.

Flow properties (i) and (ii) are direct consequences of the basic properties of the exponential, and the fact that φ is linear follows from the fact that the exponential of a bounded linear operator is itself a bounded linear operator. \square

When a flow φ is of the form $e^{tL}x$ for some bounded linear operator L , then L is said to generate the flow φ . It is easily seen that two different operators cannot generate the same flow.

Proposition 2.12 *Let φ be a linear flow on X generated by $L, M \in \mathcal{L}(X)$. Then $L = M$.*

Proof. By assumption $e^{tL}x = e^{tM}x$ for all $(t, x) \in \mathbb{R} \times X$. It follows that $Le^{tL}x = Me^{tM}x$ for all $(t, x) \in \mathbb{R} \times X$, by differentiating both sides with respect to t . In particular $Lx = Mx$ for all $x \in X$ by fixing $t = 0$, and thus $L = M$. \square

Proposition 2.11 states that every linear operator generates a linear flow in a canonical fashion. This allows one to easily construct a wide variety of flows; the flows in Examples

2.2 and 2.3 were constructed by applying Proposition 2.11 to

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \text{ and } \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

respectively. It is natural to ask if the converse is true — that is, are all linear flows generated by linear operators? This is indeed the case. An elegant proof of this fact, as outlined by my supervisor, makes use of the operator integral. The operator integral is defined by naively applying Riemann integration to continuous maps $f : \mathbb{R} \rightarrow \mathcal{L}(X)$. Given a flow φ and numbers $a, b \in \mathbb{R}$, one may then construct $\int_a^b \varphi_t dt$. As was the case with the operator exponential, the properties of the operator integral are similar to those of the Riemann integral, and $\int_a^b \varphi_t dt$ is a bounded linear operator. Operator integration is discussed further in Appendix B.

It is clear from the preceding discussion that a flow generated by some bounded linear operator is necessarily differentiable, and the derivative at $t = 0$ is the operator that generates it. The idea then is to first show that every linear flow is differentiable at $t = 0$. From there one shows that the flow generated by this derivative is in fact the original flow.

Lemma 2.13 *Let φ be a linear flow on X . Then the map $t \mapsto \varphi_t x$ is differentiable at $t = 0$ for all $x \in X$.*

Proof. Set $Y := \{x \in X : \lim_{t \rightarrow 0} \frac{1}{t}(\varphi_t x - x) \text{ exists}\}$. Note that Y cannot be empty, as $\lim_{t \rightarrow 0} \frac{1}{t}(\varphi_t 0 - 0) = 0$. It follows from the fact that φ_t is linear for each $t \in \mathbb{R}$ and the basic properties of limits that Y is in fact a subspace of X . By construction then Y is the subspace of X for which the map $t \mapsto \varphi_t$ is differentiable at $t = 0$, so the goal is to show that $Y = X$.

Define $A_n \in \mathcal{L}(X)$ by setting $A_n := n \int_0^{\frac{1}{n}} \varphi_t dt$ for each $n \in \mathbb{N}$. Then $\lim_{n \rightarrow \infty} A_n = I$, so that for all $x \in X$ the sequence $\{A_n x\}_{n \in \mathbb{N}}$ converges to x . Since Y is a subspace of a finite-dimensional space, Y must be closed. As a consequence of the above two statements, if it can be shown that $A_n x \in Y$ for all $x \in X$ and $n \in \mathbb{N}$, then $Y = X$. But this follows

from the various basic properties of the operator integral, as

$$\begin{aligned}
\lim_{s \rightarrow 0} \frac{1}{s} (\varphi_s A_n x - A_n x) &= \lim_{s \rightarrow 0} \frac{n}{s} \left(\varphi_s \int_0^{\frac{1}{n}} \varphi_t dt x - \int_0^{\frac{1}{n}} \varphi_t dt x \right) \\
&= \lim_{s \rightarrow 0} \frac{n}{s} \left(\int_0^{\frac{1}{n}} \varphi_{s+t} dt x - \int_0^{\frac{1}{n}} \varphi_t dt x \right) \\
&= \lim_{s \rightarrow 0} \frac{n}{s} \left(\int_s^{s+\frac{1}{n}} \varphi_t dt - \int_0^{\frac{1}{n}} \varphi_t dt \right) x \\
&= \lim_{s \rightarrow 0} \frac{n}{s} \left(\int_{\frac{1}{n}}^{s+\frac{1}{n}} \varphi_t dt - \int_0^s \varphi_t dt \right) x \\
&= n \left(\lim_{s \rightarrow 0} \frac{1}{s} \int_{\frac{1}{n}}^{s+\frac{1}{n}} \varphi_t dt - \lim_{s \rightarrow 0} \frac{1}{s} \int_0^s \varphi_t dt \right) x \\
&= n(\varphi_{\frac{1}{n}} - I)x
\end{aligned}$$

for all $x \in X$ and $n \in \mathbb{N}$. \square

Recalling that the idea of flows arose from the study of differential equations, and now armed with the knowledge that every linear flow is differentiable, it is possible to demonstrate that every linear flow is the exponential of a linear operator; that is, every linear flow is of the form described in Proposition 2.11.

Proposition 2.14 *Let φ be a linear flow on X . There exists an $L_\varphi \in \mathcal{L}(X)$ such that $\varphi(t, x) = e^{tL_\varphi} x$ for all $(t, x) \in \mathbb{R} \times X$.*

Proof. Set $L_\varphi := \lim_{t \rightarrow 0} \frac{1}{t} (\varphi_t - I)$. Now L_φ is linear as a consequence of the basic properties of limits, and since φ_t is linear for each $t \in \mathbb{R}$ by assumption. It then follows that $L_\varphi \in \mathcal{L}(X)$, as X is finite-dimensional.

Consider that $\lim_{t \rightarrow s} \frac{\varphi_t - \varphi_s}{t-s} = \lim_{t \rightarrow s} \frac{\varphi_{t-s} - I}{t-s} \varphi_s = \lim_{t-s \rightarrow 0} \frac{\varphi_{t-s} - I}{t-s} \varphi_s = L_\varphi \varphi_s$ for all $s, t \in \mathbb{R}$ by the previous lemma. Thus φ is differentiable with respect to t for all $t \in \mathbb{R}$ with derivative $\frac{d}{dt} \varphi_t = L_\varphi \varphi_t$. Now consider the map $t \mapsto e^{-tL_\varphi} \varphi_t$. This is differentiable with derivative $-L_\varphi e^{-tL_\varphi} \varphi_t + e^{-tL_\varphi} L_\varphi \varphi_t = 0$, so $e^{-tL_\varphi} \varphi_t = C$ for some constant $C \in \mathcal{L}(X)$, and consequently $\varphi_t = C e^{tL_\varphi}$. But then $\varphi_0 = I$, so $C = I$, and thus $\varphi_t = e^{tL_\varphi}$. The desired result follows. \square

The following theorem summarizes the relationship between linear flows and bounded linear operators.

Theorem 2.15 *Every linear flow on X can be uniquely identified with a bounded linear operator on X and conversely.*

Proof. This follows immediately from Propositions 2.11, 2.12, and 2.14. \square

The identification between linear flows and linear operators is the foundation from which the classification of linear flows can be meaningfully discussed. It will ultimately be shown that the existence of an equivalence between two linear flows will depend entirely on the relative properties of the underlying operators. In particular, operator spectra and the subspaces they induce will prove essential to determining whether or not two linear flows are homeomorphically equivalent. Furthermore, it is now possible to provide an initial classification result for linear equivalence, the strongest notion of equivalence discussed in this thesis. First it is necessary to provide a notion of similar operators.

Definition 2.16 Let $L, M \in \mathcal{L}(X)$. Then L and M are *similar* if there exists an invertible $N \in \mathcal{L}(X)$ such that $NLN^{-1} = M$.

Theorem 2.17 *Let φ and ψ be linear flows on X . Then φ and ψ are linearly equivalent if and only if L_φ and L_ψ are similar.*

Proof. Suppose first that φ and ψ are linearly equivalent. Then there exists an invertible linear $h : X \rightarrow X$ such that $h\varphi(t, x) = \psi(t, hx)$ for all $t \in \mathbb{R} \times X$. Now $h \in \mathcal{L}(X)$ since X is finite-dimensional, and $e^{tL_\varphi h^{-1}} hx = he^{tL_\varphi} x = h\varphi(t, x) = \psi(t, hx) = e^{tL_\psi} hx$ for all $(t, x) \in \mathbb{R} \times X$ as a consequence of Proposition 2.14. It follows from Proposition 2.12 that $hL_\varphi h^{-1} = L_\psi$, so L_φ and L_ψ are similar.

Conversely, suppose L_φ and L_ψ are similar. Then there exists an invertible $N \in \mathcal{L}(X)$ such that $NL_\varphi N^{-1} = L_\psi$. Set $h : X \rightarrow X$ so that $h(x) = Nx$ for all $x \in X$. Then $h(\varphi(t, x)) = Ne^{tL_\varphi} x = e^{tNL_\varphi N^{-1}} Nx = e^{tL_\psi} Nx = \psi(t, h(x))$ for all $(t, x) \in \mathbb{R} \times X$, and h is clearly an invertible linear map. Thus φ and ψ are linearly equivalent. \square

After fixing a basis for X , it is now straightforward to check if two linear flows are linearly equivalent. First find the two operators that generate the flows by differentiating. These two operators can be represented as matrices with respect to the chosen basis, and it is then a simple exercise in verifying whether or not these two matrices are similar. Conversely one can easily generate linearly equivalent linear flows using similar matrices.

The remainder of this chapter is focused on developing additional properties of flows; forms of equivalence aside from linear will be mostly ignored until the third chapter. In particular, it will prove useful to represent flows in a specific form with respect to a given basis. Now an arbitrary flow may not have the desired form with respect to an arbitrary basis. With that said, it will always be possible to find a flow that has the desired form via similar operators; in other words, given an arbitrary flow and basis, one can always find a linearly equivalent flow that is of the desired form with respect to the basis. As a consequence of Proposition 2.9, it is possible to assume that flows are always of the desired form, regardless of the type of equivalence being investigated.

2.3 Flow Invariance

Although the identification between linear flows and linear operators will prove fundamental to this discussion of the classification of linear flows, there are several other concepts that will also prove useful. The first of these is invariance.

Definition 2.18 A subspace Y of X is *invariant* under $L \in \mathcal{L}(X)$ if $L(Y) \subseteq Y$, and Y is (*flow*) *invariant* under a linear flow φ on X if it is invariant under φ_t for all $t \in \mathbb{R}$.

The notion of invariance under an operator is quite natural. In the context of flows, invariance is the idea that any path through some subspace exists entirely within that subspace; one cannot travel along a path either into or out of that subspace. Note that, while equality is not required for invariance, in the case of flow invariance equality is automatic: if $\varphi_t(Y) \subseteq Y$ for all $t \in \mathbb{R}$, then $\varphi_{-t}(Y) \subseteq Y$ for all $t \in \mathbb{R}$, and thus $Y = \varphi_t(\varphi_{-t}(Y)) \subseteq \varphi_t(Y)$ for all $t \in \mathbb{R}$.

Invariance can be used to simplify the analysis of flows. The general idea is to consider the behaviour of a linear flow along its various invariant subspaces. Every flow induces component flows along these subspaces, and by working with specific well-chosen invariant subspaces one can reduce the problem of classifying a flow to classifying certain component flows. As some flows are significantly easier to classify than others, this can greatly simplify the problem.

Definition 2.19 Let $L \in \mathcal{L}(X)$, and let Y be a subspace of X invariant under L . The map $L|_Y : Y \rightarrow Y$ given by $L|_Y(y) := Ly$ for all $y \in Y$ is the Y component of L . Similarly, let φ be a linear flow on X , and let Y be a subspace of X invariant under φ . The map $\varphi|_Y : \mathbb{R} \times Y \rightarrow Y$ with $\varphi|_Y(t, y) := \varphi(t, y)$ for all $(t, y) \in \mathbb{R} \times Y$ is the Y component of φ .

It is clear that a component of a bounded linear operator is itself a bounded linear operator. Similarly, a component of a linear flow is itself a linear flow.

It will often be the case that many invariant subspaces will be considered simultaneously. If $\{X_k\}_{k=1}^m$ is a finite collection of subspaces of X invariant under some operator L or flow φ , then L_k and φ_k will be written in place of $L|_{X_k}$ and $\varphi|_{X_k}$ respectively to simplify notation.

Given the identification between linear flows and linear operators, it is natural to examine the relationship between invariance under linear flows and under operators.

Proposition 2.20 *Let φ be a linear flow on X , and let Y be a subspace of X . Then Y is invariant under L_φ if and only if Y is invariant under φ .*

Proof. Suppose first that Y is invariant under L_φ , and fix y in Y . Now by assumption $L_\varphi y \in Y$; moreover, if $L_\varphi^k y \in Y$ for some $k \in \mathbb{N}$, then $L_\varphi^{k+1} y \in Y$. It follows by induction on k that $L_\varphi^k y \in Y$ for all $k \in \mathbb{N}$. This can be extended to \mathbb{N}_0 as $L_\varphi^0 y = Iy = y \in Y$. It follows that $\sum_{k=0}^n \frac{1}{k!} (tL_\varphi)^k y \in Y$ for all $t \in \mathbb{R}$ and $n \in \mathbb{N}$, as Y is a subspace of X . Furthermore, Y is closed since X is assumed to be finite-dimensional, and therefore $\varphi_t y = \lim_{n \rightarrow \infty} \sum_{k=0}^n \frac{1}{k!} (tL_\varphi)^k y \in Y$ for all $t \in \mathbb{R}$. Consequently Y is invariant under φ , as y was arbitrary.

Suppose now that Y is invariant under φ and consider that $L_\varphi x = [\frac{d}{dt}\varphi_t]_{t=0}x$ for all $x \in Y$. On Y in particular $[\frac{d}{dt}\varphi_t]_{t=0}y = [\frac{d}{dt}(\varphi|_Y)_t]_{t=0}y$, so $L_\varphi y = [\frac{d}{dt}(\varphi|_Y)_t]_{t=0}y$ for all $y \in Y$. It follows from the definition of the derivative that $[\frac{d}{dt}(\varphi|_Y)_t]_{t=0}y \in \mathcal{L}(Y)$, and thus $L_\varphi y \in Y$ for all $y \in Y$, so Y is invariant under L_φ . \square

As a consequence of the preceding proposition, there is no need to distinguish between invariance under a flow and the operator that generates it. As such, subspaces will often be referred to as simply invariant, without reference to a specific operator or flow.

Now for invariance to be a valuable consideration when it comes to classifying flows, it is important that invariance be preserved by equivalence, and this is indeed the case.

Proposition 2.21 *Let h be an equivalence between two linear flows φ and ψ on X and Y respectively. Further let Z be an invariant subspace of X . If $h(Z)$ is a subspace of Y , then $h(Z)$ is invariant under ψ ; moreover, $h|_Z : Z \rightarrow h(Z)$ is an equivalence between $\varphi|_Z$ and $\psi|_{h(Z)}$.*

Proof. Suppose $y \in h(Z)$. Then $y = h(z)$ for some $z \in Z$, and since Z is invariant under φ it follows that $\psi_t y = \psi_t h(z) = h(\varphi_t z) \in h(Z)$. This proves the first statement, as t and y were arbitrary. The second statement is a consequence of the fact that $h|_Z$ inherits the properties of h and $h|_Z(\varphi|_Z(t, z)) = h(\varphi(t, z)) = \psi(t, h(z)) = \psi|_{h(Z)}(t, h|_Z(z))$ for all $(t, z) \in \mathbb{R} \times Z$. \square

If h in the preceding proposition is a linear equivalence, then necessarily $h(Z)$ is a subspace of X , so the preceding proposition applies. Unfortunately the same cannot be said for weaker forms of equivalence such as homeomorphic equivalence. When considering homeomorphic equivalence it is necessary to seek out invariant subspaces for which the image under a homeomorphic equivalence is itself a subspace. Of particular interest, then, will be invariant subspaces defined by a specific property — a property that is preserved by homeomorphic equivalence. Two important examples of this are the stable and unstable subspaces of a flow.

Definition 2.22 Let φ be a linear flow on X . The *stable* subspace of φ , denoted X_φ^+ , is the set of all $x \in X$ that satisfy $\lim_{t \rightarrow +\infty} \varphi_t x = 0$. The *unstable* subspace of φ , denoted X_φ^- , is the set of all $x \in X$ that satisfy $\lim_{t \rightarrow -\infty} \varphi_t x = 0$. The stable and unstable components of φ are denoted by φ^+ and φ^- respectively.

It is a simple consequence of the basic properties of limits, along with the fact that φ_t is linear for each $t \in \mathbb{R}$, that the stable and unstable subspaces of any linear flow are indeed subspaces. These subspaces play a significant role in characterizing homeomorphic equivalence and will be discussed in increasing detail throughout the remaining sections of this chapter. First it is shown that these subspaces are always invariant. The proof is straightforward.

Proposition 2.23 *Let φ be a linear flow on X . Then X_φ^+ and X_φ^- are invariant under φ .*

Proof. Consider that $\lim_{t \rightarrow +\infty} \varphi_t \varphi_s x = \lim_{t \rightarrow +\infty} \varphi_s \varphi_t x = \varphi_s(\lim_{t \rightarrow +\infty} \varphi_t x) = 0$ for all $s \in \mathbb{R}$ and $x \in X_\varphi^+$. Thus, X_φ^+ is invariant under φ . The X_φ^- case is proved in a similar fashion. \square

As previously discussed, it is desirable that the defining property of the stable and unstable subspaces be preserved by equivalence. This is indeed the case for homeomorphic or stronger equivalence. Recall that homeomorphic equivalences are assumed to fix the origin as a consequence of Proposition 2.10.

Proposition 2.24 *Let h be a homeomorphic equivalence between two linear flows, φ and ψ , on X . Then $h(X_\varphi^+) = X_\psi^+$ and $h(X_\varphi^-) = X_\psi^-$.*

Proof. As h is continuous $\lim_{t \rightarrow +\infty} \psi_t h(x) = \lim_{t \rightarrow +\infty} h(\varphi_t x) = h(\lim_{t \rightarrow +\infty} \varphi_t x) = 0$ for all $x \in X_\varphi^+$. Thus $h(X_\varphi^+) \subseteq X_\psi^+$. On the other hand, since h^{-1} is also continuous it follows that $\lim_{t \rightarrow +\infty} \varphi_t h^{-1}(x) = \lim_{t \rightarrow +\infty} h^{-1}(\psi_t x) = h^{-1}(\lim_{t \rightarrow +\infty} \psi_t x) = 0$ for all $x \in X_\psi^+$. Consequently $X_\psi^+ \subseteq h(X_\varphi^+)$. The unstable case is proved similarly. \square

One final useful property of the stable and unstable subspaces of a flow is that these subspaces only ever intersect trivially. This can be seen by considering a flow restricted to the intersection of its stable and unstable subspaces.

Proposition 2.25 *Let φ be a linear flow on X . Then $X_\varphi^+ \cap X_\varphi^- = \{0\}$.*

Proof. First note that $f(Y \cap Z) \subseteq f(Y) \cap f(Z)$ for all subsets Y and Z of X and for all maps $f : X \rightarrow X$. Since the intersection of any two subspaces is itself a subspace, it follows that $Y \cap Z$ is invariant whenever Y and Z are invariant, and thus it is sensible to consider the component $\varphi|_{X_\varphi^+ \cap X_\varphi^-}$ of φ . Since $\varphi|_{X_\varphi^+ \cap X_\varphi^-}$ is a flow on $X_\varphi^+ \cap X_\varphi^-$, there exists some $L \in \mathcal{L}(X_\varphi^+ \cap X_\varphi^-)$ such that L generates $\varphi|_{X_\varphi^+ \cap X_\varphi^-}$. Fix $x \in X_\varphi^+ \cap X_\varphi^-$. Now by assumption $\lim_{t \rightarrow +\infty} e^{tL}x = \lim_{t \rightarrow -\infty} e^{tL}x = 0$, so there exists $T_x^+, T_x^- \in \mathbb{R}$ such that $\|e^{tL}x\| < 1$ for all $t < T_x^-$ and $\|e^{tL}x\| < 1$ for all $t > T_x^+$. Without loss of generality, it may be assumed that $T_x^- < T_x^+$. Since the map $t \mapsto \|e^{tL}x\|$ from $[T_x^-, T_x^+]$ to \mathbb{R} is a continuous function on a closed interval, it follows that $\sup_{t \in \mathbb{R}} \|e^{tL}x\| < \infty$, and this in fact holds for every $x \in X_\varphi^+ \cap X_\varphi^-$ since x was arbitrary. Then by the uniform boundedness principle $\sup_{t \in \mathbb{R}} \|e^{tL}\| < \infty$. But $\|x\| = \lim_{t \rightarrow +\infty} \|x\| = \lim_{t \rightarrow +\infty} \|e^{-tL}e^{tL}x\| \leq \sup_{t \in \mathbb{R}} \|e^{tL}\| \lim_{t \rightarrow +\infty} \|e^{tL}x\| = 0$ for all $x \in X_\varphi^+ \cap X_\varphi^-$, so it must be that $X_\varphi^+ \cap X_\varphi^- = \{0\}$. \square

The preceding property is useful in that one can meaningfully consider $X_\varphi^+ \oplus X_\varphi^-$. This leads nicely into the next section's discussion on reducibility, as direct sums play a substantial role.

2.4 Reducible Flows

The notion of component flows was introduced in the previous section, along with the idea of simplifying the problem of classifying flows to classifying their various components. For this technique to be useful, it is necessary that the classification of a flow is entirely determined by classifying a small number of easily classified components. Unfortunately flows can have numerous components; worse still, components may overlap, and they may end up being overly complex in their own right. To deal with these issues, a notion of 'recombining' component flows is introduced.

The tool for this purpose is the direct sum. Given a space X , if $\{X_k\}_{k=1}^m$ is a finite collection of subspaces of X , and if $X_k \cap (\sum_{j=1}^{k-1} X_j + \sum_{j=k+1}^m X_j) = \{0\}$ for all $k \leq m$, then one may construct $\bigoplus_{k=1}^m X_k$. If one is further given maps $f_k : X_k \rightarrow X_k$ for each $k \leq m$, then one may construct the map $\bigoplus_{k=1}^m f_k : \bigoplus_{k=1}^m X_k \rightarrow \bigoplus_{k=1}^m X_k$ defined by setting $(\bigoplus_{k=1}^m f_k)(x) := \sum_{k=1}^m f_k(x_k)$ for all $x = \sum_{k=1}^m x_k \in \bigoplus_{k=1}^m X_k$. It is easily verified that $\bigoplus_{k=1}^m L_k$ is a bounded linear operator whenever all the L_k are likewise, and similarly $\bigoplus_{k=1}^m \varphi_k$ is a linear flow whenever all the φ_k are as well.

Proposition 2.26 *Fix $m \in \mathbb{N}$, and let $\{X_k\}_{k=1}^m$ be a collection of subspaces of X such that $X_k \cap (\sum_{j=1}^{k-1} X_j + \sum_{j=k+1}^m X_j) = \{0\}$ for all $k \leq m$. Then $\bigoplus_{k=1}^m L_k \in \mathcal{L}(\bigoplus_{k=1}^m X_k)$ whenever $L_k \in \mathcal{L}(X_k)$ for each $k \leq m$. Similarly, $\bigoplus_{k=1}^m \varphi_k$ is a linear flow on $\bigoplus_{k=1}^m X_k$ whenever φ_k is a linear flow on X_k for each $k \leq m$.*

Proof. Note first that

$$\begin{aligned} \left(\bigoplus_{k=1}^m L_k\right)(ax + by) &= \sum_{k=1}^m L_k(ax_k + by_k) \\ &= \sum_{k=1}^m (aL_k x_k + bL_k y_k) \\ &= a \sum_{k=1}^m L_k x_k + b \sum_{k=1}^m L_k y_k \\ &= a \left(\bigoplus_{k=1}^m L_k\right)x + b \left(\bigoplus_{k=1}^m L_k\right)y \end{aligned}$$

for all $a, b \in \mathbb{K}$ and $x, y \in \bigoplus_{k=1}^m X_k$, so $\bigoplus_{k=1}^m L_k$ is indeed linear. Now fix $x \in \bigoplus_{k=1}^m X_k$ along with a sequence $\{x_n\}_{n \in \mathbb{N}}$ in $\bigoplus_{k=1}^m X_k$ converging to x . Then for each $k \leq m$ the sequence $\{x_{n,k}\}_{n \in \mathbb{N}}$ in X_k converges to x_k . It follows by the continuity of each L_k that $\lim_{n \rightarrow \infty} L_k x_{n,k} = L_k x_k$ for all $k \leq m$, so $\lim_{n \rightarrow \infty} \sum_{k=1}^m L_k x_{n,k} = \sum_{k=1}^m L_k x_k$, and thus $\lim_{n \rightarrow \infty} (\bigoplus_{k=1}^m L_k) x_n = (\bigoplus_{k=1}^m L_k) x$. As x and $\{x_n\}_{n \in \mathbb{N}}$ were arbitrary, $\bigoplus_{k=1}^m L_k$ is continuous, and therefore $\bigoplus_{k=1}^m L_k \in \mathcal{L}(\bigoplus_{k=1}^m X_k)$.

It follows from the above that $\bigoplus_{k=1}^m \varphi_k$ is linear for each fixed $t \in \mathbb{R}$, and an argument similar to the preceding continuity argument demonstrates that $\bigoplus_{k=1}^m \varphi_k$ is also continuous. Now $(\bigoplus_{k=1}^m \varphi_k)(0, x) = \sum_{k=1}^m \varphi_k(0, x_k) = \sum_{k=1}^m x_k = x$ for all $x \in \bigoplus_{k=1}^m X_k$ and

$$\begin{aligned} \left(\bigoplus_{k=1}^m \varphi_k \right) (s + t, x) &= \sum_{k=1}^m \varphi_k(s + t, x_k) \\ &= \sum_{k=1}^m \varphi_k(s, \varphi_k(t, x_k)) \\ &= \left(\bigoplus_{k=1}^m \varphi_k \right) \left(s, \sum_{k=1}^m \varphi_k(t, x_k) \right) \\ &= \left(\bigoplus_{k=1}^m \varphi_k \right) \left(s, \left(\bigoplus_{k=1}^m \varphi_k \right) (t, x) \right) \end{aligned}$$

for all $s, t \in \mathbb{R}$ and $x \in \bigoplus_{k=1}^m X_k$. Thus $\bigoplus_{k=1}^m \varphi_k$ is a linear flow on $\bigoplus_{k=1}^m X_k$. \square

Note that the various X_k are invariant under these constructions, so that the various L_k and φ_k are components of $\bigoplus_{k=1}^m L_k$ and $\bigoplus_{k=1}^m \varphi_k$ respectively. It is constructions such as these that allow one to classify a flow based on its components as the following lemma demonstrates.

Lemma 2.27 *Fix $m \in \mathbb{N}$, and let $\{X_k\}_{k=1}^m$ and $\{Y_k\}_{k=1}^m$ be collections of subspaces of X such that $X_k \cap (\sum_{j=1}^{k-1} X_j + \sum_{j=k+1}^m X_j) = \{0\}$ for all $k \leq m$ and similarly for $\{Y_k\}_{k=1}^m$. If φ_k and ψ_k are (homeomorphically or linearly) equivalent linear flows on X_k and Y_k respectively for each $k \leq m$, then $\bigoplus_{k=1}^m \varphi_k$ and $\bigoplus_{k=1}^m \psi_k$ are (respectively homeomorphically or linearly) equivalent flows on $\bigoplus_{k=1}^m X_k$ and $\bigoplus_{k=1}^m Y_k$ respectively.*

Proof. Let $h_k : X_k \rightarrow Y_k$ be equivalences between φ_k and ψ_k for each $k \leq m$. It is easily verified that $\bigoplus_{k=1}^m h_k$ has inverse $\bigoplus_{k=1}^m h_k^{-1}$. Now for all $(t, x) \in \mathbb{R} \times \bigoplus_{k=1}^m X_k$ necessarily $h_k(\varphi_k(t, x_k)) = \psi_k(t, h_k(x_k))$, implying that $\sum_{k=1}^m h_k(\varphi_k(t, x_k)) = \sum_{k=1}^m \psi_k(t, h_k(x_k))$, and thus $(\bigoplus_{k=1}^m h_k)((\bigoplus_{k=1}^m \varphi_k)(t, x)) = (\bigoplus_{k=1}^m \psi_k)(t, (\bigoplus_{k=1}^m h_k)(x))$. Now if all the h_k are continuous (respectively linear), then $\bigoplus_{k=1}^m h_k$ is also continuous (respectively linear) as per

the argument in the previous proposition. Similarly, if the h_k^{-1} are continuous (respectively linear), then $\bigoplus_{k=1}^m h_k^{-1}$ is continuous (respectively linear). \square

Given a linear flow, it is worthwhile to consider various direct sums of its component flows. Of particular interest will be when a flow is completely represented by the direct sum of some collection of component flows. This type of situation arises when one finds invariant subspaces $\{X_k\}_{k=1}^m$ such that $X_k \cap (\sum_{j=1}^{k-1} X_j + \sum_{j=k+1}^m X_j) = \{0\}$ for all $k \leq m$, while also satisfying the equation $\sum_{k=1}^m X_k = X$. In such a situation $\bigoplus_{k=1}^m X_k$ can be identified naturally with X . Whenever $X = \bigoplus_{k=1}^m X_k$ is written, it should be assumed that $X_k \cap (\sum_{j=1}^{k-1} X_j + \sum_{j=k+1}^m X_j) = \{0\}$ for all $k \leq m$ and $\sum_{k=1}^m X_k = X$. Furthermore, as such constructions are only of interest in the context of invariant subspaces, it should also be assumed that X_k is invariant for all $k \leq m$.

Proposition 2.28 *Let $X = \bigoplus_{k=1}^m X_k$, let $L \in \mathcal{L}(X)$, and let φ be a linear flow on X . If $\{X_k\}_{k=1}^m$ is a collection of invariant subspaces of X under L , then $L = \bigoplus_{k=1}^m L_k$. Similarly, if $\{X_k\}_{k=1}^m$ is a collection of invariant subspaces of X under φ , then $\varphi = \bigoplus_{k=1}^m \varphi_k$.*

Proof. This follows from the definition of $\bigoplus_{k=1}^m L_k$ and $\bigoplus_{k=1}^m \varphi_k$. For the operator case, note that $Lx = L \sum_{k=1}^m x_k = \sum_{k=1}^m Lx_k = \sum_{k=1}^m L_k x_k = (\bigoplus_{k=1}^m L_k)x$ for all $x \in X$. Similarly, $\varphi(t, x) = \varphi(t, \sum_{k=1}^m x_k) = \sum_{k=1}^m \varphi(t, x_k) = \sum_{k=1}^m \varphi_k(t, x_k) = (\bigoplus_{k=1}^m \varphi_k)(t, x)$ for all $(t, x) \in \mathbb{R} \times X$ as a consequence of the linearity of φ and φ_k for all $k \leq m$. \square

Note that the ordering of $\bigoplus_{k=1}^m X_k$ is irrelevant, as are the orderings of both $\bigoplus_{k=1}^m L_k$ and $\bigoplus_{k=1}^m \varphi_k$ as long as they are consistent with the ordering of $\bigoplus_{k=1}^m X_k$.

Lemma 2.27 and Proposition 2.28 combined demonstrate how one can classify flows based on their components.

Theorem 2.29 *Let φ and ψ be linear flows on $X = \bigoplus_{k=1}^m X_k = \bigoplus_{k=1}^m Y_k$ where $\{X_k\}_{k=1}^m$ and $\{Y_k\}_{k=1}^m$ are collections of subspaces of X invariant under φ and ψ respectively. If φ_k and ψ_k (here $\varphi_k = \varphi|_{X_k}$ and $\psi_k = \psi|_{Y_k}$) are (homeomorphically or linearly) equivalent for all $k \leq m$, then φ and ψ are (respectively homeomorphically or linearly) equivalent.*

Proof. This follows directly from Lemma 2.27 and Proposition 2.28. \square

As per the preceding note, the specific ordering of $\bigoplus_{k=1}^m \psi_k$ is irrelevant. If necessary, it can be reordered prior to applying Theorem 2.29.

To see the value of Theorem 2.29, consider the respective stable and unstable subspaces of two flows, φ and ψ . As a consequence of Proposition 2.25, showing $X_\varphi^+ + X_\varphi^- = X$ and $X_\psi^+ + X_\psi^- = X$ is sufficient to demonstrate that $X_\varphi^+ \oplus X_\varphi^- = X$ and $X_\psi^+ \oplus X_\psi^- = X$ respectively. Under these conditions, if one can show that φ^+ and φ^- are homeomorphically equivalent to ψ^+ and ψ^- respectively, then it immediately follows from Theorem 2.29 that φ and ψ are homeomorphically equivalent. It turns out that checking whether or not φ^+ and ψ^+ (and similarly φ^- and ψ^-) are homeomorphically equivalent is fairly straightforward, as will be seen in Section 3.2, so flows for which $X_\varphi^+ \oplus X_\varphi^- = X$ have their own special designation.

Definition 2.30 Let φ be a linear flow on some $X \neq \{0\}$. If $X = X_\varphi^+ \oplus X_\varphi^-$, then φ is *hyperbolic*. If $X_\varphi^+ \oplus X_\varphi^- = \{0\}$, then φ is *central*.

Note that, alternatively, φ is hyperbolic if $\varphi = \varphi^+ \oplus \varphi^-$ for nontrivial φ . Theorem 2.29 is of limited use if one cannot find a collection of invariant subspaces $\{X_k\}_{k=1}^m$ so that $X = \bigoplus_{k=1}^m X_k$. To this end a notion of reducibility is introduced.

Definition 2.31 $L \in \mathcal{L}(X)$ is *reducible* if there exist nontrivial invariant subspaces Y and Z of X such that $X = Y \oplus Z$; otherwise, L is *irreducible*. Similarly, a linear flow φ on X is *reducible* if there exist nontrivial invariant subspaces Y and Z of X such that $X = Y \oplus Z$; otherwise, φ is *irreducible*.

Note that L and φ are reducible if $L = M \oplus N$ and $\varphi = \psi \oplus \gamma$ respectively for some nontrivial bounded linear operators M and N and nontrivial linear flows ψ and γ . More generally, L and φ are reducible if $L = \bigoplus_{k=1}^m L_k$ and $\varphi = \bigoplus_{k=1}^m \varphi_k$ respectively for some $m > 1$ with nontrivial bounded linear operators L_k and linear flows φ_k .

Irreducible flows essentially operate on invariant subspaces that do not contain any further nontrivial invariant subspaces. In the previous section the relationship between operator invariance and flow invariance was investigated; in particular, Proposition 2.20 states that a flow is invariant on some subspace if and only if the operator that generates it is invariant on that subspace. A similar result holds for reducibility.

Proposition 2.32 *Let φ be a linear flow on X . Then φ is reducible if and only if L_φ is reducible.*

Proof. Suppose first that L_φ is reducible. Then there exist nontrivial subspaces Y and Z of X such that Y and Z are invariant under L_φ and $X = Y \oplus Z$. It follows immediately from Proposition 2.20 that both Y and Z are invariant under φ as well, so φ is reducible. The converse is similar. \square

As a consequence of the above proposition, φ is irreducible if and only if L_φ is irreducible. This proposition will prove useful in the next section as part of the discussion of irreducible flows. Similar to invariance, as a consequence of this proposition, spaces will often be referred to as reducible or irreducible without referencing a specific operator or flow, as it will generally be clear from the context.

There are some basic results that provide invariant decomposition of spaces given some operator. In light of Proposition 2.32, such results can be exploited to find decompositions for linear flows.

Proposition 2.33 *Let φ be a linear flow on X with $\mathbb{K} = \mathbb{C}$ such that L_φ has at least two distinct eigenvalues. Then φ is reducible.*

Proof. For each eigenvalue λ_k of L_φ , let $X_k := \ker(L_\varphi - \lambda_k I)^d$. By assumption these subspaces are nontrivial. It is easily verified that X_k is invariant under L_φ for all $k \leq m$, and it follows from a standard result of operator theory [7] that $X = \bigoplus_{k=1}^m X_k$, and so $L_\varphi = \bigoplus_{k=1}^m L_k$. As a consequence of the previous proposition, $\varphi = \bigoplus_{k=1}^m \varphi_k$ with nontrivial φ_k , and thus φ is reducible. \square

It follows immediately from the preceding proposition that for every complex irreducible flow φ , it must hold that $\sigma(L_\varphi) = \{\lambda\}$ for some $\lambda \in \mathbb{C}$. Abusing terminology somewhat, the unique λ associated with an irreducible φ will be referred to as the eigenvalue of φ . The situation is slightly more complicated when $\mathbb{K} = \mathbb{R}$, but the related real result is not needed in this thesis.

With the idea of simplifying flows by reducing them to their components, it is reasonable to attempt to reduce a flow as much as possible. Now any decomposition consisting entirely

of irreducible components by definition cannot be further reduced. It turns out that such a decomposition can be found for every linear flow. The requirement that the subspaces be nontrivial in Definition 2.31 is important here, as every flow would be a reduction of itself, so there would otherwise be no such thing as an irreducible flow. Requiring that the reduction be nontrivial forces the subspaces to each have their dimension be strictly less than that of the initial space.

Proposition 2.34 *Let φ be a nontrivial linear flow on X . There exists a finite collection $\{\varphi_k\}_{k=1}^m$ of nontrivial irreducible components of φ such that $\varphi = \bigoplus_{k=1}^m \varphi_k$.*

Proof. Let $\{\varphi_k\}_{k=1}^m$ be a collection of nontrivial components of φ satisfying $\varphi = \bigoplus_{k=1}^m \varphi_k$ and with the property that, for some fixed $n \in \mathbb{N}$, each φ_k is irreducible whenever $d_k > n$ where d_k is the dimension of X_k . Define a new finite collection $\{\tilde{\varphi}_j\}_{j=1}^{\tilde{m}}$ of components of φ based on $\{\varphi_k\}_{k=1}^m$ as follows. Include φ_k in $\{\tilde{\varphi}_j\}_{j=1}^{\tilde{m}}$ for all $k \leq m$ such that $d_k \neq n$. Moreover include all irreducible φ_k with $d_k = n$. For reducible φ_k with $d_k = n$, reduce them to nontrivial components $\varphi_{k,1}$ and $\varphi_{k,2}$ and include those as the final elements of $\{\tilde{\varphi}_j\}_{j=1}^{\tilde{m}}$. It is clear by construction that $\varphi = \bigoplus_{j=1}^{\tilde{m}} \tilde{\varphi}_j$. Now the elements of $\{\tilde{\varphi}_j\}_{j=1}^{\tilde{m}}$ with $d_j > n$ are exactly the elements of $\{\varphi_k\}_{k=1}^m$ with $d_k > n$, and are by assumption all irreducible. Moreover, the elements of $\{\tilde{\varphi}_j\}_{j=1}^{\tilde{m}}$ with $d_j = n$ are exactly the irreducible elements of $\{\varphi_k\}_{k=1}^m$ with $d_k = n$, as the reducible elements of $\{\varphi_k\}_{k=1}^m$ with $d_k = n$ were all reduced to nontrivial elements of $\{\tilde{\varphi}_j\}_{j=1}^{\tilde{m}}$ such that $d_j < n$. Thus the elements of $\{\tilde{\varphi}_j\}_{j=1}^{\tilde{m}}$ with $d_j = n$ are also all irreducible. Consequently, $\{\tilde{\varphi}_j\}_{j=1}^{\tilde{m}}$ is a finite collection of nontrivial components of φ satisfying $\varphi = \bigoplus_{j=1}^{\tilde{m}} \tilde{\varphi}_j$ and with the property that, for some fixed $n \in \mathbb{N}$, each $\tilde{\varphi}_j$ is irreducible whenever $d_j > n - 1$.

Now the collection $\{\tilde{\varphi}_j\}$ satisfies the above construction for $n = d$ trivially. By repeatedly applying the above procedure d times, one acquires a finite collection $\{\varphi_k\}_{k=1}^m$ of nontrivial components of φ satisfying $\varphi = \bigoplus_{k=1}^m \varphi_k$ and with the property that φ_k is irreducible whenever $d_k > 0$. But the elements of $\{\varphi_k\}_{k=1}^m$ are nontrivial so $d_k > 0$ for all $k \leq m$, and consequently $\{\varphi_k\}_{k=1}^m$ is a finite collection of nontrivial irreducible components of φ such that $\varphi = \bigoplus_{k=1}^m \varphi_k$. \square

With all the work that has gone into reducing flows to irreducible components, one might

hope that irreducible components are indeed simpler to work with than general linear flows. To demonstrate that this is indeed the case, irreducible flows are discussed in detail in the next section.

2.5 Irreducible Flows

So far flows have been discussed without fixing a basis or scalar field for the underlying normed space. It is often more elegant to work without a fixed basis, but it will become necessary in future sections for linear flows to be in certain desirable forms with respect to some basis. Now a consequence of Theorem 2.17 is that two linear flows are linearly equivalent if and only if one flow is the same as the other after an appropriate change of basis. Since linear equivalence is the strongest form of equivalence considered in this thesis, it follows that the choice of basis has no effect when it comes to considering diffeomorphic and homeomorphic equivalence. As such, one can always choose to work with a basis for which the generating operator is in (real or complex) Jordan canonical form (see [5] and [7]) and this motivates the following definition.

Definition 2.35 Let $L \in \mathcal{L}(X)$ with $\mathbb{K} = \mathbb{C}$ (respectively \mathbb{R}). Then L is of λ -Jordan-type if its matrix with respect to an appropriate basis consists of a single complex (respectively real) Jordan block with eigenvalue λ (respectively conjugate pair of eigenvalues $\{\lambda, \bar{\lambda}\}$). In other words for $\mathbb{K} = \mathbb{C}$, with respect to an appropriate basis, L is of the form

$$J_d(\lambda) = \begin{bmatrix} \lambda & 1 & 0 & & 0 & 0 \\ 0 & \lambda & 1 & & 0 & 0 \\ 0 & 0 & \lambda & & 0 & 0 \\ & & & \ddots & & \\ 0 & 0 & 0 & & \lambda & 1 \\ 0 & 0 & 0 & & 0 & \lambda \end{bmatrix}$$

for some $\lambda \in \mathbb{C}$. The situation is similar when $\mathbb{K} = \mathbb{R}$. In this case L is of the same form as

above with respect to an appropriate basis for some $\lambda \in \mathbb{R}$. Alternatively L is of the form

$$\tilde{J}_d(\lambda) = \begin{bmatrix} \Re\lambda & \Im\lambda & 1 & 0 & 0 & 0 \\ -\Im\lambda & \Re\lambda & 0 & 1 & 0 & 0 \\ 0 & 0 & \Re\lambda & \Im\lambda & 0 & 0 \\ 0 & 0 & -\Im\lambda & \Re\lambda & 0 & 0 \\ & & & & \ddots & \\ 0 & 0 & 0 & 0 & \Re\lambda & \Im\lambda \\ 0 & 0 & 0 & 0 & -\Im\lambda & \Re\lambda \end{bmatrix}$$

with respect to an appropriate basis for some $\lambda \in \mathbb{C} \setminus \mathbb{R}$.

The complex case is clearly much easier to work with. The problem with real Jordan blocks is that one cannot split conjugate pairs of eigenvalues while working solely with real linear transformations, so real Jordan blocks with nonreal eigenvalues end up being much more unwieldy. Still, real Jordan blocks fundamentally behave similarly to complex Jordan blocks, so many results that hold for complex linear operators also hold for real ones. The preferred strategy is to directly prove a result for the complex case, and then extend the result to the real case via complexification. The concept of complexification (along with the concept of realification) is discussed in more detail in Appendix C.

This previous definition is useful in that it completely characterizes irreducible flows once a basis is fixed.

Proposition 2.36 *Let $L \in \mathcal{L}(X)$. Then L is irreducible if and only if it is of λ -Jordan-type for some $\lambda \in \mathbb{C}$.*

Proof. For simplicity, assume $\mathbb{K} = \mathbb{C}$.

Fix an appropriate basis for X so that L is in Jordan canonical form. Explicitly, L is of

the form

$$\begin{bmatrix} \lambda_1 & \delta_1 & 0 & & 0 & 0 \\ 0 & \lambda_2 & \delta_2 & & 0 & 0 \\ 0 & 0 & \lambda_3 & & 0 & 0 \\ & & & \ddots & & \\ 0 & 0 & 0 & & \lambda_{d-1} & \delta_{d-1} \\ 0 & 0 & 0 & & 0 & \lambda_d \end{bmatrix}$$

with respect to this basis, where $\lambda_n \in \sigma(L)$ for all $n \leq d$ and, for all $n < d$, where $\delta_n = 0$ whenever $\lambda_n \neq \lambda_{n+1}$ and $\delta_n \in \{0, 1\}$ otherwise. If any of the δ_n are in fact zero, then L is clearly reducible. Suppose now that L is irreducible. It follows from the contrapositive of the preceding argument that $\delta_n = 1$ for all $n < d$. Moreover, as necessarily $\delta_n = 0$ whenever $\lambda_n \neq \lambda_{n+1}$, it must further be the case that $\lambda_1 = \lambda_n$ for all $n \leq d$. Consequently, L is a single Jordan Block with eigenvalue λ_1 , so L is of λ_1 -Jordan-type. This proves the ‘only if’ part of the proposition.

To prove the ‘if’ part, suppose that L is reducible, in which case there exist invariant subspaces X_1 and X_2 of X such that $X = X_1 \oplus X_2$. Fix an ordered basis $\{b_1, \dots, b_{d_1}, \tilde{b}_1, \dots, \tilde{b}_{d_2}\}$ where $\{b_1, \dots, b_{d_1}\}$ and $\{\tilde{b}_1, \dots, \tilde{b}_{d_2}\}$ are bases of X_1 and X_2 respectively. With respect to this basis L is clearly block diagonal with at least two blocks. It follows that the Jordan canonical form of L must consist of at least two blocks, and thus L cannot be of λ -Jordan-type for any $\lambda \in \mathbb{C}$. The argument when $\mathbb{K} = \mathbb{R}$ is similar if somewhat unwieldy. \square .

This result should not be too surprising when $\mathbb{K} = \mathbb{C}$ in light of Proposition 2.33. Considering the preceding proposition, it is not unreasonable to refer to the eigenvalue of an irreducible linear flow when $\mathbb{K} = \mathbb{R}$ as well, with the understanding that irreducible linear flows on \mathbb{R} really correspond to the conjugate pair $\{\lambda, \bar{\lambda}\}$ for $\lambda \in \mathbb{C} \setminus \mathbb{R}$.

Consequently every irreducible linear flow is generated by a single Jordan Block. But every linear flow is generated by an operator in Jordan canonical form with respect to an appropriate basis. Since the Jordan blocks of this operator correspond to irreducible components of the flow, the uniqueness of the Jordan decomposition forces a form of uniqueness on irreducible decompositions of the flow.

Theorem 2.37 *Every nontrivial linear flow φ on X has a decomposition $\varphi = \bigoplus_{k=1}^m \varphi_k$ for some $m \in \mathbb{N}$ where φ_k is a nontrivial irreducible component of φ for each $k \leq m$. Moreover, if $\bigoplus_{k=1}^l \psi_k$ is another such decomposition, then $l = m$ and (reordering $\{\psi_k\}_{k=1}^m$ as necessary) φ_k and ψ_k are linearly equivalent for all $k \leq m$.*

Proof. Proposition 2.34 demonstrates the existence of such a decomposition. For each $k \leq m$, let X_k be the invariant subspace of X associated with φ_k , and let L_k be the bounded linear operator that generates φ_k . Further, fix a basis \mathcal{B}_k for X_k for each $k \leq m$. By Proposition 2.36, L_k is a single Jordan Block of dimension d_k and eigenvalue λ_k for each $k \leq m$. It follows by taking the basis elements of each \mathcal{B}_k in order to form a basis \mathcal{B} for X that L_φ is in Jordan canonical form with respect to that basis.

One may perform the same procedure on $\{\psi_j\}_{j=1}^l$ to form a second basis $\tilde{\mathcal{B}}$. It follows from the uniqueness of the Jordan canonical form that $m = l$ and, upon reordering $\{\psi_k\}_{k=1}^m$ and $\tilde{\mathcal{B}}_k$ as necessary, the matrices of L_φ and L_ψ with respect to \mathcal{B} and $\tilde{\mathcal{B}}$ respectively must be identical. Now if $\{b_j\}_{j=1}^d$ and $\{\tilde{b}_j\}_{j=1}^d$ are the elements of the bases \mathcal{B} and $\tilde{\mathcal{B}}$ respectively in order, then the change of basis determined by the map $b_j \mapsto \tilde{b}_j$ for all $j \leq d$ is clearly a linear equivalence between φ_k and ψ_k when restricted to X_k for all $k \leq m$. This completes the proof. \square

As a consequence of the above theorem, given two irreducible decompositions, the irreducible components of each can be identified with each other in a natural way. This identification preserves important properties such as the dimension of each component (more accurately, the dimension of its associated invariant subspace) and in the case of $\mathbb{K} = \mathbb{C}$ the eigenvalue of each component. More generally, irreducible decompositions are unique up to a change of basis, so irreducible decompositions are unique with respect to properties of operators that are independent of a specific choice of basis.

Having characterized irreducible operators with respect to an appropriate choice of basis, it is worthwhile to consider the form of irreducible linear flows with respect to that basis. It follows from Proposition 2.36 that every irreducible linear flow on X with $\mathbb{K} = \mathbb{C}$ is of

the form

$$e^{tJ_d(\lambda)} = e^{t\lambda} e^{tJ_d(0)} = e^{t\lambda} \begin{bmatrix} 1 & t & \frac{1}{2}t^2 & \frac{1}{(d-2)!}t^{d-2} & \frac{1}{(d-1)!}t^{d-1} \\ 0 & 1 & t & \frac{1}{(d-3)!}t^{d-3} & \frac{1}{(d-2)!}t^{d-2} \\ 0 & 0 & 1 & \frac{1}{(d-4)!}t^{d-4} & \frac{1}{(d-3)!}t^{d-3} \\ & & & \ddots & \\ 0 & 0 & 0 & 1 & t \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

for some $\lambda \in \mathbb{C}$ with respect to an appropriate basis. A similar (though again less pleasant) situation arises when $\mathbb{K} = \mathbb{R}$, in which case the flow is either of the form described above for some $\lambda \in \mathbb{R}$ or is of the form $e^{t\Re(\lambda)} e^{t\tilde{J}(\lambda)} x$ for some $\lambda \in \mathbb{C} \setminus \mathbb{R}$ where $e^{t\tilde{J}(\lambda)}$ is similar to the above matrix, except that each $a_{i,j}$ in the matrix is replaced with the 2×2 block

$$a_{i,j} \cdot \begin{bmatrix} \cos(t\Im\lambda) & \sin(t\Im\lambda) \\ -\sin(t\Im\lambda) & \cos(t\Im\lambda) \end{bmatrix}$$

with the powers and factorials ranging only to $\frac{d}{2} - 1$ instead of $d - 1$. Written in this form, it is not hard to make some basic statements concerning the behaviour of irreducible linear flows.

Proposition 2.38 *Let φ be an irreducible linear flow on X with eigenvalue $\lambda \in \mathbb{C}$. If $\Re\lambda < 0$, then $X = X_\varphi^+$. Similarly, if $\Re\lambda > 0$, then $X = X_\varphi^-$. Finally, if $\Re\lambda = 0$, then $X_\varphi^+ = X_\varphi^- = \{0\}$.*

Proof. Fix a basis for X , and assume $\mathbb{K} = \mathbb{C}$. As in Proposition 2.36, the $\mathbb{K} = \mathbb{R}$ case is similar but more unwieldy due to the $\sin(t\Im\lambda)$ and $\cos(t\Im\lambda)$ terms that appear, and it will be omitted.

Suppose first that $\Re\lambda < 0$, and fix $x \in X$. Then φ has the form described above with respect to an appropriate basis. Now the product $e^{tJ}x$ can be viewed as a vector consisting of polynomials with respect to t . It follows that $\|e^{t\lambda} e^{tJ}x\| \leq \sum_{j=1}^d \|e^{t\Re\lambda} p_j(t)\|$ where the $p_j(t)$ are for each $j \leq d$ these polynomials. Clearly $\lim_{t \rightarrow +\infty} e^{t\Re\lambda} = 0$ and it follows that $\lim_{t \rightarrow +\infty} e^{t\Re\lambda} p_j(t) = 0$ for all $j \leq d$ as the exponential behaviour will dominate the polynomial behaviour regardless of the specific polynomials that occur. Thus $\lim_{t \rightarrow +\infty} \varphi(t, x) = 0$.

As x was arbitrary, it follows that $X = X_\varphi^+$. A similar argument demonstrates the second statement.

Suppose now that $\Re\lambda = 0$, and further let $x \in X$ such that $\lim_{t \rightarrow +\infty} \varphi(t, x) = 0$. In this case there is no exponential term, and for this limit to hold, it must be the case that $p_j(t)$ is constant zero for all $j \leq d$. But this can only be the case if $x = 0$. Thus $X_\varphi^+ = \{0\}$. A similar argument works for X_φ^- . \square

This result leads to the construction of a third useful invariant subspace of X to go along with the stable and unstable subspaces — a subspace generated by the parts of the flow that are neither stable nor unstable.

Definition 2.39 Let φ be a flow on X with irreducible decomposition $\bigoplus_{k=1}^m \varphi_k$. If K is the subset of $\{1, \dots, m\}$ consisting of those indices for which the irreducible component is neither stable nor unstable, then the *central subspace* of X , denoted X_φ^0 is given by $X_\varphi^0 = \bigoplus_{k \in K} X_k$, and the *central part* of φ , denoted φ^0 is given by $\bigoplus_{k \in K} \varphi_k$.

It follows from the construction of the central part, along with Theorem 2.37, that the behaviour of a linear flow is completely determined by its behaviour along its stable, unstable, and central subspaces.

Theorem 2.40 Let φ be a linear flow on X . Then $X = X_\varphi^- \oplus X_\varphi^0 \oplus X_\varphi^+$ and similarly $\varphi = \varphi^- \oplus \varphi^0 \oplus \varphi^+$.

Proof. This follows directly from the preceding definition and Proposition 2.38 via an irreducible decomposition of φ . \square

There are a couple of additional handy results that will be proved before moving on to discuss homeomorphic and diffeomorphic classification. To start with, it is desirable to have a norm estimate for flows. This norm estimate will be constructed based on a norm estimate for irreducible flows. The following proposition is a necessary step toward this goal.

Proposition 2.41 Let φ be an irreducible linear flow on X with $\mathbb{K} = \mathbb{C}$. Then, for any fixed $\epsilon \in \mathbb{R}^+$, with respect to an appropriate basis φ is of the form $e^{t\lambda} e^{t\epsilon J_d(0)} x$, where λ is the eigenvalue of φ .

Proof. Now L_φ can be written in the form of a single Jordan block with eigenvalue λ with respect to an appropriate basis by Proposition 2.36. It follows that $L_\varphi = \lambda I + J_d(0)$. But then $e^{tL_\varphi}x = e^{t\lambda I + tJ_d(0)}x = e^{t\lambda I}e^{tJ_d(0)}x = e^{t\lambda}e^{tJ_d(0)}x$ for all $(t, x) \in \mathbb{R} \times X$ and the result follows for $\epsilon = 1$.

To allow for arbitrary values of $\epsilon \in \mathbb{R}^+$, consider $T(\epsilon) := \text{diag}(1, \epsilon, \epsilon^2, \dots, \epsilon^{d-1})$ with respect to the chosen basis. Now $T^{-1}(\epsilon)L_\varphi T(\epsilon) = \lambda I + T^{-1}(\epsilon)J_d(0)T(\epsilon)$. But it is easily seen that $T^{-1}(\epsilon)J_d(0)T(\epsilon) = \epsilon J_d(0)$. As $T(\epsilon)$ is clearly an invertible bounded linear operator on X , the similarity transformation $T^{-1}(\epsilon)L_\varphi T(\epsilon)$ can be viewed as a change of basis, and it follows that $\varphi(t, x)$ is of the form $e^{t\lambda}e^{t\epsilon J_d(0)}x$ for all $(t, x) \in \mathbb{R} \times X$ with respect to the new basis. \square

It is often useful to consider an irreducible flow in the above form with $\epsilon = 1$. The ϵ is only needed for the norm estimate, so that $\epsilon \|J_d(0)\|$ can be made arbitrarily small.

Lemma 2.42 *Let φ be an irreducible linear flow on X with $\mathbb{K} = \mathbb{C}$ where λ is the eigenvalue of φ , and fix $r \in \mathbb{R}$ large enough that $r > \Re\lambda$. Then $\|\varphi(t, x)\| \leq e^{rt}\|x\|$ holds for all $t \geq 0$ and $x \in X$, and consequently $\|\varphi_t\| \leq e^{rt}$ for all $t \geq 0$.*

Proof. Fix $\delta \in \mathbb{R}^+$ sufficiently small that $r - \delta > \Re\lambda$. Then $\varphi(t, x) = e^{t\lambda}e^{t\epsilon J_d(0)}x$ under an appropriate basis with $\epsilon \in \mathbb{R}^+$ is chosen small enough that $\epsilon \|J_d(0)\| \leq \delta$, and consequently $\|\varphi(t, x)\| \leq e^{t\Re\lambda} \|e^{t\epsilon J_d(0)}\| \|x\| \leq e^{t(r-\delta)} e^{t\epsilon \|J_d(0)\|} \|x\| \leq e^{t(r-\delta)} e^{t\delta} \|x\| = e^{rt} \|x\|$ for all $t \geq 0$ and $x \in X$. \square

One might hope to extend this result to arbitrary linear flows by considering an irreducible decomposition. Unfortunately there is no guarantee that an arbitrary norm will behave well when a space is decomposed via an irreducible decomposition of a flow. With that said, given an arbitrary norm, one can always construct a new norm that does behave well when a space is decomposed in this fashion.

Proposition 2.43 *Let φ be a linear flow on X , and fix $r \in \mathbb{R}$ large enough that $r > \Re\lambda$ for all $\lambda \in \sigma(L_\varphi)$. Then there exists a norm $\|\cdot\|_D$ on X such that $\|\varphi(t, x)\|_D \leq e^{rt}\|x\|_D$ for all $t \geq 0$ and $x \in X$.*

Proof. Suppose first that $\mathbb{K} = \mathbb{C}$, and let $\bigoplus_{k=1}^m \varphi_k$ be an irreducible decomposition of φ . By the preceding lemma $\|\varphi_k(t, x_k)\| \leq e^{rt}\|x_k\|$ for all $t \geq 0$ and $x_k \in X_k$ and $k \leq m$. Now

define the map $\|\cdot\|_D : X \rightarrow \mathbb{R}$ by setting $\|x\|_D = \sum_{k=1}^m \|x_k\|$. It is easily verified that $\|\cdot\|_D$ is a norm on X ; moreover, $\|\varphi(t, x)\|_D = \sum_{k=1}^m \|\varphi_k(t, x_k)\| \leq \sum_{k=1}^m e^{rt} \|x_k\| = e^{rt} \|x\|_D$ for all $t \geq 0$ and $x \in X$.

In the case where $\mathbb{K} = \mathbb{R}$, consider the complexification $\varphi_{\mathbb{C}}$ of φ . As $\varphi_{\mathbb{C}}$ and φ have the same eigenvalues, by the above argument there exists a norm $\|\cdot\|_{\tilde{D}}$ on $X_{\mathbb{C}}$ which satisfies the property for $\varphi_{\mathbb{C}}$. Define $\|\cdot\|_D$ to be the restriction of $\|\cdot\|_{\tilde{D}}$ to X . It is easily verified that $\|\cdot\|_D$ is indeed a norm on X and $\|\varphi(t, x)\|_D = \|\varphi_{\mathbb{C}}(t, x)\|_{\tilde{D}} \leq e^{rt} \|x\|_{\tilde{D}} = e^{rt} \|x\|_D$ for all $t \geq 0$ and $x \in X$. \square

The other useful result that will be needed in the upcoming discussion of flow classifications is the invertibility of certain matrices. To start with, recall the matrix $T(\epsilon)$ from the proof of Proposition 2.41. This construction appears frequently (with ϵ replaced with t) in the process of discussing homeomorphic equivalence of nonhyperbolic flows, so it is worthwhile to give it a formal definition.

Definition 2.44 A *t-rescaling*, denoted $T_d(t)$, is the bounded linear operator on \mathbb{K}^d given by $T_d(t) := \text{diag}(1, t, t^2, \dots, t^{d-1})$.

Note that a $T_d(t)$ is invertible for all non-zero $t \in \mathbb{R}$, with $T_d^{-1}(t) = T_d(t^{-1})$. This construction ends up being useful in many contexts. In particular, consider the matrix form of an irreducible complex flow φ with respect to some appropriate basis so that φ_t is of the form

$$e^{t\lambda} \begin{bmatrix} 1 & t & \frac{1}{2}t^2 & & \frac{1}{(d-2)!}t^{d-2} & \frac{1}{(d-1)!}t^{d-1} \\ 0 & 1 & t & & \frac{1}{(d-3)!}t^{d-3} & \frac{1}{(d-2)!}t^{d-2} \\ 0 & 0 & 1 & & \frac{1}{(d-4)!}t^{d-4} & \frac{1}{(d-3)!}t^{d-3} \\ & & & \ddots & & \\ 0 & 0 & 0 & & 1 & t \\ 0 & 0 & 0 & & 0 & 1 \end{bmatrix}$$

for all $t \in \mathbb{R}$. Given any $z \in \mathbb{Z}$ and constant $d \times d$ matrix M , one may introduce t into M with the diagonal t pattern similar to the above matrix — that is, with the element at position (j_1, j_2) multiplied by $t^{z+j_2-j_1}$ — by considering $t^z T_d^{-1}(t) M T_d(t)$. Conversely, given any $d \times d$ matrix $M(t)$ with this diagonal t pattern, one may apply $t^{-z} T_d(t) M(t) T_d^{-1}(t)$ to

remove the t pattern, resulting in a constant matrix. As an example of the usefulness of t -rescalings, the notion is used in the proof of the following proposition.

Proposition 2.45 *Let $n \in \mathbb{N}$ and let $l \in \mathbb{N}_0$ with $l \leq n$. Then the matrix*

$$M_{n,l} = \begin{bmatrix} \frac{1}{n!} & \frac{1}{(n+1)!} & \frac{1}{(n+2)!} & \frac{1}{(n+l)!} \\ \frac{1}{(n-1)!} & \frac{1}{n!} & \frac{1}{(n+1)!} & \frac{1}{(n+l-1)!} \\ \frac{1}{(n-2)!} & \frac{1}{(n-1)!} & \frac{1}{n!} & \frac{1}{(n+l-2)!} \\ & & & \ddots \\ \frac{1}{(n-l)!} & \frac{1}{(n-l+1)!} & \frac{1}{(n-l+2)!} & \frac{1}{n!} \end{bmatrix}$$

is invertible.

Proof. Instead of working with $M_{n,l}$ directly, consider $t^n T_{l+1}^{-1}(t) M_{n,l} T_{l+1}(t)$ for some nonzero $t \in \mathbb{R}$. It is easily seen that

$$\begin{bmatrix} p(t) \\ p'(t) \\ \vdots \\ p^{(l)}(t) \end{bmatrix} = \begin{bmatrix} c_0 \frac{1}{n!} t^n + c_1 \frac{1}{(n+1)!} t^{n+1} + \cdots + c_l \frac{1}{(n+l)!} t^{n+l} \\ c_0 \frac{1}{(n-1)!} t^{n-1} + c_1 \frac{1}{n!} t^n + \cdots + c_l \frac{1}{(n+l-1)!} t^{n+l-1} \\ \vdots \\ c_0 \frac{1}{(n-l)!} t^{n-l} + c_1 \frac{1}{(n-l+1)!} t^{n-l+1} + \cdots + c_l \frac{1}{n!} t^n \end{bmatrix} = 0$$

for any $c \in \mathbb{C}^{l+1}$ such that $t^n T_{l+1}^{-1}(t) M_{n,l} T_{l+1}(t) c = 0$. Now consider the l^{th} degree polynomial $q(t) := c_0 \frac{1}{n!} + c_1 \frac{1}{(n+1)!} t + \cdots + c_l \frac{1}{(n+l)!} t^l$. Now $0 = p(t) = t^n q(t)$, so $q(t) = 0$ as $t \neq 0$. But then $0 = p'(t) = nt^{n-1} q(t) + t^n q'(t) = t^n q'(t)$, so $q'(t) = 0$ as $t \neq 0$. Continuing in this fashion, $q^{(j)}(t) = 0$ for all $j \leq l$. It follows that q is constant zero, and thus $c = 0$. This shows that $t^n T_{l+1}^{-1}(t) M_{n,l} T_{l+1}(t)$ is invertible for all nonzero $t \in \mathbb{R}$. But for all nonzero $t \in \mathbb{R}$ clearly $t^n T_{l+1}^{-1}(t)$ and $T_{l+1}(t)$ are both invertible. Therefore $M_{n,l}$ is also invertible. \square .

It is easily seen that these factorial-type matrices arise as submatrices of the matrix form of an irreducible flow described on the previous page; see also Section 4.2.

3 Classification Theorems

This chapter builds on the ideas and results of the previous chapter to provide characterizations for several notions of flow equivalence. So far only linear equivalence has been considered in detail (Theorem 2.17) — diffeomorphic and homeomorphic equivalence have yet to be considered. It turns out that the stronger notions of flow equivalence display a substantial degree of rigidity; diffeomorphic equivalence in fact coincides with linear equivalence. The situation is not the same for homeomorphic equivalence. Characterizing homeomorphic equivalence is substantially more challenging, so that its discussion is split into two parts. A characterization for homeomorphic equivalence of hyperbolic flows is first provided, prior to the discussion of the general case.

3.1 Diffeomorphic Equivalence

Proposition 2.9 demonstrates that linear (that is, finest) equivalence implies diffeomorphic equivalence which in turn implies homeomorphic (that is, coarsest) equivalence. This simply reflects the fact that every linear map is its own best linear approximation and is thus differentiable, and every differentiable map is necessarily continuous; that is, $\mathcal{L}(X) \subset \mathcal{D}(X) \subset \mathcal{C}(X)$, where $\mathcal{D}(X)$ and $\mathcal{C}(X)$ are the spaces of differentiable and continuous maps on X respectively. These inclusions are in general strict, and it is tempting to extend this additional fact to the chain of equivalences — that is, one might assume that linear equivalence is strictly finer than diffeomorphic equivalence and so on. It turns out that this is not the case. The following lemma demonstrates that diffeomorphic equivalence is at least as fine as linear equivalence by generating a linear equivalence from the derivative of a diffeomorphic equivalence.

Lemma 3.1 *Let φ and ψ be linear flows on X . If φ and ψ are diffeomorphically equivalent, then they are linearly equivalent.*

Proof. Let h be a diffeomorphic equivalence between φ and ψ , and let H be the derivative of h at 0, so that $H := D_0h$. It follows immediately that $H \in \mathcal{L}(X)$. As a consequence of the chain rule $D_0hD_0h^{-1} = D_{h^{-1}(0)}hD_0h^{-1} = D_0(h \circ h^{-1}) = D_0I = I$ and similarly $D_0h^{-1}D_0h = I$. Thus H is invertible.

Now $h \circ e^{tL_\varphi} = e^{tL_\psi} \circ h$ for all $t \in \mathbb{R}$ since h is an equivalence. Again as a consequence of the chain rule $He^{tL_\varphi} = D_0h \circ e^{tL_\varphi} = D_0e^{tL_\psi} \circ h = e^{tL_\psi}H$ for all $t \in \mathbb{R}$. It follows that $e^{tHL_\varphi H^{-1}} = He^{tL_\varphi}H^{-1} = e^{tL_\psi}$ for all $t \in \mathbb{R}$. But then $L_\psi = HL_\varphi H^{-1}$ by Proposition 2.12, so L_φ is similar to L_ψ . Thus φ and ψ are linearly equivalent by Theorem 2.17. \square

This lemma in combination with Proposition 2.9 demonstrates that diffeomorphic equivalence and linear equivalence are in fact identical, so diffeomorphic equivalence may be characterized in the same fashion as linear equivalence.

Theorem 3.2 *Let φ and ψ be linear flows on X . The following are equivalent:*

- (i) φ and ψ are diffeomorphically equivalent;
- (ii) φ and ψ are linearly equivalent;
- (iii) L_φ and L_ψ are similar.

Proof. (i) \implies (ii) follows immediately from the previous lemma. (ii) \implies (i) follows from Proposition 2.9. Finally, (ii) \iff (iii) is just Theorem 2.17. \square

When it comes to equivalences stronger than homeomorphic, diffeomorphic equivalence is as fine as it gets; homeomorphic equivalence classes may break into several diffeomorphic equivalence classes, but smoother equivalence does nothing to break up those classes further. The above theorem essentially reduces the problem of linearly, diffeomorphically, and homeomorphically classifying flows to the linear and homeomorphic cases. Since linear equivalence has already been characterized, all that remains is to characterize homeomorphic equivalence.

Before examining homeomorphic equivalence in detail, it is worthwhile to point out that diffeomorphic equivalence is in fact strictly finer than homeomorphic equivalence. This is true even in the case of hyperbolic linear flows, and moreover even when the space is one-dimensional. It is not hard to construct an example of two linear flows that are homeomorphically but not diffeomorphically equivalent.

Example 3.3 Consider flows $e^t x$ and $e^{3t} x$ on \mathbb{R} . The induced linear operators in this case have eigenvalues 1 and 3 respectively, so the flows cannot be diffeomorphically equivalent.

On the other hand, the map $h : \mathbb{R} \rightarrow \mathbb{R}$ given by $h(x) = x^3$ is clearly a homeomorphism with

$$h(e^t x) = e^{3t} x^3 = e^{3t} h(x)$$

for all $x \in \mathbb{R}$ and $t \in \mathbb{R}$. It follows that the two flows are homeomorphically equivalent. \diamond

3.2 Homeomorphic Equivalence of Hyperbolic Flows

Homeomorphic equivalence has so far only been discussed briefly, namely in Proposition 2.24, where it was demonstrated that a homeomorphic equivalence between two flows preserves their respective stable and unstable subspaces. These subspaces were discussed in more detail in the previous chapter. This section will build on Proposition 2.24 in conjunction with various properties of the stable and unstable subspaces toward developing a characterization of homeomorphic equivalence. Note that this section is based heavily on Chapter 13 of [1].

In particular recall Proposition 2.38 — irreducible components of a linear flow φ are either stable, unstable, or central if their eigenvalue λ satisfies $\Re\lambda < 0$, $\Re\lambda > 0$, or $\Re\lambda = 0$ respectively. Consequently, φ^+ and φ^- consist of all the components of φ for which the eigenvalue has a negative or positive real part respectively. It follows that $\dim X_\varphi^-$ is just the sum of the sizes of the stable component flows, and similarly for $\dim X_\varphi^+$. In this fashion Propositions 2.24 and 2.38 can be combined to produce a partial characterization of homeomorphic equivalence.

Proposition 3.4 *Let φ and ψ be two linear flows on X . If φ and ψ are homeomorphically equivalent, then the following hold:*

- (i) $\dim X_\varphi^+ = \dim X_\psi^+$ and $\dim X_\varphi^- = \dim X_\psi^-$;
- (ii) L_φ and L_ψ have the same number of eigenvalues (counting algebraic multiplicities) with negative real part and the same number of eigenvalues (counting algebraic multiplicities) with positive real part.

Proof. Since homeomorphisms preserve dimension [4], it follows from Proposition 2.24 that $\dim X_\varphi^+ = \dim h(X_\varphi^+) = \dim X_\psi^+$ and similarly that $\dim X_\varphi^- = \dim X_\psi^-$.

For (ii), consider first φ^+ . It follows from Proposition 2.38 that the irreducible components of φ in φ^+ are exactly the components for which the real part of their eigenvalue is negative. Given an irreducible decomposition, let $K \subseteq \{1, \dots, m\}$ be the indexes of these components, so $\varphi^+ = \bigoplus_{k \in K} \varphi_k$. As a consequence of Proposition 2.36 the φ_k correspond to the Jordan blocks of L_φ of size d_k for which the eigenvalue has negative real part. It follows that the number of eigenvalues with negative real part (counting algebraic multiplicities) is the dimension of φ^+ ; that is, the number of eigenvalues with negative real part (counting algebraic multiplicities) is $\dim X_\varphi^+$. A similar statement holds for the number of eigenvalues of negative real part for L_ψ , and thus it follows from (i) that L_φ and L_ψ must have the same number of eigenvalues (counting algebraic multiplicities) with negative real part. A similar argument for eigenvalues with positive real part completes the proof of (ii). \square

This may not seem like a strong foundation from which to construct a characterization of homeomorphic equivalence — after all, knowledge of the dimensions of a couple of subspaces can hardly be considered a substantial insight. Indeed, it is not particularly challenging to construct two linear flows that are not homeomorphically equivalent despite having appropriately-sized stable and unstable subspaces.

Example 3.5 Fix a basis for \mathbb{C}^3 , consider the flows generated by

$$L = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & i \end{bmatrix} \quad \text{and} \quad M = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 2i \end{bmatrix},$$

and fix the point $(t, x) = (\pi, (0, 0, 1)) \in \mathbb{R} \times \mathbb{C}^3$. Further, let h be any homeomorphism of \mathbb{C}^3 . On one hand $h(e^{tM}x) = h((0, 0, e^{2i\pi})) = h((0, 0, 1)) = h(x)$, but on the other hand $e^{tL}h(x) = h(x)$ only if $h(x) = 0$. Now consider $y = (0, 0, 2)$. Since h is invertible $h(y) \neq 0$, and thus $e^{tL}h(y) \neq h(y)$. Unfortunately, as in the case of x , y satisfies $h(e^{tM}y) = h(y)$. It follows that h is not an equivalence, let alone a homeomorphic equivalence. Since h was an arbitrary homeomorphism, the flows generated by L and M are not homeomorphically equivalent, despite each having one-dimensional stable and unstable subspaces. \diamond

Beyond showing that the converse of Proposition 3.4 does not hold in general, the preceding example serves to illustrate how periodic portions of a flow can make even basic

equivalence impossible — note that the continuity of h was never required in the previous example. The periodic aspects of two flows must line up in a very specific way to even allow for the possibility of homeomorphic equivalence. This will be discussed in more detail in future sections.

With that said, the flows in the preceding example are clearly not hyperbolic. Somewhat surprisingly, the converse of Proposition 3.4 actually holds for hyperbolic linear flows; in other words, identically-sized stable and unstable subspaces imply homeomorphic equivalence in this case. As such, homeomorphic equivalence of hyperbolic linear flows can be completely characterized based solely on the dimensions of stable and unstable subspaces. Proving this fact rigorously is non-trivial, but the idea behind the proof is relatively straightforward.

It turns out that the nontrivial paths of irreducible hyperbolic linear flows are all spirals about the origin; there is always an exponential component that ends up dominating the behaviour of the flow. This exponential component arises from the real part of the eigenvalue of the induced operator (since the real part is nonzero) and the sign of the real part determines whether the spirals are directed toward or away from the origin. More generally, the nontrivial paths of a stable or unstable flow are also (somewhat more complicated) spirals toward or away from the origin.

The idea then is that one can homeomorphically straighten the nontrivial paths of the stable (respectively unstable) part of a flow and then unstraighten those paths into the nontrivial paths of the stable (respectively unstable) part of another flow, assuming the two parts are the same size. The process of straightening paths consists of demonstrating that every nontrivial path intersects a unit sphere at exactly one point. Assuming this is indeed possible, one may take any point in $X \setminus \{0\}$ to that sphere while storing the time t required to get there. One then may proceed by following for time t the path of the flow generated by either $-I$ or I (whose paths consist of straight lines to or from the origin respectively) as appropriate from the sphere.

Recall Proposition 2.43, which states that for every linear flow φ and $r \in \mathbb{R}$ satisfying $r > \Re \lambda$ for all $\lambda \in \sigma(L_\varphi)$ there exists a norm $\|\cdot\|_D$ such that $\|\varphi(t, x)\|_D \leq e^{rt}\|x\|_D$ for all $t \geq 0$ and $x \in X$. One demonstrates that every nontrivial path of a stable or unstable linear

flow — that is, a flow with $\Re\lambda < 0$ or $\Re\lambda > 0$ for all $\lambda \in \sigma(L_\varphi)$ respectively — intersects a unit sphere at exactly one point using this norm and its associated unit sphere.

Lemma 3.6 *Let φ be a stable or unstable linear flow on X . There exists a norm $\|\cdot\|_D$ on X so that the map $\tilde{\varphi} : \mathbb{R} \times \mathbb{S}_D \rightarrow X \setminus \{0\}$ given by $\tilde{\varphi}(t, x) = \varphi(t, x)$ for all $(t, x) \in \mathbb{R} \times \mathbb{S}_D$ is a homeomorphism.*

Proof. Assume first that φ is stable. Then $\Re\lambda < 0$ for all $\lambda \in \sigma(L_\varphi)$, and fix $r \in \mathbb{R}^+$ such that $\lambda < -r$ for all $\lambda \in \sigma(L_\varphi)$. Proposition 2.43 provides a norm $\|\cdot\|_D$ on X with the property that $\|\varphi(t, x)\|_D \leq e^{-rt}\|x\|_D$ for all $t \geq 0$ and $x \in X$, and this norm is in fact the desired norm, although that remains to be proved. Now it follows that $\|e^{tL_\varphi}\|_D \leq e^{-rt}$ for all $t \geq 0$. But then $\|x\|_D = \|e^{tL_\varphi}e^{-tL_\varphi}x\|_D \leq e^{-rt}\|e^{-tL_\varphi}x\|_D$ for all $t \geq 0$ and $x \in X$, so $e^{rt}\|x\|_D \leq \|e^{-tL_\varphi}x\|_D$ for all $t \geq 0$ and $x \in X$. Replacing t with $-t$ in the preceding equation results in the reverse inequality $\|\varphi(t, x)\|_D \geq e^{-rt}\|x\|_D$ for all $t \leq 0$ and $x \in X$. To summarize, the following inequalities hold:

- (i) $\|e^{tL_\varphi}x\|_D \leq e^{-rt}\|x\|_D$ for all $t \geq 0$ and $x \in X$ and
- (ii) $\|e^{tL_\varphi}x\|_D \geq e^{-rt}\|x\|_D$ for all $t \leq 0$ and $x \in X$.

Fix $x \in X \setminus \{0\}$. As a consequence of (i), one may choose a $t \in \mathbb{R}$ sufficiently large that $\|e^{tL_\varphi}x\|_D < 1$. Similarly, one may choose a $t \in \mathbb{R}$ sufficiently small that $\|e^{tL_\varphi}x\|_D > 1$ as a consequence of (ii). It follows from the continuity of φ that $\|e^{tL_\varphi}x\|_D = 1$ for some $t \in \mathbb{R}$. Since x was arbitrary, every non-trivial path of φ intersects \mathbb{S}_D at at least one point. Now fix $x \in \mathbb{S}_D$. The above two inequalities imply that $\|e^{tL_\varphi}x\|_D < 1$ for all $t > 0$ and $\|e^{tL_\varphi}x\|_D > 1$ for $t < 0$. Since x was again arbitrary, it follows that every non-trivial path of φ intersects \mathbb{S}_D at at most one point, so $\|e^{tL_\varphi}x\|_D = 1$ for exactly one $t \in \mathbb{R}$ for every $x \in X \setminus \{0\}$.

Again, fix $x \in X \setminus \{0\}$, and fix $t \in \mathbb{R}$ such that $\|e^{tL_\varphi}x\|_D = 1$. It follows that $\tilde{\varphi}(-t, \varphi(t, x)) = \varphi(-t, \varphi(t, x)) = x$ for $(-t, \varphi(t, x)) \in \mathbb{R} \times \mathbb{S}_D$. Consequently $\tilde{\varphi}$ is surjective, since x was arbitrary. On the other hand, consider points $(s, x), (t, y) \in \mathbb{R} \times \mathbb{S}_D$ such that $\tilde{\varphi}(s, x) = \tilde{\varphi}(t, y)$. Then $\|e^{-sL_\varphi}\tilde{\varphi}(s, x)\|_D = \|\varphi(-s, \varphi(s, x))\|_D = \|x\|_D = 1$ and $\|e^{-tL_\varphi}\tilde{\varphi}(s, x)\|_D = \|e^{tL_\varphi}\tilde{\varphi}(t, y)\|_D = \|\varphi(-t, \varphi(t, y))\|_D = \|y\|_D = 1$, and, by the above,

it must be that $s = t$. It follows that $x = \varphi(-s, \varphi(s, x)) = \varphi(-t, \varphi(t, y)) = y$ since $\varphi(s, x) = \tilde{\varphi}(s, x) = \tilde{\varphi}(t, y) = \varphi(t, y)$. Thus $\tilde{\varphi}$ is injective, as (s, x) and (t, y) were arbitrary.

It remains to show that $\tilde{\varphi}^{-1}$ is continuous, as the continuity of $\tilde{\varphi}$ follows immediately from the continuity of φ . Consider a sequence $\{x_n\}_{n \in \mathbb{N}} \subseteq X \setminus \{0\}$ converging to some $x \in X \setminus \{0\}$. Necessarily there exists a sequence $\{(t_n, y_n)\}_{n \in \mathbb{N}} \subseteq \mathbb{R} \times \mathbb{S}_D$ with $x_n = \tilde{\varphi}(t_n, y_n)$ for all $n \in \mathbb{N}$.

Now consider a subsequence $\{(t_k, y_k)\}$ of $\{(t_n, y_n)\}_{n \in \mathbb{N}}$. It follows from (i) and (ii) that $\|x_k\|_D = \|e^{t_k L_\varphi} y_k\|_D \leq e^{-rt_k} \|y_k\|_D = e^{-rt_k}$ whenever $t_k \geq 0$ and similarly it follows that $\|x_k\|_D = \|e^{t_k L_\varphi} y_k\|_D \geq e^{-rt_k} \|y_k\|_D = e^{-rt_k}$ whenever $t_k \leq 0$, where $\{x_k\}$ is the induced subsequence of $\{x_n\}_{n \in \mathbb{N}}$. As $\{x_k\}$ is a subsequence of $\{x_n\}_{n \in \mathbb{N}}$ it also converges to $x \in X \setminus \{0\}$, and thus for sufficiently large k it must be that $\|x_k\|_D$ is greater than, say, $\frac{1}{2}\|x\|_D$. But since $\|x_k\|_D \leq e^{-rt_k}$ for all $t_k \geq 0$ there must exist an $M \in \mathbb{R}^+$ such that $t_k \leq M$ for sufficiently large k . Similarly, it must be that $\|x_k\|_D$ is less than, say, $2\|x\|_D$ for sufficiently large k , and since $\|x_k\|_D \geq e^{-rt_k}$ for all $t_k \leq 0$ there must also exist an $m \in \mathbb{R}^-$ such that $t_k \geq m$ for sufficiently large k . Consequently $\{t_k\}$ is bounded.

Since \mathbb{S}_D is compact, and since $\{t_k\}$ is bounded, there exists a further subsequence $\{(t_j, y_j)\}$ of $\{(t_k, y_k)\}$ that is convergent to some $(t, y) \in \mathbb{R} \times \mathbb{S}_D$. But the continuity of φ forces $(t, y) = \tilde{\varphi}^{-1}(x)$. Since the subsequence $\{(t_k, y_k)\}$ was arbitrary, every subsequence of $\{(t_n, y_n)\}_{n \in \mathbb{N}} = \{\tilde{\varphi}^{-1}(x_n)\}_{n \in \mathbb{N}}$ has a further subsequence converging to $\tilde{\varphi}^{-1}(x)$, and consequently $\{\tilde{\varphi}^{-1}(x_n)\}_{n \in \mathbb{N}}$ converges to $\tilde{\varphi}^{-1}(x)$. As $\{x_n\}_{n \in \mathbb{N}}$ was itself arbitrary, $\tilde{\varphi}^{-1}$ is continuous. This completes the proof for the stable case.

Suppose now that φ is unstable. Under an appropriate choice of $r \in \mathbb{R}^+$ Proposition 2.43 can be applied to the stable flow $e^{t(-L_\varphi)}x$. It follows as above that $\|e^{t(-L_\varphi)}x\|_D \leq e^{-rt}\|x\|_D$ for all $t \geq 0$ and $x \in X$, and similarly $\|e^{t(-L_\varphi)}x\|_D \geq e^{-rt}\|x\|_D$ for all $t \leq 0$ and $x \in X$. But then $\|e^{tL_\varphi}x\|_D \leq e^{rt}\|x\|_D$ for all $t \leq 0$ and $x \in X$, and similarly $\|e^{tL_\varphi}x\|_D \geq e^{rt}\|x\|_D$ for all $t \geq 0$ and $x \in X$. These two inequalities can be applied as in the stable case to complete the proof for the unstable case. \square

Applying the inverse of $\tilde{\varphi}$ as defined in the preceding lemma amounts to compressing the nontrivial paths of a stable or unstable flow down to the points of an appropriate unit sphere. It is tempting to now simply apply the preceding lemma to the flows generated by

$-I$ and I and then compose the two resulting homeomorphisms, but this approach has a potential flaw: the unit sphere used for the first homeomorphism may not be the same as the unit sphere used for the second one. Ultimately, since every unit sphere is homeomorphic (as a consequence of the fact that all norms are equivalent for finite dimensional normed spaces) this does not pose a significant problem. In fact the problem can be completely circumvented, as it is easily seen that the flows generated by $-I$ and I intersect any unit sphere at exactly one point. This is due to the fact that for any norm $\|e^{\pm tI}x\| = e^{\pm t}\|x\|$ for all $(t, x) \in \mathbb{R} \times X$.

Lemma 3.7 *For any norm $\|\cdot\|$ on X the maps $h_+, h_- : \mathbb{R} \times \mathbb{S} \rightarrow X \setminus \{0\}$ given by $h_+(t, x) = e^{-tI}x$ and $h_-(t, x) = e^{tI}x$ respectively for all $(t, x) \in \mathbb{R} \times \mathbb{S}$ are homeomorphisms.*

Proof. It is easily verified that h_+ and h_- are clearly bijections where h_+^{-1} and h_-^{-1} are given by $h_+^{-1}(x) = (-\ln\|x\|, e^{-\ln\|x\|I}x)$ and $h_-^{-1}(x) = (\ln\|x\|, e^{-\ln\|x\|I}x)$ respectively for all $x \in X \setminus \{0\}$. Moreover, h_+^{-1} and h_-^{-1} are clearly both continuous, and h_+ and h_- are both continuous as a restriction of the flows $e^{-tI}x$ and $e^{tI}x$. \square

At this point, since Lemma 3.7 works for any unit sphere, Lemmas 3.6 and 3.7 can be combined to produce a homeomorphism of the nontrivial paths of a stable or unstable flow into the nontrivial paths of either of the flows generated by $-I$ and I . This homeomorphism can be extended to the entire space by simply fixing the origin, although the extension may fail to be continuous at that point. It turns out that continuity at the origin depends on the relative directions of the two flows. The extended map will fail to be continuous at the origin if the nontrivial paths of the two flows are in opposite directions — that is, one flow has nontrivial paths directed toward the origin while the nontrivial paths of the other flow are directed away from the origin. For example, the extension of the homeomorphism between a stable flow and the flow generated by I will fail to be bicontinuous at the origin. On the other hand, the extended map will be bicontinuous at the origin if the nontrivial paths of the two flows are in the same direction; better yet, the extended map in this case turns out to be a homeomorphic equivalence.

Proposition 3.8 *Every stable or unstable flow on X is homeomorphically equivalent to the flow generated by $-I$ or I respectively.*

Proof. Suppose that φ is a stable flow on X . Further let $\|\cdot\|_D$ be the norm on X guaranteed to exist by Lemma 3.6, and consider the map $h : X \rightarrow X$ given by $h(0) = 0$ and $h(x) = (h_+ \circ \tilde{\varphi}^{-1})(x)$ for all $x \in X \setminus \{0\}$, where $\tilde{\varphi}$ and h_+ are the homeomorphisms guaranteed by Lemmas 3.6 and 3.7 respectively. It follows from these lemmas that $h|_{X \setminus \{0\}}$ is a homeomorphism, and it is further clear that h is a bijection. It remains to be shown that h is bicontinuous at 0 and is an equivalence.

Fix $\epsilon \in \mathbb{R}^+$, and consider that $h(x) = e^{-\tau(x)I}\chi(x)$ for all $x \in X \setminus \{0\}$, where $\tau : X \rightarrow \mathbb{R}$ and $\chi : X \rightarrow \mathbb{S}_D$ are the component functions of $\tilde{\varphi}^{-1} : X \rightarrow \mathbb{R} \times \mathbb{S}_D$. Now $\|e^{-tI}x\|_D = e^{-t}$ for all $t \in \mathbb{R}$ and $x \in \mathbb{S}_D$. As such, there exists a $t_\epsilon \in \mathbb{R}^+$ such that $e^{-tI}x \in B_\epsilon(0)$ for all $t \in (t_\epsilon, +\infty)$ and $x \in \mathbb{S}_D$. Consequently, if it is possible to choose $\delta \in \mathbb{R}^+$ such that $\|x\|_D \in B_\delta(0) \setminus \{0\}$ implies $\tau(x) \in (t_\epsilon, +\infty)$, then $\|x\|_D \in B_\delta(0) \setminus \{0\}$ further forces $h(x) = e^{-\tau(x)I}\chi(x) \in B_\epsilon(0)$, as $\chi(x) \in \mathbb{S}_D$ for all $x \in X \setminus \{0\}$.

It turns out that it is always possible to choose such a δ . Fix $\delta = \|e^{-t_\epsilon L_\varphi}\|_D^{-1}$. Then $\|e^{-t_\epsilon L_\varphi}x\|_D < 1$ for all $x \in B_\delta(0)$, and, applying Proposition 2.43 with an appropriate choice of $r \in \mathbb{R}^+$, it follows that $\|e^{(t-t_\epsilon)L_\varphi}x\|_D \leq \|e^{tL_\varphi}\|_D \|e^{-t_\epsilon L_\varphi}x\|_D < e^{-rt} \leq 1$ for all $t \geq 0$ and $x \in B_\delta(0)$. The strict inequality in the above equation means that in particular $\|e^{(t-t_\epsilon)L_\varphi}x\|_D \neq 1$ for all $t \in \mathbb{R}_0^+$ and $x \in B_\delta(0)$. Since $\|e^{-\tau(x)L_\varphi}x\|_D = 1$ from the definition of τ , it must be that $-\tau(x) \neq t - t_\epsilon$ for all $t \geq 0$ and $x \in B_\delta(0) \setminus \{0\}$. It follows that $\tau(x) \in (t_\epsilon, +\infty)$ for all $x \in B_\delta(0) \setminus \{0\}$, and by the above argument, and since ϵ was arbitrary, h is continuous. Continuity of h^{-1} is similar.

To see that $h(\varphi(t, x)) = e^{-tI}h(x)$, consider that this is trivially true for all $t \in \mathbb{R}$ when $x = 0$. Using τ and χ as defined above

$$\begin{aligned} h(\varphi(t, x)) &= h(\varphi(t, e^{\tau(x)L_\varphi}\chi(x))) \\ &= h(e^{(t+\tau(x))L_\varphi}\chi(x)) \\ &= e^{-(t+\tau(x))I}\chi(x) \\ &= e^{-tI}e^{-\tau(x)I}\chi(x) \\ &= e^{-tI}h(x) \end{aligned}$$

for all $(t, x) \in \mathbb{R} \times X \setminus \{0\}$. Thus h is a homeomorphic equivalence. The unstable case proceeds similarly with some minor adjustments, as in (the proof of) Lemma 3.6. \square

Proposition 3.8, along with Lemmas 3.6 and 3.7, are constructive — the desired maps are built explicitly as part of the proofs. As such, it is tempting to attempt to construct the homeomorphic equivalence of Proposition 3.8 directly given a specific stable or unstable flow. A problem arises here in that the equivalence constructed in Proposition 3.8 is based on the inverse of the homeomorphism constructed in Lemma 3.6, rather than the homeomorphism directly, and providing an explicit representation of that inverse can be non-trivial. The following example demonstrates this, and it also shows how this issue can be overcome.

Example 3.9 Consider the stable flow φ generated by

$$L_\varphi = \begin{bmatrix} -2 & 1 \\ 0 & -2 \end{bmatrix}$$

on \mathbb{C}^2 with the standard basis and norm. It is unnecessary to work with a specially constructed norm in this case (as the flow is irreducible, so the required norm estimate follows directly from Lemma 2.42) and by Lemma 3.6 the map $\tilde{\varphi} : \mathbb{R} \times \mathbb{S} \rightarrow X \setminus \{0\}$ given by $\tilde{\varphi}(t, x) = \varphi(t, x)$ is a homeomorphism. The inverse then is given by $\tilde{\varphi}^{-1} = (\tau(x), \chi(x))$ for some $\tau : X \setminus \{0\} \rightarrow \mathbb{R}$ and $\chi : X \setminus \{0\} \rightarrow \mathbb{S}$. χ is given by $\chi(x) = e^{-\tau(x)L_\varphi}x$, but what is τ ? One could attempt to solve for τ using the norm, as

$$\begin{aligned} e^{-\tau(x)L_\varphi}x &= e^{2\tau(x)}e^{-\tau(x)J(0)}x \\ &= e^{2\tau(x)} \begin{bmatrix} 1 & -\tau(x) \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \\ &= e^{2\tau(x)}(x_1 - \tau(x)x_2, x_2) \end{aligned}$$

so $e^{4\tau(x)}\|(x_1 - \tau(x)x_2, x_1)\|^2 = \|e^{-\tau(x)L_\varphi}x\|^2 = 1$. This works out to be the product of a polynomial and an exponential and cannot be easily solved algebraically. Instead of working from φ to $e^{-tI}x$, consider working from $e^{-tI}x$ to φ — that is, consider h^{-1} rather than h . In this case the inverse map is coming from $e^{-tI}x$, and in particular $\tau(x) = -\ln\|x\|$ in this case, as $\|e^{-\ln\|x\|I}x\| = 1$ for all nonzero $x \in X$. Thus h^{-1} is given by $h(0) = 0$ and $h^{-1}(x) = e^{-\ln\|x\|L_\varphi}e^{-\ln\|x\|I}x = e^{-\ln\|x\|}e^{-\ln\|x\|L_\varphi}x$ for all nonzero $x \in X$. \diamond

At this point it is possible to combine Propositions 3.4 and 3.8 to fully characterize homeomorphic equivalence of hyperbolic linear flows.

Theorem 3.10 *Let φ and ψ be hyperbolic linear flows on X . The following are equivalent:*

- (i) φ and ψ are homeomorphically equivalent;
- (ii) $\dim X_\varphi^+ = \dim X_\psi^+$;
- (iii) $\dim X_\varphi^- = \dim X_\psi^-$.

Proof. (i) \implies (ii) and (i) \implies (iii) follow from Proposition 3.4. (ii) \iff (iii) is a simple consequence of the fact that $X_\varphi^+ \oplus X_\varphi^- = X = X_\psi^+ \oplus X_\psi^-$ as φ and ψ are both hyperbolic. To see that both (ii) and (iii) imply (i), suppose $d_1 = \dim X_\varphi^+ = \dim X_\psi^+$ and $d_2 = \dim X_\varphi^- = \dim X_\psi^-$. Now φ^+ and ψ^+ are both homeomorphically equivalent to the flow generated by $-I_{d_1}$ by Proposition 3.8, so by transitivity there exists a homeomorphic equivalence h^+ between φ^+ and ψ^+ . Similarly there exists a homeomorphic equivalence h^- between φ^- and ψ^- . As φ and ψ are both hyperbolic, $\varphi = \varphi^+ \oplus \varphi^-$ and $\psi = \psi^+ \oplus \psi^-$, and it follows from Theorem 2.29 that φ and ψ are homeomorphically equivalent. \square

Theorem 3.10 can easily be modified to characterize homeomorphic equivalence of hyperbolic linear flows based on properties of the induced operators rather than the unstable and stable subspaces.

Corollary 3.11 *Two hyperbolic linear flows φ and ψ on X are homeomorphically equivalent if and only if the number of eigenvalues (counting algebraic multiplicity) of L_φ with negative (respectively positive) real part is the same as the number of eigenvalues (counting algebraic multiplicity) of L_ψ with negative (respectively positive) real part.*

Proof. This follows immediately from Theorem 3.10 and the proof of Proposition 3.4. \square

3.3 Homeomorphic Equivalence of General Linear Flows

With homeomorphic equivalence fully characterized for hyperbolic linear flows, it makes sense to step back and consider homeomorphic equivalence of general linear flows. Many results of the previous section do not require a hyperbolic flow, and thus still apply in the general case. For example, Proposition 3.4 guarantees that the stable subspaces of

two homeomorphically equivalent linear flows have the same dimension, and this is also true for the unstable subspaces. Conversely, Proposition 3.8 guarantees that if the stable subspaces of two linear flows have the same dimension, then the stable parts of the two flows are homeomorphically equivalent, and the situation is similar for the unstable subspaces. This is enough to completely characterize homeomorphic equivalence between flows that are hyperbolic, as such flows are completely determined by their stable and unstable subspaces.

The problem that arises in the case of general linear flows is their central subspaces. It is an immediate consequence of Proposition 3.4 that the central subspaces of two homeomorphically equivalent linear flows must have the same dimension. Unlike the stable and unstable subspaces, the converse does not hold in general; that is, even if the central subspaces of two linear flows have the same dimension, the central parts of the two flows need not be homeomorphically equivalent. Recall Example 3.5, for instance. In that example it was shown that the linear flows on \mathbb{C}^3 generated by

$$L = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & i \end{bmatrix} \quad \text{and} \quad M = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 2i \end{bmatrix}$$

were not homeomorphically equivalent. But the stable and unstable parts of these flows are clearly homeomorphically equivalent, so it must be that the central parts are not. But it is also clear that the central subspaces of the two flows have the same dimension.

As such, Theorem 3.10 cannot be trivially extended to general linear flows. With that said, it turns out that there is a substantial degree of rigidity when it comes to the central parts of two homeomorphically equivalent linear flows. It will be proved in the next chapter that, given any linear flow φ on X with $\mathbb{K} = \mathbb{C}$, there exists a family $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ of subspaces of X such that for any homeomorphically equivalent linear flow ψ on X with equivalence h the following two properties hold:

- (i) $h(\text{BC}_{n,t}(\varphi)) = \text{BC}_{n,t}(\psi)$ for all $n \in \mathbb{N}_0$ and $t \in \mathbb{R}^+$ and
- (ii) the dimension of $\text{BC}_{n,t}(\varphi)$ is the number of irreducible components of φ of dimension greater than n with eigenvalue either 0 or $\frac{z2\pi i}{t}$ for some $z \in \mathbb{Z}$, and similarly for $\text{BC}_{n,t}(\psi)$.

Specifically, this is Theorem 4.12. Now since homeomorphisms preserve dimension [4], (i) guarantees that the number of irreducible components of ψ of size greater than n with eigenvalue either 0 or $\frac{z2\pi i}{t}$ for some nonzero $z \in \mathbb{Z}$ is the same as that of φ . Note that both 0 and $\frac{z2\pi i}{t}$ for all $z \in \mathbb{Z}$ and $t \in \mathbb{R}^+$ lie on $i\mathbb{R}$, so only irreducible components of the central part of a flow are ever counted by elements of this family. Conversely, every irreducible component of the central part of a flow is counted by some element of this family, as any $\lambda \in i\mathbb{R} \setminus \{0\}$ can be written as $\pm \frac{2\pi i}{t}$ by setting $t = \frac{2\pi}{|\lambda|}$. In fact, if the dimensions of the elements of $\{BC_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ are known, then it is possible to almost completely determine the irreducible decomposition of φ^0 by taking advantage of property (ii).

The general idea behind this is straightforward, though the details can become finicky. To start with, since φ is finite-dimensional, one can always select a t so that for all irreducible components the eigenvalue does not satisfy $\lambda = \frac{z2\pi i}{t}$ for any nonzero $z \in \mathbb{Z}$. In this way only irreducible components with eigenvalue 0 will be counted by $BC_{n,t}(\varphi)$ for this choice of t . Now $BC_{n,t}(\varphi)$ is trivial for all $n \geq d$, as it is impossible for an irreducible component to have dimension greater than that of the total space. As such $\dim BC_{d-1,t}(\varphi)$ is not simply the number of irreducible components of dimension greater than $d-1$; rather, $\dim BC_{d-1,t}(\varphi)$ is the number of irreducible components of dimension exactly d with eigenvalue 0. But then $\dim BC_{d-2,t}(\varphi) - \dim BC_{d-1,t}(\varphi)$ is the number of irreducible components of dimension exactly $d-1$ with eigenvalue 0, $\dim BC_{d-3,t}(\varphi) - \dim BC_{d-2,t}(\varphi)$ is the number of irreducible components of dimension exactly $d-2$ with eigenvalue 0, and so on. Continuing in this fashion one can determine the number of irreducible components of every dimension with eigenvalue 0 until reaching $\dim BC_{0,t}(\varphi)$, the number of irreducible components of any dimension with eigenvalue 0.

At this point one can select a different t — a t for which an irreducible component with a nonzero eigenvalue is counted by $BC_{n,t}(\varphi)$, at least for sufficiently small $n \in \mathbb{N}_0$. Taking care to subtract the number of irreducible components of each dimension with eigenvalue 0 from the new totals (as these components appear no matter the choice of t) one can then determine the number of irreducible components of each dimension with this new eigenvalue. This procedure is repeated for various values of t until the dimensions of all the irreducible components of φ^0 are known. Some care is needed when it comes to choosing each new value

of t , as multiple nonzero eigenvalues can line up with the same t . An example illustrates this technique in action.

Example 3.12 Suppose φ is a linear flow on X with $\mathbb{K} = \mathbb{C}$. As X is finite-dimensional, φ has a finite collection of eigenvalues, say $\{0, -i, i, 2i, 3i, -4i, 5i, 6i\}$. As irreducible components of φ with eigenvalue 0 will always contribute nontrivially to $\text{BC}_{n,t}(\varphi)$ for all $t \in \mathbb{R}^+$ and small $n \in \mathbb{N}_0$, the first goal is to choose a $t_0 \in \mathbb{R}^+$ so that only the irreducible components with eigenvalue 0 can possibly contribute to $\text{BC}_{n,t_0}(\varphi)$ for any $n \in \mathbb{N}_0$. Consider $t_0 = \frac{2\pi}{7}$. If φ_k is an irreducible component of φ with eigenvalue λ_k , then φ_k can contribute nontrivially to $\text{BC}_{n,t_0}(\varphi)$ only if either $\lambda_k = 0$ or $\lambda_k = \frac{z2\pi i}{t_0} = z7i$ for some nonzero $z \in \mathbb{Z}$. The latter case is clearly impossible here, as otherwise $|\lambda_k| = |7z| \geq 7$.

Now φ must have at least one irreducible component with eigenvalue 0, so it must be the case that $\dim \text{BC}_{n,t_0}(\varphi) > 0$ for some small $n \in \mathbb{N}_0$, say $\dim \text{BC}_{0,t_0}(\varphi) = 1$ and $\dim \text{BC}_{n,t_0}(\varphi) = 0$ for all $n \geq 1$. In this case φ has exactly one irreducible component of dimension one with eigenvalue 0 and no other irreducible components of any other dimension with eigenvalue 0. Now fix $t_1 = \frac{2\pi}{6}$. Clearly the irreducible component with eigenvalue 0 contributes to $\text{BC}_{n,t_1}(\varphi)$ (at least when $n = 0$) but what other eigenvalues potentially contribute? In this case the irreducible components with eigenvalue $6i$ also contribute for sufficiently small $n \in \mathbb{N}_0$ as $6i = \frac{6 \cdot 2\pi i}{2\pi} = \frac{2\pi i}{t_1}$. No other irreducible components φ_k can contribute, for otherwise $|\lambda_k| = |6z| \geq 6$. Again, φ must have at least one irreducible component with eigenvalue $6i$, so it must be the case that $\dim \text{BC}_{n,t_0}(\varphi) > 0$ for some small $n \in \mathbb{N}_0$, say $\dim \text{BC}_{0,t_0}(\varphi) = 2$ and $\dim \text{BC}_{n,t_0}(\varphi) = 0$ for all $n \geq 1$. It follows in this case that φ also has exactly one irreducible component of dimension one with eigenvalue $6i$ (recall that the 2 is counting the irreducible component with eigenvalue 0 as well) and no other irreducible components of any other dimension with eigenvalue $6i$.

One continues in this fashion by setting $t_2 = \frac{2\pi}{5}$. Similar to above, only the irreducible components with eigenvalue either 0 or with norm greater than or equal to five can potentially contribute to $\text{BC}_{n,t_2}(\varphi)$. Now by construction irreducible components with eigenvalue $5i$ definitely contribute, at least for small n , as $5i = \frac{5 \cdot 2\pi i}{2\pi} = \frac{2\pi i}{t_2}$. In this case irreducible components with eigenvalue $6i$ cannot possibly contribute as that would imply $\frac{6}{5} = z$ for some nonzero $z \in \mathbb{Z}$. Now fix $t_3 = \frac{2\pi}{4}$. Repeating the previous arguments, one finds that

only the irreducible component with eigenvalue 0 and any irreducible components with eigenvalue $-4i$ can possibly contribute to $BC_{n,t_3}(\varphi)$.

Suppose $BC_{n,t_2}(\varphi)$ and $BC_{n,t_3}(\varphi)$ are such that φ has exactly one irreducible component of dimension 1 with eigenvalues $5i$ and $-4i$ respectively and no irreducible components of any other dimension with either eigenvalue. Now fix $t_4 = \frac{2\pi}{3}$. Following the preceding argument, the irreducible component with eigenvalue 0 and any irreducible component with eigenvalue $3i$ can possibly contribute to $BC_{n,t_3}(\varphi)$. No other irreducible components can possibly contribute, with the exception of irreducible components with eigenvalue $6i$, as $6i = \frac{2 \cdot 3 \cdot 2\pi i}{2\pi} = \frac{2 \cdot 2\pi i}{t_4}$. Now suppose $\dim BC_{0,t_4}(\varphi) = 3$ and $\dim BC_{n,t_4}(\varphi) = 0$ for all $n \geq 1$. Since it is already known that φ has exactly one irreducible component of dimension one with eigenvalue 0 and exactly one irreducible component of dimension one with eigenvalue $6i$, it must be the case that φ also has exactly one irreducible component of dimension one with eigenvalue $3i$ and no irreducible components of any other dimension with that eigenvalue.

Now fix $t_5 = \frac{2\pi}{2}$. It is easily verified in the same fashion as above that the irreducible components of φ that contribute to $BC_{n,t_5}(\varphi)$ are exactly those components with eigenvalues 0, $6i$, $-4i$, and $2i$. Suppose in this case that $\dim BC_{0,t_5}(\varphi) = 6$, $\dim BC_{1,t_5}(\varphi) = 3$, $\dim BC_{2,t_5}(\varphi) = 2$, $\dim BC_{3,t_5}(\varphi) = 2$, and $\dim BC_{n,t_5}(\varphi) = 0$ for all $n \geq 4$. Since in particular $\dim BC_{4,t_5}(\varphi) = 0$, there are no irreducible components of φ with dimension greater than 4 and eigenvalue either 0, $6i$, $-4i$, or $2i$. Since $\dim BC_{3,t_5}(\varphi) = 2$, it follows that there are exactly two irreducible components of φ of dimension 4 with eigenvalue either 0, $6i$, $-4i$, or $2i$. But it is already known that the only irreducible components of φ with eigenvalue either 0, $6i$, or $-4i$ have dimension 1. Consequently, φ must have exactly two irreducible components of dimension 4 with eigenvalue $2i$. Now consider that $\dim BC_{2,t_5}(\varphi) = 2$, so φ has two irreducible components of dimension greater than 3 with eigenvalue either 0, $6i$, $-4i$, or $2i$. As it is already known that φ has two irreducible components of dimension 4 with eigenvalue $2i$, it must be that there are no irreducible components of dimension 3 with eigenvalue $2i$ (or 0, $6i$, or $-4i$). Moving on to $n = 1$, as the only irreducible components that contribute to $BC_{1,t_5}(\varphi)$ are the two irreducible components of dimension 4 with eigenvalue $2i$, and as it is already known that there are no irreducible components

of dimension 2 with eigenvalue either 0, $6i$, or $-4i$, it must be the case that there is exactly one irreducible component of dimension 2 with eigenvalue $2i$. At this point there are three known irreducible components with eigenvalue $2i$ and exactly one irreducible component for eigenvalues 0, $6i$, and $-4i$ respectively. Since $\dim \text{BC}_{0,t_5}(\varphi) = 6$, it follows that there are no irreducible components of φ of dimension 1 with eigenvalue $2i$.

Finally, fix $t_6 = 2\pi$. In this case every eigenvalue of φ contributes to $\text{BC}_{n,t_6}(\varphi)$ for small $n \in \mathbb{N}_0$. But this is not a problem since the irreducible components of every dimension and all of the eigenvalues except for i and $-i$ are already known at this point. Suppose that $\dim \text{BC}_{0,t_5}(\varphi) = 10$, $\dim \text{BC}_{1,t_5}(\varphi) = 4$, $\dim \text{BC}_{2,t_5}(\varphi) = 2$, $\dim \text{BC}_{3,t_5}(\varphi) = 2$, and $\dim \text{BC}_{n,t_5}(\varphi) = 0$ for all $n \geq 4$. As per the preceding discussion, φ must have exactly one irreducible component of dimension 2 with eigenvalue either i or $-i$, and exactly three irreducible components of dimension 1 with eigenvalue either i or $-i$. Unfortunately, one cannot determine which of these irreducible components have eigenvalue i specifically. Still, it has been shown in this example how knowing $\dim \text{BC}_{n,t}(\varphi)$ for all $n \in \mathbb{N}_0$ and $t \in \mathbb{R}^+$ is sufficient to completely determine the number of irreducible components of the central part of φ along with their respective dimensions and eigenvalues up to complex conjugation. \diamond

The arguments of the preceding example can easily be reapplied to other flows as long as $\mathbb{K} = \mathbb{C}$, as the existence of $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ will only be proved for complex φ in Chapter 4.

Lemma 3.13 *Let φ be a flow on X with $\mathbb{K} = \mathbb{C}$, let $\bigoplus_{k=1}^m \varphi_k$ be the irreducible decomposition of φ^0 , and set $b_{n,t} = \dim \text{BC}_{n,t}(\varphi)$ for all $n \in \mathbb{N}_0$ and $t \in \mathbb{R}^+$. Then m , d_k , and λ_k (up to complex conjugation) are completely determined by the family $\{b_{n,t}\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ for all $k \leq m$, where d_k and λ_k are the dimension and eigenvalue respectively of φ_k .*

Proof. As X is finite-dimensional, $\sigma(L_\varphi)$ is finite, so $\sigma(L_\varphi) \cap i\mathbb{R}$ is finite. Discard the eigenvalue 0 (if it appears) along with exactly one eigenvalue for each conjugate pair of eigenvalues in $\sigma(L_\varphi) \cap i\mathbb{R}$ (if such pairs appear) and arrange the remaining eigenvalues descending in norm, thus yielding $\{\tilde{\lambda}_j\}_{j=1}^l$. In this fashion $\tilde{\lambda}_1$ is (one of) the nonzero eigenvalue(s) of L_φ with maximum norm while $\tilde{\lambda}_l$ is (one of) the nonzero eigenvalue(s) of L_φ with minimum norm. From this construct $\{t_j\}_{j=0}^l$ by setting $t_0 = \frac{2\pi}{|\tilde{\lambda}_1|+1}$ and $t_j = \frac{2\pi}{|\tilde{\lambda}_j|}$

for all $1 \leq j \leq l$. If it is possible to completely determine m , d_k , and λ_k for all $k \leq m$ from the various $b_{n,t}$ with $n \in \mathbb{N}_0$ and $t \in \{t_j\}_{j=0}^l$ taken in order, then the proposition follows.

With $\{t_j\}_{j=0}^l$ constructed as above, b_{n,t_0} is the number of irreducible components of φ^0 (note that only $\lambda \in i\mathbb{R}$ can possibly satisfy $\lambda = \frac{z2\pi i}{t}$ for some $z \in \mathbb{Z}$ no matter the choice of $t \in \mathbb{R}^+$) of dimension greater than n with eigenvalue 0, and b_{n,t_j} is the number of irreducible components of φ^0 of dimension greater than n with eigenvalue λ in the set $\{0, \pm\tilde{\lambda}_1, \dots, \pm\tilde{\lambda}_j\} \cap \{z|\tilde{\lambda}_j|i : z \in \mathbb{Z}\}$ for $1 \leq j \leq l$. First suppose that $\lambda = \frac{z2\pi i}{t_0}$ for some nonzero $z \in \mathbb{Z}$ implies that $\lambda = z(|\tilde{\lambda}_1| + 1)i$ for some nonzero $z \in \mathbb{Z}$ by construction. It follows that $|\lambda| = n(|\tilde{\lambda}_1| + 1)$ for some $n \in \mathbb{N}$, so it must be that $|\lambda| \geq |\tilde{\lambda}_1| + 1$. But by construction $|\tilde{\lambda}_j| \leq |\tilde{\lambda}_1|$ for all $j \leq l$. Thus only irreducible components of dimension greater than n with eigenvalue 0 are counted by b_{n,t_0} . Now suppose $\lambda = \frac{z2\pi i}{t_j}$ for some nonzero $z \in \mathbb{Z}$ and $1 \leq j \leq l$. Then by construction $\lambda = z|\tilde{\lambda}_j|i$ for some nonzero $z \in \mathbb{Z}$. As $0 = 0|\tilde{\lambda}_j|$, this is sufficient to demonstrate the right-hand side of the intersection. To get the left-hand side, take the norm of the previous equation to get that $|\lambda| = n|\tilde{\lambda}_j|$ for some $n \in \mathbb{N}$, so $|\lambda| \geq |\tilde{\lambda}_j|$. This inequality only holds for $\lambda = \tilde{\lambda}_J$ with $J \leq j$ by construction.

In particular this means that b_{n,t_j} is the number of irreducible components of φ^0 of dimension greater than n with eigenvalue in some subset of $\{0, \pm\tilde{\lambda}_1, \dots, \pm\tilde{\lambda}_j\}$ for each $1 \leq j \leq l$. One may now completely determine the number of irreducible components of φ^0 along with their various dimensions and eigenvalues up to complex conjugation by induction on j . The initial step is to determine the number of irreducible components of φ^0 with eigenvalue 0 along with their various dimensions. To this end, consider b_{n,t_0} for $n \in \mathbb{N}_0$. For each $n \in \mathbb{N}_0$, since b_{n,t_0} and b_{n+1,t_0} are the number of irreducible components of φ^0 of dimension greater than n and $n+1$ respectively with eigenvalue 0, it follows that $b_{n,t_0} - b_{n+1,t_0}$ is the number of irreducible components of φ^0 of dimension exactly $n+1$ with eigenvalue 0. In this way the number of irreducible components of φ^0 of dimension n with eigenvalue 0 is completely determined for all $n \in \mathbb{N}$. Note that these numbers may be all zero if $0 \notin \sigma(L_\varphi)$.

Suppose now that the number of irreducible components of φ^0 of dimension n with eigenvalue $\lambda \in \{0, \pm\tilde{\lambda}_1, \dots, \pm\tilde{\lambda}_j\}$ is known for all $n \in \mathbb{N}$ for some fixed $1 \leq j < l$. It is then possible to determine the number of irreducible components of φ^0 of dimension n

with eigenvalue $\pm\tilde{\lambda}_{j+1}$ for all $n \in \mathbb{N}$ by considering the various b_{n,t_j} . Similar to above $b_{n,t_{j+1}} - b_{n+1,t_{j+1}}$ is the number of irreducible components of φ^0 of dimension exactly $n+1$ with eigenvalue in $\{0, \pm\tilde{\lambda}_1, \dots, \pm\tilde{\lambda}_{j+1}\} \cap \{z|\tilde{\lambda}_{j+1}|i : z \in \mathbb{Z}\}$ for all $n \in \mathbb{N}$. One then, for each $n \in \mathbb{N}_0$, determines the number of irreducible components of φ^0 of dimension exactly $n+1$ with eigenvalue $\tilde{\lambda}_{j+1}$ by subtracting the number of irreducible components of φ^0 of dimension exactly $n+1$ with eigenvalue in $\{0, \pm\tilde{\lambda}_1, \dots, \pm\tilde{\lambda}_j\} \cap \{z|\tilde{\lambda}_{j+1}|i : z \in \mathbb{Z}\}$, known by assumption from $b_{n,t_{j+1}} - b_{n+1,t_{j+1}}$.

Once the number of irreducible components of φ^0 of size n with eigenvalue λ is known for all $n \in \mathbb{N}$ and $\lambda \in \sigma(L_\varphi)$, then one immediately gets the total number of irreducible components of φ^0 . \square

Unfortunately the preceding procedure cannot distinguish between irreducible components of the same dimension with conjugate eigenvalues. This is due to the fact that, if $\lambda = \frac{z2\pi i}{t}$ for some nonzero $z \in \mathbb{Z}$, then $\bar{\lambda} = \frac{-z2\pi i}{t}$. This might seem like a let-down after all the work that went into distinguishing the various irreducible component flows in the first place. But for homeomorphic equivalence, this turns out to be a good thing, as two irreducible linear flows with complex conjugate eigenvalues are easily seen to be homeomorphically equivalent.

Proposition 3.14 *Let φ and ψ be irreducible linear flows on X and Y respectively where X and Y are normed spaces over \mathbb{C} with $\dim X = \dim Y$. If $\lambda_\varphi = \overline{\lambda_\psi}$, then φ and ψ are homeomorphically equivalent.*

Proof. Begin by fixing a basis $\{b_n\}_{n=1}^d$ for X such that φ is of the form $e^{t\lambda_\varphi} e^{tJ_d(0)}x$. This is possible by Proposition 2.41. Similarly fix a basis $\{\hat{b}_n\}_{n=1}^d$ for Y such that ψ is of the form $e^{t\lambda_\psi} e^{tJ_d(0)}x$. Now consider $h : X \rightarrow Y$ given by $h(x) = h(\sum_{n=1}^d c_n b_n) = \sum_{n=1}^d \bar{c}_n \hat{b}_n =: \bar{x}$ for all $x \in X$. Clearly h is invertible. Furthermore, h is a homeomorphism of X since the map $c \mapsto \bar{c}$ is continuous. Finally $h(e^{t\lambda_\varphi} e^{tJ_d(0)}x) = \overline{e^{t\lambda_\varphi} e^{tJ_d(0)}x} = e^{t\overline{\lambda_\varphi}} e^{tJ_d(0)}\bar{x} = e^{t\lambda_\psi} e^{tJ_d(0)}h(x)$ for all $(t, x) \in \mathbb{R} \times X$, as $J_d(0)$ is simply a Jordan block for eigenvalue 0 and as such has no nonreal entries. Thus h is a homeomorphic equivalence between $e^{t\lambda_\varphi} e^{tJ_d(0)}x$ and $e^{t\lambda_\psi} e^{tJ_d(0)}x$, and the result follows. \square

Of course, if φ and ψ are as in the previous lemma with $\lambda_\varphi = \lambda_\psi$, then they are not

only homeomorphically equivalent but in fact linearly equivalent via the homeomorphism $h(x) = h(\sum_{n=1}^d c_n b_n) = \sum_{n=1}^d c_n \hat{b}_n$.

Returning to the discussion of $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$, consider property (ii). If φ and ψ are two homeomorphically equivalent linear flows, then since homeomorphisms preserve dimension, $\dim \text{BC}_{n,t}(\varphi) = \dim h(\text{BC}_{n,t}(\varphi)) = \dim \text{BC}_{n,t}(\psi)$ for all $n \in \mathbb{N}_0$ and $t \in \mathbb{R}^+$. But this means that applying the previously described procedure to either flow will result in the same decomposition up to reordering and complex conjugation of individual irreducible components. It turns out that this is sufficient to completely characterize homeomorphic equivalence of linear flows. The following theorem and its corollaries are the main results of this thesis. It is necessary to use complexifications $\varphi_{\mathbb{C}}$ and $\psi_{\mathbb{C}}$ of φ and ψ respectively here as the family $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ is only constructed for complex flows φ ; that is, flows on X with $\mathbb{K} = \mathbb{C}$. Of course $\varphi_{\mathbb{C}} = \varphi$ and $\psi_{\mathbb{C}} = \psi$ for flows on X with $\mathbb{K} = \mathbb{C}$.

Theorem 3.15 *Let φ and ψ be linear flows on X . Then $\varphi_{\mathbb{C}}$ and $\psi_{\mathbb{C}}$ are homeomorphically equivalent if and only if the following three conditions hold:*

- (i) $\dim X_{\varphi_{\mathbb{C}}}^+ = \dim X_{\psi_{\mathbb{C}}}^+$;
- (ii) $\dim X_{\varphi_{\mathbb{C}}}^- = \dim X_{\psi_{\mathbb{C}}}^-$;
- (iii) *if $\bigoplus_{k=1}^m \varphi_k$ and $\bigoplus_{k=1}^m \psi_k$ are irreducible decompositions of $\varphi_{\mathbb{C}}^0$ and $\psi_{\mathbb{C}}^0$ respectively, then (reordering $\bigoplus_{k=1}^m \psi_k$ as necessary) $d_{\varphi_k} = d_{\psi_k}$ and either $\lambda_{\varphi_k} = \lambda_{\psi_k}$ or $\lambda_{\varphi_k} = \overline{\lambda_{\psi_k}}$ for every $k \leq m$.*

Proof. Suppose first that the three conditions hold. It follows from (i) and (ii) that there exist homeomorphic equivalences h^+ between $\varphi_{\mathbb{C}}^+$ and $\psi_{\mathbb{C}}^+$ and h^- between $\varphi_{\mathbb{C}}^-$ and $\psi_{\mathbb{C}}^-$ by Theorem 3.10. Consider property (iii). Now for φ_k and ψ_k with $\lambda_{\varphi_k} = \lambda_{\psi_k}$ it is clear that φ_k and ψ_k are homeomorphically (in fact, linearly) equivalent. For φ_k and ψ_k with $\lambda_{\varphi_k} = \overline{\lambda_{\psi_k}}$ it follows from Proposition 3.14 that φ_k and ψ_k are homeomorphically equivalent. Thus (iii) guarantees the existence of a homeomorphic equivalence h^0 between $\varphi_{\mathbb{C}}^0$ and $\psi_{\mathbb{C}}^0$. Consequently $h^+ \oplus h^0 \oplus h^-$ is a homeomorphic equivalence between $\varphi_{\mathbb{C}}$ and $\psi_{\mathbb{C}}$. This completes the proof of the ‘if’ case.

Suppose now that $\varphi_{\mathbb{C}}$ and $\psi_{\mathbb{C}}$ are homeomorphically equivalent. Properties (i) and (ii) follow directly from Proposition 3.4. To see that property (iii) also holds, consider

that $\dim \text{BC}_{n,t}(\varphi_{\mathbb{C}}) = \dim h(\text{BC}_{n,t}(\varphi)) = \dim \text{BC}_{n,t}(\psi)$ for all $n \in \mathbb{N}_0$ and $t \in \mathbb{R}^+$ by the construction of $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ and since homeomorphisms preserve dimension. Property (iii) then follows as applying Lemma 3.13 to $\varphi_{\mathbb{C}}$ and $\psi_{\mathbb{C}}$ completely determines the irreducible components of these flows up to conjugate eigenvalues. This completes the proof of the ‘only if’ part. \square

This theorem can be made more elegant by considering the cases where $\mathbb{K} = \mathbb{C}$ and $\mathbb{K} = \mathbb{R}$ separately. The $\mathbb{K} = \mathbb{C}$ case in particular follows almost directly from the previous theorem, and it improves on property (iii) of the theorem by using the realifications of the central subspaces.

Corollary 3.16 *Let φ and ψ be linear flows on X with $\mathbb{K} = \mathbb{C}$. Then φ and ψ are homeomorphically equivalent if and only if the following three conditions hold:*

- (i) $\dim X_{\varphi}^+ = \dim X_{\psi}^+$;
- (ii) $\dim X_{\varphi}^- = \dim X_{\psi}^-$;
- (iii) $\varphi_{\mathbb{R}}^0$ and $\psi_{\mathbb{R}}^0$ are linearly equivalent.

Proof. As $\varphi_{\mathbb{C}} = \varphi$ and $\psi_{\mathbb{C}} = \psi$, properties (i) and (ii) of this theorem are equivalent to properties (i) and (ii) of Theorem 3.15. Thus it suffices to show that property (iii) of this corollary is equivalent to property (iii) of Theorem 3.15. Note that $\varphi_{\mathbb{C}}^0 = \varphi^0$ and $\psi_{\mathbb{C}}^0 = \psi^0$, let $\bigoplus_{k=1}^m \varphi_k$ and $\bigoplus_{k=1}^m \psi_k$ be irreducible decomposition of φ^0 and ψ^0 respectively, and let $(\varphi_k)_{\mathbb{R}}$ and $(\psi_k)_{\mathbb{R}}$ be the realifications of φ_k and ψ_k respectively. Now if φ_k is of dimension d_k with eigenvalue 0, then $(\varphi_k)_{\mathbb{R}}$ is the direct sum of two irreducible components of $\varphi_{\mathbb{R}}^0$, each of dimension d_k with eigenvalue 0 and conversely. If φ_k is of dimension d_k with eigenvalue $\lambda_k \in i\mathbb{R} \setminus \{0\}$, then $(\varphi_k)_{\mathbb{R}}$ is an irreducible component of dimension $2d_k$ with conjugate eigenvalue pair $\{\lambda_k, \overline{\lambda_k}\}$ and conversely. The situation is the same for the components ψ_k . The desired result then follows by noting that $(\varphi_k)_{\mathbb{R}}$ and $(\psi_k)_{\mathbb{R}}$ have identical conjugate eigenvalue pairs, even if the eigenvalue of φ_k is not identical but merely conjugate to the eigenvalue of ψ_k . \square

Oftentimes working in complex spaces is easier and provides more elegant proofs and results than working in real spaces. For instance, \mathbb{C} is algebraically closed while \mathbb{R} is not.

All differentiable complex-valued functions on \mathbb{C} are analytic, but the same certainly cannot be said for real-valued functions on \mathbb{R} . Even in this thesis, the family $\{B_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ will only be constructed for flows φ on X with $\mathbb{K} = \mathbb{C}$, as dealing with real Jordan blocks with nonreal eigenvalues is a hassle. Surprisingly, the characterization of homeomorphic equivalence of real flows — that is, flows on X with $\mathbb{K} = \mathbb{R}$ — is actually more elegant than the complex case.

Corollary 3.17 *Let φ and ψ be linear flows on X with $\mathbb{K} = \mathbb{R}$. Then φ and ψ are homeomorphically equivalent if and only if the following three conditions hold:*

- (i) $\dim X_\varphi^+ = \dim X_\psi^+$;
- (ii) $\dim X_\varphi^- = \dim X_\psi^-$;
- (iii) φ^0 and ψ^0 are linearly equivalent.

Proof. The ‘if’ part is clear as properties (i) and (ii) guarantee homeomorphic equivalence between φ^+ and ψ^+ and between φ^- and ψ^- as per the proof of Theorem 3.15, while property (iii) guarantees homeomorphic equivalence between φ^0 and ψ^0 . For the ‘only if’ part, note that homeomorphic equivalence between φ and ψ guarantees that properties (i) and (ii) hold as an immediate consequence of Proposition 3.4, while also guaranteeing $\varphi_{\mathbb{C}}$ and $\psi_{\mathbb{C}}$ are homeomorphically equivalent. Thus it suffices to show that property (iii) of Theorem 3.15 implies property (iii) of this corollary.

Let $\bigoplus_{k=1}^m \varphi_k$ and $\bigoplus_{k=1}^m \psi_k$ be irreducible decompositions of $\varphi_{\mathbb{C}}^0$ and $\psi_{\mathbb{C}}^0$ respectively. Now each φ_k of dimension d_k with eigenvalue $\lambda_k \in i\mathbb{R} \setminus \{0\}$ is generated by an irreducible component of φ^0 of dimension $2d_k$ with conjugate eigenvalue pair $\{\lambda_k, \overline{\lambda_k}\}$, so the irreducible components φ_k of dimension d_k with eigenvalue $\lambda_k \in i\mathbb{R} \setminus \{0\}$ can be paired off exactly with the irreducible components φ_j of dimension $d_j = d_k$ with eigenvalue $\lambda_j = \overline{\lambda_k}$. The situation is the same for the irreducible components of $\psi_{\mathbb{C}}^0$. Reorder the φ_k and ψ_k as necessary so that the pair of irreducible components of $\varphi_{\mathbb{C}}$ corresponding to a single irreducible component of φ^0 are all together while preserving property (iii) of Theorem 3.15. Now if $n_{E,d,\varphi}$ is the number of irreducible components of $\varphi_{\mathbb{C}}^0$ of dimension d with eigenvalue in $E \subset \mathbb{C}$, then $n_{\{\lambda, \overline{\lambda}\}, d, \varphi} = 2n_{\{\lambda\}, d, \varphi} = 2n_{\{\overline{\lambda}\}, d, \varphi}$ by the preceding argument. The situation is the same

for $\psi_{\mathbb{C}}^0$. Now property (iii) of Theorem 3.15 is itself equivalent to the requirement that $n_{\{\lambda, \bar{\lambda}\}, d, \varphi} = n_{\{\lambda, \bar{\lambda}\}, d, \psi}$ for all $d \in \mathbb{N}$ and $\lambda \in i\mathbb{R} \setminus \{0\}$ and consequently $n_{\{\lambda\}, d, \varphi} = n_{\{\lambda\}, d, \psi}$ for all $d \in \mathbb{N}$ and $\lambda \in i\mathbb{R} \setminus \{0\}$. But this means that $\bigoplus_{k=1}^m \psi_k$ may be reordered so that $d_{\varphi_k} = d_{\psi_k}$ and $\lambda_{\varphi_k} = \lambda_{\psi_k}$ for all $k \leq m$; moreover, this reordering may be done so that pairs of φ_k generated by a single irreducible component of φ^0 line up with pairs of ψ_k generated by a single irreducible component of ψ^0 . Let $\bigoplus_{j=1}^l \hat{\varphi}_j$ and $\bigoplus_{j=1}^l \hat{\psi}_j$ be irreducible decompositions of φ^0 and ψ^0 respectively ordered based on the ordering of $\bigoplus_{k=1}^m \varphi_k$ and $\bigoplus_{k=1}^m \psi_k$; that is, for instance, $\hat{\varphi}_1$ generates either the first or the first pair of irreducible component(s) of $\varphi_{\mathbb{C}}^0$ as necessary, $\hat{\varphi}_2$ generates either the second or second pair of irreducible component(s) as necessary, and so on. Then by construction $d_{\hat{\varphi}_j} = d_{\hat{\psi}_j}$ and $\lambda_{\hat{\varphi}_j} = \lambda_{\hat{\psi}_j}$ for all $j \leq l$. It follows that φ^0 and ψ^0 are linearly equivalent. \square

Before closing out this chapter, it is worthwhile to consider characterizing the homeomorphic equivalence of linear flows in terms of their generating operators.

Corollary 3.18 *Let φ and ψ be linear flows on X with $\mathbb{K} = \mathbb{C}$. Then φ and ψ are homeomorphically equivalent if and only if the following three conditions hold:*

- (i) L_{φ} and L_{ψ} have the same number of eigenvalues with negative real part;
- (ii) L_{φ} and L_{ψ} have the same number of eigenvalues with positive real part;
- (iii) For each $n \in \mathbb{N}$ and $\lambda \in i\mathbb{R}$, L_{φ} and L_{ψ} have the same number of Jordan blocks of dimension n corresponding to λ or $\bar{\lambda}$.

Proof. This follows from Corollary 3.16. First note that φ^+ is the direct sum of all irreducible components of φ_k^+ with eigenvalue λ_k such that $\Re \lambda_k < 0$. Thus $\dim X_{\varphi^+}$ is the sum of the dimensions of the irreducible components φ_k^+ of φ^+ . But each irreducible component is a Jordan block $J_{d_k}(\lambda_k)$ with respect to an appropriate choice of basis. The situation is similar for ψ^+ , and it follows that property (i) of this corollary is equivalent to property (i) of Corollary 3.16. A similar argument demonstrates that property (ii) of this corollary is equivalent to property (ii) of Corollary 3.16. Finally, property (iii) of this corollary is also clearly equivalent to property (iii) of Corollary 3.16 by considering each irreducible component of φ^0 and ψ^0 as a single Jordan block. \square

The real case of the preceding corollary is similar.

Corollary 3.19 *Let φ and ψ be linear flows on X with $\mathbb{K} = \mathbb{R}$. Then φ and ψ are homeomorphically equivalent if and only if the following three conditions hold:*

- (i) *L_φ and L_ψ have the same number of eigenvalues with negative real part;*
- (ii) *L_φ and L_ψ have the same number of eigenvalues with positive real part;*
- (iii) *For each $n \in \mathbb{N}$ and $\lambda \in i\mathbb{R}$, L_φ and L_ψ have the same number of Jordan blocks of dimension n corresponding to λ .*

Proof. This corollary is proved exactly as the proof of Corollary 3.18, but using Corollary 3.17 in place of Corollary 3.16. \square

After proving the existence of the family $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ for flows φ on X with $\mathbb{K} = \mathbb{C}$ in Chapter 4, Chapter 5 includes a comparison between the complex and real classifications of linear flows on normed spaces with small dimension, from which it will be apparent that the complex situation is, well, more complex.

4 Nonhyperbolic Flows

A complete characterization of homeomorphic equivalence of linear flows was presented in Section 3.3. In demonstrating that characterization, the existence of a certain family of subspaces $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ was assumed. This family had two significant properties: every $\text{BC}_{n,t}(\varphi)$ was preserved by homeomorphic equivalence, and the dimension of each $\text{BC}_{n,t}(\varphi)$ was the number of irreducible component flows φ_k of dimension $d_k > n$ with eigenvalue λ_k satisfying either $\lambda_k = 0$ or $\lambda_k = \frac{z2\pi i}{t}$ for some nonzero $z \in \mathbb{Z}$. The purpose of this chapter is to prove the existence of this family. This chapter is heavily based on [9].

4.1 Basic Constructions

The family $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ of a given linear flow φ is constructed by combining simpler subspaces in various permutations. The desired properties of this family can then be demonstrated based on the properties of these simpler subspaces. The fundamental building blocks that will be used in this construction are the strong and weak centres of a given flow.

Definition 4.1 Let φ be a linear flow on X . The *weak centre* of φ on a subset Y of X , denoted $\text{WC}(\varphi, Y)$, is the set of all $y \in Y$ with the following property:

For every sequence $\{t_n\}_{n \in \mathbb{N}}$ in \mathbb{R} with $\lim_{n \rightarrow \infty} |t_n| = +\infty$ there exists a sequence $\{y_n\}_{n \in \mathbb{N}}$ in Y converging to y such that the sequence $\{\varphi(t_n, y_n)\}_{n \in \mathbb{N}}$ is bounded.

The *strong centre* of φ on Y , denoted $\text{SC}(\varphi, Y)$, is defined similarly to the weak centre, except that the sequence $\{\varphi(t_n, y_n)\}_{n \in \mathbb{N}}$ is required to converge to zero.

Clearly strong centres are contained in weak centres. It is a straightforward exercise to verify that, when Y is a subspace of X , the strong centre and weak centre of a linear flow are always themselves subspaces of that subspace. The specific nature of these subspaces will be discussed in detail in the next section, but their crucial aspect is that, under the right conditions, they end up being approximately half the size of the initial subspace.

Example 4.2 Consider the linear flow φ on \mathbb{R}^2 generated by

$$L = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \text{ so that } \varphi(t, x) = e^{tL}x = \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}.$$

Given any $x_1 \in \mathbb{R}$, and given any sequence $\{t_n\}_{n \in \mathbb{N}}$ in \mathbb{R} such that $\lim_{n \rightarrow \infty} |t_n| = +\infty$, the sequence $\{(x_1, -\frac{x_1}{t_n})\}_{n \in \mathbb{N}}$ converges to $(x_1, 0)$; moreover, as $\varphi(t_n, (x_1, -\frac{x_1}{t_n})) = (0, -\frac{x_1}{t_n})$ by direct calculation via the above matrix, $\lim_{n \rightarrow \infty} \varphi(t_n, (x_1, -\frac{x_1}{t_n})) = \lim_{n \rightarrow \infty} (0, -\frac{x_1}{t_n}) = (0, 0)$, so by definition $(x_1, 0) \in \text{SC}(\varphi, \mathbb{R}^2)$. As x_1 was also arbitrary, $\text{span}\{(1, 0)\} \subseteq \text{SC}(\varphi, \mathbb{R}^2)$.

On the other hand, suppose $(x_1, x_2) \in \text{WC}(\varphi, \mathbb{R}^2)$, and again fix a sequence $\{t_n\}_{n \in \mathbb{N}}$ in \mathbb{R} with $\lim_{n \rightarrow \infty} |t_n| = +\infty$. By definition there exists a sequence $\{(x_{1,n}, x_{2,n})\}_{n \in \mathbb{N}}$ in \mathbb{R}^2 converging to (x_1, x_2) such that $\{\varphi(t_n, (x_{1,n}, x_{2,n}))\}_{n \in \mathbb{N}}$ is bounded. In particular then the sequence $\{x_{1,n} + t_n x_{2,n}\}_{n \in \mathbb{N}}$, the first coordinate of the previous sequence, is also bounded. Since $\{t_n\}_{n \in \mathbb{N}}$ is unbounded, the sequence $\{x_{2,n}\}_{n \in \mathbb{N}}$ must converge to zero, so $x_2 = 0$. Thus $\text{WC}(\varphi, \mathbb{R}^2) \subseteq \text{span}\{(1, 0)\}$. Since also $\text{SC}(\varphi, \mathbb{R}^2) \subseteq \text{WC}(\varphi, \mathbb{R}^2)$, it therefore follows that $\text{SC}(\varphi, \mathbb{R}^2) = \text{WC}(\varphi, \mathbb{R}^2) = \text{span}\{(1, 0)\}$. \diamond

In the preceding example the strong and weak centres ended up being the same subspace, but this is not always the case.

Example 4.3 Consider the flow φ on \mathbb{R} given by $\varphi(t, x) = x$ for all $(t, x) \in \mathbb{R} \times \mathbb{R}$, and consider the sequence $\{x_n\}_{n \in \mathbb{N}}$ in \mathbb{R} given by $x_n = x$ for all $n \in \mathbb{N}$. Since φ_t is the identity for all $t \in \mathbb{R}$, the induced sequence $\{\varphi(t_n, x_n)\}_{n \in \mathbb{N}}$ is simply $\{x_n\}_{n \in \mathbb{N}}$ irrespective of the choice of $\{t_n\}_{n \in \mathbb{N}}$. This sequence is clearly bounded, so $x \in \text{WC}(\varphi, \mathbb{R})$. As x was arbitrary, it follows that $\text{WC}(\varphi, \mathbb{R}) = \mathbb{R}$. On the other hand, given any sequence $\{x_n\}_{n \in \mathbb{N}}$ in \mathbb{R} converging to some $x \in \mathbb{R}$, the induced sequence $\{\varphi(t_n, x_n)\}_{n \in \mathbb{N}}$ is simply $\{x_n\}_{n \in \mathbb{N}}$, as previously. It follows that for x to be in $\text{SC}(\varphi, \mathbb{R})$, it must be that $x = 0$, so $\text{SC}(\varphi, \mathbb{R}) = \{0\} \neq \mathbb{R} = \text{WC}(\varphi, \mathbb{R})$. \diamond

Since strong and weak centres on a subspace are always themselves subspaces, one may consider constructions involving combinations of strong and weak centres. For instance, one could consider the strong centre of a flow on the weak centre of the weak centre of that flow. The iterated centres of a flow are formed by combinations of this type.

Definition 4.4 Let φ be a linear flow on X , and let $n \in \mathbb{N}_0$. Then n may be uniquely written in the form $n = \delta_l 2^{l-1} + \dots + \delta_1 2^0$ where the $\delta_j \in \{0, 1\}$ (δ_l must be 1, except in

the case where $l = 1$ to allow for $n = 0$) are the digits of the binary representation of n , written $n = \langle \delta_l \delta_{l-1} \dots \delta_1 \rangle$. The n^{th} iterated centre of φ , denoted $\text{IC}_n(\varphi)$ is given by

$$\text{IC}_n(\varphi) := \tilde{\delta}_l(\varphi, \tilde{\delta}_{l-1}(\varphi, \dots \tilde{\delta}_2(\varphi, \tilde{\delta}_1(\varphi, X))))),$$

where $\tilde{\delta}_j(\varphi, Y) = \text{WC}(\varphi, Y)$ if $\delta_j = 0$ and $\tilde{\delta}_j(\varphi, Y) = \text{SC}(\varphi, Y)$ if $\delta_j = 1$ for any subspace Y of X .

Since strong and weak centres on a subspace are always themselves subspaces, and since the iterated centres are finite combinations of these constructions, it follows that the iterated centres are always subspaces as well. The iterated centres for $n \neq 0$ are made up of every finite combination of strong and weak centres that ends with a strong centre. The pattern is explicitly

n^{th} iterated centre	n in binary	Construction
$\text{IC}_0(\varphi)$	$\langle 0 \rangle$	$\text{WC}(\varphi, X)$
$\text{IC}_1(\varphi)$	$\langle 1 \rangle$	$\text{SC}(\varphi, X)$
$\text{IC}_2(\varphi)$	$\langle 10 \rangle$	$\text{SC}(\varphi, \text{WC}(\varphi, X))$
$\text{IC}_3(\varphi)$	$\langle 11 \rangle$	$\text{SC}(\varphi, \text{SC}(\varphi, X))$
$\text{IC}_4(\varphi)$	$\langle 100 \rangle$	$\text{SC}(\varphi, \text{WC}(\varphi, \text{WC}(\varphi, X)))$
$\text{IC}_5(\varphi)$	$\langle 101 \rangle$	$\text{SC}(\varphi, \text{WC}(\varphi, \text{SC}(\varphi, X)))$
$\text{IC}_6(\varphi)$	$\langle 110 \rangle$	$\text{SC}(\varphi, \text{SC}(\varphi, \text{WC}(\varphi, X)))$
$\text{IC}_7(\varphi)$	$\langle 111 \rangle$	$\text{SC}(\varphi, \text{SC}(\varphi, \text{SC}(\varphi, X)))$
$\text{IC}_8(\varphi)$	$\langle 1000 \rangle$	$\text{SC}(\varphi, \text{WC}(\varphi, \text{WC}(\varphi, \text{WC}(\varphi, X))))$

and so on. The idea behind iterated centres is that each repeated weak and strong centre reduces the size of the iterated centre, so that $\text{IC}_n(\varphi)$ becomes trivial for sufficiently large n . One can then work backwards until $\text{IC}_n(\varphi)$ is nontrivial, revealing the largest irreducible components of φ , and then continue to work backwards to find the next largest irreducible components, and so on.

Example 4.5 Consider Example 4.3. By definition $\text{IC}_0(\varphi) = \mathbb{R}$. Furthermore, it is clear that $\text{SC}(\varphi, \{0\}) = \{0\}$ and $\text{WC}(\varphi, \{0\}) = \{0\}$ (in fact, this is true for any φ) and it follows that $\text{IC}_n(\varphi) = \{0\}$ for all $n \in \mathbb{N}$, since a strong centre appears in the construction of all of these iterated centres.

Now consider Example 4.2 denoting the flow by ψ in this case. Again it immediately follows from the construction of IC_0 and IC_1 that $IC_0(\psi) = IC_1(\psi) = \text{span}\{(1, 0)\}$. Furthermore, $IC_2(\psi) = SC(\psi, WC(\psi, \mathbb{R}^2)) = SC(\psi, \text{span}\{(1, 0)\}) = \{0\}$ for the same reasons that $SC(\varphi, \mathbb{R})$ is trivial. More generally, $IC_n(\psi)$ will always end up with either $\text{span}\{(1, 0)\}$ or $\{0\}$ before the final strong centre, so that $IC_n(\psi) = \{0\}$ for all $n \geq 2$. \diamond

Although close to the goal, iterated centres are not yet enough to count the various central irreducible component flows as ultimately desired. The primary remaining issue is that iterated centres have no way of distinguishing components by their eigenvalue; two components of the same dimension and eigenvalue should be and are indistinguishable, but for homeomorphic equivalence to work as described in Section 3.3 it is necessary to distinguish two components of the same dimension but with different eigenvalues. This leads to the final construction of this section.

Definition 4.6 Let φ be a linear flow on X . The family of subspaces $\{BC_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ is given by

$$BC_{n,t}(\varphi) = IC_n(\varphi) \cap \ker(\varphi_t - I)$$

for all $n \in \mathbb{N}$ and $t \in \mathbb{R}^+$.

It is clear that this family can be constructed for any linear flow. Since iterated centres are all subspaces, and since these kernels are all clearly subspaces, it follows that all of these intersections are subspaces. As such, it is correct to refer to $\{BC_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ as a family of subspaces. Demonstrating that $\{BC_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ has the two desired properties discussed previously will be done by building upon properties of strong and weak centres. Having established properties of strong and weak centres, one may consider the consequences of those properties when strong and weak centres are combined to form iterated centres. From there, one considers the properties of $\ker(\varphi_t - I)$ in conjunction with the properties of iterated centres. This process is demonstrated in the proof of the following two propositions. First it will be shown that these constructions all behave well with respect to flow decomposition. Then it will be shown that these constructions also behave well with respect to homeomorphisms of the space.

Proposition 4.7 *Let φ be a linear flow on X with irreducible decomposition $\bigoplus_{k=1}^m \varphi_k$. The following all hold:*

- (i) $\text{WC}(\varphi, \bigoplus_{k=1}^m Y_k) = \bigoplus_{k=1}^m \text{WC}(\varphi_k, Y_k)$ for any collection $\{Y_k\}_{k=1}^m$ of subspaces of X with $Y_k \subseteq X_k$ for all $k \leq m$;
- (ii) $\text{SC}(\varphi, \bigoplus_{k=1}^m Y_k) = \bigoplus_{k=1}^m \text{SC}(\varphi_k, Y_k)$ for any collection $\{Y_k\}_{k=1}^m$ of subspaces of X with $Y_k \subseteq X_k$ for all $k \leq m$;
- (iii) $\text{IC}_n(\varphi) = \bigoplus_{k=1}^m \text{IC}_n(\varphi_k)$ for all $n \in \mathbb{N}_0$;
- (iv) $\ker(\varphi_t - I) = \bigoplus_{k=1}^m \ker((\varphi_k)_t - I)$ for all $t \in \mathbb{R}^+$;
- (v) $\text{BC}_{n,t}(\varphi) = \bigoplus_{k=1}^m \text{BC}_{n,t}(\varphi_k)$ for all $n \in \mathbb{N}_0$ and $t \in \mathbb{R}^+$.

Proof. (i) and (ii) are consequences of the fact that a sequence is bounded or converges to zero if and only if all its components are bounded or converge to zero respectively. For each $k \leq m$ let Y_k be a subspace of X_k , and suppose $x \in \text{WC}(\varphi, \bigoplus_{k=1}^m Y_k)$. Fix a sequence $\{t_n\}_{n \in \mathbb{N}}$ in \mathbb{R} so that $\lim_{n \rightarrow \infty} |t_n| = +\infty$. Then there exists a sequence $\{x_n\}_{n \in \mathbb{N}}$ in $\bigoplus_{k=1}^m Y_k$ converging to x such that the induced sequence $\{\varphi(t_n, x_n)\}_{n \in \mathbb{N}}$ is bounded. But then the sequence $\{\varphi_k(t_n, x_{n,k})\}_{n \in \mathbb{N}}$ is bounded for each $k \leq m$. Since each sequence $\{x_{n,k}\}_{n \in \mathbb{N}}$ in Y_k converges to x_k , and since $\{t_n\}_{n \in \mathbb{N}}$ was arbitrary, it follows that $x_k \in \text{WC}(\varphi_k, Y_k)$ for every $k \leq m$, and thus $x \in \bigoplus_{k=1}^m \text{WC}(\varphi_k, Y_k)$.

Conversely, suppose $x \in \bigoplus_{k=1}^m \text{WC}(\varphi_k, Y_k)$, and again fix a sequence $\{t_n\}_{n \in \mathbb{N}}$ in \mathbb{R} so that $\lim_{n \rightarrow \infty} |t_n| = +\infty$. Then for each $k \leq m$ there exists a sequence $\{x_{n,k}\}_{n \in \mathbb{N}}$ in Y_k converging to x_k such that the induced sequences $\{\varphi_k(t_n, x_{n,k})\}_{n \in \mathbb{N}}$ are bounded. These sequences generate a sequence $\{x_n\}_{n \in \mathbb{N}}$ in Y converging to x such that the induced sequence $\{\varphi(t_n, x_n)\}_{n \in \mathbb{N}}$ is bounded. Since $\{t_n\}_{n \in \mathbb{N}}$ was arbitrary, $x \in \text{WC}(\varphi, \bigoplus_{k=1}^m Y_k)$.

Since x was arbitrary, $\text{WC}(\varphi, \bigoplus_{k=1}^m Y_k) = \bigoplus_{k=1}^m \text{WC}(\varphi_k, Y_k)$, and as Y was also arbitrary this is true for any collection $\{Y_k\}_{k=1}^m$ of subspaces of X with $Y_k \subseteq X_k$ for all $k \leq m$. The proof of (ii) is similar.

Recall Definition 4.4, in which n is written in the form $\delta_l 2^{l-1} + \dots + \delta_1 2^0$, where the δ_j are the digits of the binary representation of n . Every value of n has a corresponding value for l (though several n may share the same l value) so (iii) may be proved by induction on l .

In the case where $l = 1$, then $\tilde{\delta}_1(\varphi, X)$ is either $\text{WC}(\varphi, X)$ or $\text{SC}(\varphi, X)$, so as a consequence of (i) and (ii) $\tilde{\delta}_1(\varphi, X) = \bigoplus_{k=1}^m \tilde{\delta}_1(\varphi_k, X_k)$. Now suppose that, for some fixed L , the equality $\tilde{\delta}_L(\varphi, \tilde{\delta}_{L-1}(\varphi, \dots \tilde{\delta}_2(\varphi, \tilde{\delta}_1(\varphi, X)))) = \bigoplus_{k=1}^m \tilde{\delta}_L(\varphi_k, \tilde{\delta}_{L-1}(\varphi_k, \dots \tilde{\delta}_2(\varphi_k, \tilde{\delta}_1(\varphi_k, X_k))))$ holds for any combination of $\tilde{\delta}_j$, and fix some combination $\tilde{\delta}_{L+1}(\varphi, \tilde{\delta}_L(\varphi, \dots \tilde{\delta}_2(\varphi, \tilde{\delta}_1(\varphi, X))))$ of $\tilde{\delta}_j$. Now $\tilde{\delta}_{L+1}(\varphi, \tilde{\delta}_L(\varphi, \dots \tilde{\delta}_2(\varphi, \tilde{\delta}_1(\varphi, X)))) = \tilde{\delta}_{L+1}(\varphi, \bigoplus_{k=1}^m \tilde{\delta}_L(\varphi_k, \dots \tilde{\delta}_2(\varphi_k, \tilde{\delta}_1(\varphi_k, X_k))))$ by assumption. But then $\tilde{\delta}_{L+1}(\varphi, \dots \tilde{\delta}_2(\varphi, \tilde{\delta}_1(\varphi, X))) = \bigoplus_{k=1}^m \tilde{\delta}_{L+1}(\varphi_k, \dots \tilde{\delta}_2(\varphi_k, \tilde{\delta}_1(\varphi_k, X_k)))$ as a consequence of (i) and (ii) since $\tilde{\delta}_L(\varphi_k, \dots \tilde{\delta}_2(\varphi_k, \tilde{\delta}_1(\varphi_k, X_k))) \subseteq X_k$ for all $k \leq m$. This argument works for any combination of $\tilde{\delta}_j$ with $j \leq L + 1$. It follows by induction over l that $\text{IC}_n(\varphi) = \bigoplus_{k=1}^m \text{IC}_n(\varphi_k)$.

(iv) is similar to (i) and (ii). Fix $t \in \mathbb{R}^+$, and suppose $x \in \ker(\varphi_t - I)$. Then $\varphi(t, x) = x$, so that $\varphi_i(t, x_k) = x_k$ for all $k \leq m$, and thus $x_k \in \ker((\varphi_k)_t - I)$ for all $k \leq m$. The reverse situation is similar, and (iv) follows as t was arbitrary.

(v) is a straightforward consequence of (iii) and (iv), as it is straightforward to verify that $(\bigoplus_{k=1}^m \text{IC}_n(\varphi_k)) \cap (\bigoplus_{k=1}^m \ker((\varphi_k)_t - I)) = \bigoplus_{k=1}^m (\text{IC}_n(\varphi_k) \cap \ker((\varphi_k)_t - I))$ for all $n \in \mathbb{N}_0$ and $t \in \mathbb{R}^+$. \square

The previous proposition is immensely useful, as determining the properties of these constructions on an arbitrary linear flow can be reduced to the properties of these constructions on its irreducible component flows. The next proposition demonstrates the first of the two required properties of the family $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}, t \in \mathbb{R}^+}$ in a similar fashion to the previous proposition.

Proposition 4.8 *Let h be a homeomorphic equivalence between two linear flows φ and ψ on X . The following all hold:*

- (i) $h(\text{WC}(\varphi, Y)) = \text{WC}(\psi, h(Y))$ for any subspace Y of X ;
- (ii) $h(\text{SC}(\varphi, Y)) = \text{SC}(\psi, h(Y))$ for any subspace Y of X ;
- (iii) $h(\text{IC}_n(\varphi)) = \text{IC}_n(\psi)$ for all $n \in \mathbb{N}_0$;
- (iv) $h(\ker(\varphi_t - I)) = \ker(\psi_t - I)$ for all $t \in \mathbb{R}^+$;
- (v) $h(\text{BC}_{n,t}(\varphi)) = \text{BC}_{n,t}(\psi)$ for all $n \in \mathbb{N}_0$ and $t \in \mathbb{R}^+$.

Proof. Suppose that $x \in h(\text{WC}(\varphi, Y))$. Then there exists a $y \in \text{WC}(\varphi, Y)$ satisfying $x = h(y)$. Fix a sequence $\{t_n\}_{n \in \mathbb{N}}$ in \mathbb{R} with $\lim_{n \rightarrow \infty} |t_n| = +\infty$. Then there exists a sequence $\{y_n\}_{n \in \mathbb{N}}$ in Y converging to y such that $\{\varphi(t_n, y_n)\}_{n \in \mathbb{N}}$ is bounded. Construct a sequence $\{x_n\}_{n \in \mathbb{N}}$ in $h(Y)$ by taking $x_n = h(y_n)$. Continuity of h guarantees that $\{x_n\}_{n \in \mathbb{N}}$ converges to x . Now since $\{\varphi(t_n, y_n)\}_{n \in \mathbb{N}}$ is bounded, it is entirely contained in some compact ball $\overline{B_r(0)}$ for some $r \in \mathbb{R}^+$. It follows again from the continuity of h that $\{h(\varphi(t_n, y_n))\}_{n \in \mathbb{N}}$ is contained in the compact set $h(\overline{B_r(0)})$. Thus $\{h(\varphi(t_n, y_n))\}_{n \in \mathbb{N}}$ is bounded. But $h(\varphi(t_n, y_n)) = \psi(t_n, h(y_n)) = \psi(t_n, x_n)$ for all $n \in \mathbb{N}$, so $\{x_n\}_{n \in \mathbb{N}}$ is a sequence in $h(Y)$ converging to x with the property that the induced sequence $\{\psi(t_n, x_n)\}_{n \in \mathbb{N}}$ is bounded. Since $\{t_n\}_{n \in \mathbb{N}}$ was arbitrary, $x \in \text{WC}(\psi, h(Y))$. The converse argument is similar, and thus (i) holds.

(ii) proceeds as in (i) except that in this case $\{\varphi(t_n, y_n)\}_{n \in \mathbb{N}}$ converges to zero. Defining $\{x_n\}_{n \in \mathbb{N}}$ in $h(Y)$ as in the proof of (i), it follows that $\psi(t_n, x_n) = \psi(t_n, h(y_n)) = h(\varphi(t_n, y_n))$, and thus $\{\psi(t_n, x_n)\}_{n \in \mathbb{N}}$ converges to zero as well. The rest of (ii) follows as in (i).

(iii) is proved using induction as in Proposition 4.7 - (iii). It follows from (i) and (ii) that $h(\tilde{\delta}_1(\varphi, X)) = \tilde{\delta}_1(\psi, X)$. Suppose that, for any combination of δ_j of length L , $h(\tilde{\delta}_L(\varphi, \tilde{\delta}_{L-1}(\varphi, \dots, \tilde{\delta}_2(\varphi, \tilde{\delta}_1(\varphi, X)))) = \tilde{\delta}_L(\psi, \tilde{\delta}_{L-1}(\psi, \dots, \tilde{\delta}_2(\psi, \tilde{\delta}_1(\psi, X))))$. Then by (i) and (ii) $h(\tilde{\delta}_{L+1}(\varphi, \tilde{\delta}_L(\varphi, \dots, \tilde{\delta}_2(\varphi, \tilde{\delta}_1(\varphi, X)))) = \tilde{\delta}_{L+1}(\psi, h(\tilde{\delta}_L(\varphi, \dots, \tilde{\delta}_2(\varphi, \tilde{\delta}_1(\varphi, X))))$ for any combination of $\tilde{\delta}_j$ of length $L + 1$. But then it follows from the previous assumption that $h(\tilde{\delta}_{L+1}(\varphi, \tilde{\delta}_L(\varphi, \dots, \tilde{\delta}_2(\varphi, \tilde{\delta}_1(\varphi, X)))) = \tilde{\delta}_{L+1}(\psi, \tilde{\delta}_L(\psi, \dots, \tilde{\delta}_2(\psi, \tilde{\delta}_1(\psi, X))))$. (iii) then follows by induction over l .

Fix $t \in \mathbb{R}^+$, and suppose $x \in h(\ker(\varphi_t - I))$. Then $x = h(y)$ for some $y \in \ker(\varphi_t - I)$. But then $\psi_t(x) = \psi_t(h(y)) = h(\varphi_t(y)) = h(y) = x$, so $x \in \ker(\psi_t - I)$. The converse argument is similar, proving (iv), as t was arbitrary.

Since h is, in particular, a bijection, $h(U \cap V) = h(U) \cap h(V)$ for all sets U, V in X . (v) then follows from (iii) and (iv). \square

Demonstrating the other desired property of $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$ concerning the dimensions of the various $\text{BC}_{n,t}(\varphi)$ requires more effort. First the structure of strong and weak centres of certain subspaces will be investigated. This investigation will then be extended to iterated centres and finally to $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$, culminating in Theorem 4.12.

4.2 Strong and Weak Centres

Understanding strong and weak centres is the first step toward confirming the final property of $\text{BC}_{n,t}(\varphi)$. To this end, several assumptions will be made throughout the next two sections. In light of Proposition 4.7 it will be assumed that any given flow φ is irreducible. Furthermore, it will be assumed that φ is central — that is, it will be assumed that the flow's eigenvalue lies on the imaginary axis — unless explicitly stated otherwise. The reasoning behind this assumption will become clear in the next section. Perhaps most importantly, a basis $\{b_1, \dots, b_d\}$ for X will be fixed, and it will be assumed φ has matrix form $e^{t\lambda} e^{tJ_d(0)x}$ with respect to that basis. Written explicitly, φ_t will be of the form

$$\varphi_t = e^{t\lambda} \begin{bmatrix} 1 & t & \frac{1}{2}t^2 & \frac{1}{(d-2)!}t^{d-2} & \frac{1}{(d-1)!}t^{d-1} \\ 0 & 1 & t & \frac{1}{(d-3)!}t^{d-3} & \frac{1}{(d-2)!}t^{d-2} \\ 0 & 0 & 1 & \frac{1}{(d-4)!}t^{d-4} & \frac{1}{(d-3)!}t^{d-3} \\ & & & \ddots & \\ 0 & 0 & 0 & 1 & t \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

with respect to that basis. This assumption is justified by Proposition 2.41.

To simplify notation, a chain of subspaces $K_0 \subseteq K_1 \subseteq \dots \subseteq K_d$ is defined by setting K_0 to be the trivial subspace $\{0\}$, and by setting $K_j = \text{span}\{b_1, \dots, b_j\}$ for all $0 < j \leq d$. In this section it will be shown that $\text{SC}(\varphi, K_j) = K_{\lfloor \frac{j}{2} \rfloor}$ and $\text{WC}(\varphi, K_j) = K_{\lceil \frac{j}{2} \rceil}$ for all $j \leq d$. Recall that, for any $r \in \mathbb{R}$, $\lfloor r \rfloor$ denotes the largest integer not larger than r , while $\lceil r \rceil$ denotes the smallest integer not smaller than r . Noting that $K_d = X$, the various K_j for $j \leq d$ will be sufficient to fully describe $\text{IC}_n(\varphi)$ for all $n \in \mathbb{N}_0$ in the next section. Proving that strong and weak centres behave in this fashion is long but straightforward. Note that this proof uses t -rescalings $T_d(t) := \text{diag}(1, t, t^2, \dots, t^{d-1})$ as discussed in Section 2.5.

Theorem 4.9 *Let φ be an irreducible central linear flow on X with $\mathbb{K} = \mathbb{C}$. Then the following both hold for all $j \leq d$:*

- (i) $\text{SC}(\varphi, K_j) = K_{\lfloor \frac{j}{2} \rfloor}$;
- (ii) $\text{WC}(\varphi, K_j) = K_{\lceil \frac{j}{2} \rceil}$.

Proof. First note that $\text{SC}(\varphi, K_j) = \text{SC}(\varphi|_{K_j}, K_j)$ and $\text{WC}(\varphi, K_j) = \text{WC}(\varphi|_{K_j}, K_j)$ for all $j \leq d$. Consequently, it may be assumed without loss of generality that $j = d$. Note also that $K_d = X$, and that the result is trivial if $d = 0$.

Fix a basis for X so that φ is in the matrix form described above. The proof of this theorem is split over a number of cases. The first case considered is when d is even, so there exists some $l \in \mathbb{N}$ satisfying $d = 2l$. Now φ_t may be written as

$$\varphi_t = e^{t\lambda} \left[\begin{array}{c|c} T_l^{-1}(t)AT_l(t) & t^l T_l^{-1}(t)BT_l(t) \\ \hline O & T_l^{-1}(t)AT_l(t) \end{array} \right],$$

where λ is the eigenvalue of φ and A and B are real $l \times l$ matrices of the form

$$\left[\begin{array}{ccc|cc} 1 & 1 & \frac{1}{2} & \frac{1}{(l-2)!} & \frac{1}{(l-1)!} \\ 0 & 1 & 1 & \frac{1}{(l-3)!} & \frac{1}{(l-2)!} \\ 0 & 0 & 1 & \frac{1}{(l-4)!} & \frac{1}{(l-3)!} \\ & & \ddots & & \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{array} \right] \text{ and } \left[\begin{array}{ccccc} \frac{1}{l!} & \frac{1}{(l+1)!} & \frac{1}{(l+2)!} & \frac{1}{(d-2)!} & \frac{1}{(d-1)!} \\ \frac{1}{(l-1)!} & \frac{1}{l!} & \frac{1}{(l+1)!} & \frac{1}{(d-3)!} & \frac{1}{(d-2)!} \\ \frac{1}{(l-2)!} & \frac{1}{(l-1)!} & \frac{1}{l!} & \frac{1}{(d-4)!} & \frac{1}{(d-3)!} \\ & & \ddots & & \\ \frac{1}{2} & \frac{1}{6} & \frac{1}{24} & \frac{1}{l!} & \frac{1}{(l+1)!} \\ 1 & \frac{1}{2} & \frac{1}{6} & \frac{1}{(l-1)!} & \frac{1}{l!} \end{array} \right]$$

respectively. Note that A and B are independent of t — that is, they are both constant matrices.

First it will be shown that $\text{WC}(\varphi, X) \subseteq K_l$. Suppose $x \in \text{WC}(\varphi, X)$. With respect to the chosen basis x is of the form (y_1, y_2) for some $y_1, y_2 \in \mathbb{C}^l$. Fix a sequence $\{t_n\}_{n \in \mathbb{N}}$ in \mathbb{R} such that $\lim_{n \rightarrow \infty} |t_n| = +\infty$. It may be assumed without loss of generality the $t_n \neq 0$ for all $n \in \mathbb{N}$. As $x \in \text{WC}(\varphi, X)$ there exist sequences $\{y_{1,n}\}_{n \in \mathbb{N}}$ and $\{y_{2,n}\}_{n \in \mathbb{N}}$ in \mathbb{C}^l converging to y_1 and y_2 respectively such that $\{e^{t_n \lambda} (T_l^{-1}(t_n)AT_l(t_n)y_{1,n} + t_n^l T_l^{-1}(t_n)BT_l(t_n)y_{2,n})\}_{n \in \mathbb{N}}$ is bounded. It follows immediately that $\{T_l^{-1}(t_n)AT_l(t_n)y_{1,n} + t_n^l T_l^{-1}(t_n)BT_l(t_n)y_{2,n}\}_{n \in \mathbb{N}}$ is bounded, as $|e^{t_n \lambda}| = 1$ for all $n \in \mathbb{N}$ since $\lambda \in i\mathbb{R}$. Now $\{t_n^{-l} T_l(t_n)\}_{n \in \mathbb{N}}$ clearly converges to zero, so it follows that $\{t_n^{-l} AT_l(t_n)y_{1,n} + BT_l(t_n)y_{2,n}\}_{n \in \mathbb{N}}$ converges to zero. Consider $\{t_n^{-l} AT_l(t_n)y_{1,n}\}_{n \in \mathbb{N}}$. As the sequence $\{y_{1,n}\}_{n \in \mathbb{N}}$ converges to y_1 , it is in particular bounded. As $\{t_n^{-l} T_l(t_n)\}_{n \in \mathbb{N}}$ converges to zero, and as A is independent of n , $\{t_n^{-l} AT_l(t_n)y_{1,n}\}_{n \in \mathbb{N}}$ converges to zero, so $\{BT_l(t_n)y_{2,n}\}_{n \in \mathbb{N}}$ must also converge to zero. Now B is a matrix of the form described by Proposition 2.45, so B is invertible, and thus $\{T_l(t_n)y_{2,n}\}_{n \in \mathbb{N}}$

converges to zero. As $\{T_l^{-1}(t_n)\}_{n \in \mathbb{N}}$ is clearly bounded (in fact, convergent) it must be that $\{y_{2,n}\}_{n \in \mathbb{N}}$ converges to zero. But by construction $\{y_{2,n}\}_{n \in \mathbb{N}}$ converges to y_2 , so by uniqueness of limits $y_2 = 0$, and consequently $x \in K_l$. Then $\text{WC}(\varphi, X) \subseteq K_l$ as x was arbitrary.

Next it will be shown that $K_l \subseteq \text{SC}(\varphi, X)$. Suppose $x \in K_l$. Then x is of the form $(y, 0)$ for some $y \in \mathbb{C}^l$ with respect to the chosen basis. Fix a sequence $\{t_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} |t_n| = +\infty$, and define a sequence $\{x_n\}_{n \in \mathbb{N}}$ in X given by $x_n = (y, 0)$ for $n < N$ and $x_n = (y, -t_n^{-l} T_l^{-1}(t_n) B^{-1} A T_l(t_n) y)$ otherwise with respect to the chosen basis, where $N \in \mathbb{N}$ is chosen sufficiently large that $t_n \neq 0$ for all $n \geq N$. Note that B^{-1} exists as was seen previously. Also note that, for all $n \geq N$, each entry of the $l \times l$ matrix $t_n^{-l} T_l^{-1}(t_n) B^{-1} A T_l(t_n)$ is the product of some constant real number (independent of n) with t_n^{-j} for some $j \in \mathbb{N}$. Consequently $\{-t_n^{-l} T_l^{-1}(t_n) B^{-1} A T_l(t_n) y\}_{n \geq N}$ converges to zero, and thus $\{x_n\}_{n \in \mathbb{N}}$ converges to $(y, 0) = x$. But by construction,

$$\varphi(t_n, x_n) = \varphi_{t_n} \left[\frac{y}{-t_n^{-l} T_l^{-1}(t_n) B^{-1} A T_l(t_n) y} \right] = \left[\frac{0}{-e^{t_n \lambda} t_n^{-l} T_l^{-1}(t_n) A B^{-1} A T_l(t_n) y} \right]$$

for all $n \geq N$. It follows similarly to above that $\{-t_n^{-l} T_l^{-1}(t_n) A B^{-1} A T_l(t_n) y\}_{n \geq N}$ converges to zero. Now $|e^{t_n \lambda}| = 1$ for all $n \in \mathbb{N}$ since $\lambda \in i\mathbb{R}$, and consequently the sequence $\{-e^{t_n \lambda} t_n^{-l} T_l^{-1}(t_n) A B^{-1} A T_l(t_n) y\}_{n \geq N}$ converges to zero. It follows that $\{\varphi(t_n, x_n)\}_{n \in \mathbb{N}}$ converges to zero, and thus $x \in \text{SC}(\varphi, X)$ as $\{t_n\}_{n \in \mathbb{N}}$ was arbitrary. Then $K_l \subseteq \text{SC}(\varphi, X)$ as x was arbitrary.

By the preceding arguments, $\text{WC}(\varphi, X) \subseteq K_l$ and $K_l \subseteq \text{SC}(\varphi, X)$. But necessarily $\text{SC}(\varphi, X) \subseteq \text{WC}(\varphi, X)$, so $\text{WC}(\varphi, X) = \text{SC}(\varphi, X) = K_l$. As $l = \lfloor \frac{d}{2} \rfloor = \lceil \frac{d}{2} \rceil$, this completes the proof of the theorem for even d .

Now consider the case where d is odd, so $d = 2l + 1$ for some $l \in \mathbb{N}_0$. It is easily seen that the theorem holds for $d = 1$. In this case, φ is of the form $e^{tir} x$ for some $r \in \mathbb{R}$. But then $\|\varphi(t, x)\| = \|x\|$. Consequently, given any fixed x in X , any sequence $\{t_n\}_{n \in \mathbb{N}}$ in \mathbb{R} such that $\lim_{n \rightarrow \infty} |t_n| = +\infty$, and any sequence $\{x_n\}_{n \in \mathbb{N}}$ in X converging to x , it must be that the sequence $\{\varphi(t_n, x_n)\}_{n \in \mathbb{N}}$ is bounded but can only converge to zero if $x = 0$. Suppose now that $d \geq 3$ so that $l \geq 1$. This case is similar to the even case but somewhat more complicated, and it is now necessary to handle the strong and weak centres separately.

First it will be shown that $\text{WC}(\varphi, X) = K_{\lceil \frac{d}{2} \rceil} = K_{l+1}$. Similarly to the even case, φ_t may be written as

$$\varphi_t = e^{t\lambda} \left[\begin{array}{c|c} T_l^{-1}(t)\tilde{A}T_{l+1}(t) & t^{l+1}T_l^{-1}(t)\tilde{B}T_l(t) \\ \hline \tilde{O} & tT_{l+1}^{-1}(t)\tilde{C}T_l(t) \end{array} \right],$$

where λ is the eigenvalue of φ and \tilde{A} , \tilde{B} , \tilde{O} , and \tilde{C} are real $l \times (l+1)$, $l \times l$, $(l+1) \times (l+1)$, and $(l+1) \times l$ matrices respectively satisfying

$$\left[\begin{array}{c|c} \tilde{A} & \tilde{B} \\ \hline \tilde{O} & \tilde{C} \end{array} \right] = \left[\begin{array}{cc|cccc} 1 & 1 & \frac{1}{2} & \frac{1}{(l-1)!} & \frac{1}{l!} & \frac{1}{(l+1)!} & \frac{1}{(l+2)!} & \frac{1}{(l+3)!} & \frac{1}{(d-2)!} & \frac{1}{(d-1)!} \\ 0 & 1 & 1 & \frac{1}{(l-2)!} & \frac{1}{(l-1)!} & \frac{1}{l!} & \frac{1}{(l+1)!} & \frac{1}{(l+2)!} & \frac{1}{(d-3)!} & \frac{1}{(d-2)!} \\ 0 & 0 & 1 & \frac{1}{(l-3)!} & \frac{1}{(l-2)!} & \frac{1}{(l-1)!} & \frac{1}{l!} & \frac{1}{(l+1)!} & \frac{1}{(d-4)!} & \frac{1}{(d-3)!} \\ & & \ddots & & & & & & & \\ 0 & 0 & 0 & 1 & \frac{1}{2} & \frac{1}{6} & \frac{1}{24} & \frac{1}{120} & \frac{1}{(l+1)!} & \frac{1}{(l+2)!} \\ 0 & 0 & 0 & 1 & 1 & \frac{1}{2} & \frac{1}{6} & \frac{1}{24} & \frac{1}{l!} & \frac{1}{(l+1)!} \\ \hline 0 & 0 & 0 & 0 & 1 & 1 & \frac{1}{2} & \frac{1}{6} & \frac{1}{(l-1)!} & \frac{1}{l!} \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & \frac{1}{2} & \frac{1}{(l-2)!} & \frac{1}{(l-1)!} \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & \frac{1}{(l-3)!} & \frac{1}{(l-2)!} \\ & & \ddots & & & & & & & \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array} \right].$$

First it will be shown that $\text{WC}(\varphi, X) \subseteq K_{l+1}$. Suppose $x \in \text{WC}(\varphi, X)$. With respect to the chosen basis x is of the form (y_1, y_2) for some $y_1 \in \mathbb{C}^{l+1}$ and $y_2 \in \mathbb{C}^l$. Fix a sequence $\{t_n\}_{n \in \mathbb{N}}$ in \mathbb{R} such that $\lim_{n \rightarrow \infty} |t_n| = +\infty$. It may be assumed without loss of generality the $t_n \neq 0$ for all $n \in \mathbb{N}$. As $x \in \text{WC}(\varphi, X)$ there exist sequences $\{y_{1,n}\}_{n \in \mathbb{N}}$ in \mathbb{C}^{l+1} and $\{y_{2,n}\}_{n \in \mathbb{N}}$ in \mathbb{C}^l converging to y_1 and y_2 respectively such that the induced sequence $\{e^{t_n \lambda} (T_l^{-1}(t_n)\tilde{A}T_{l+1}(t_n)y_{1,n} + t_n^{l+1}T_l^{-1}(t_n)\tilde{B}T_l(t_n)y_{2,n})\}_{n \in \mathbb{N}}$ is bounded. It follows immediately that $\{T_l^{-1}(t_n)\tilde{A}T_{l+1}(t_n)y_{1,n} + t_n^{l+1}T_l^{-1}(t_n)\tilde{B}T_l(t_n)y_{2,n}\}_{n \in \mathbb{N}}$ is bounded, as $|e^{t_n \lambda}| = 1$ for all $n \in \mathbb{N}$ since $\lambda \in i\mathbb{R}$. Now $\{t_n^{-(l+1)}T_l(t_n)\}_{n \in \mathbb{N}}$ clearly converges to zero, so it follows that $\{t_n^{-(l+1)}\tilde{A}T_{l+1}(t_n)y_{1,n} + \tilde{B}T_l(t_n)y_{2,n}\}_{n \in \mathbb{N}}$ converges to zero. Consider $\{t_n^{-(l+1)}\tilde{A}T_{l+1}(t_n)y_{1,n}\}_{n \in \mathbb{N}}$. As the sequence $\{y_{1,n}\}_{n \in \mathbb{N}}$ converges to y_1 , it is in particular bounded. As $\{t_n^{-(l+1)}T_{l+1}(t_n)\}_{n \in \mathbb{N}}$ converges to zero, and as \tilde{A} is independent of n , it must be that $\{t_n^{-(l+1)}\tilde{A}T_{l+1}(t_n)y_{1,n}\}_{n \in \mathbb{N}}$ converges to zero, so $\{\tilde{B}T_l(t_n)y_{2,n}\}_{n \in \mathbb{N}}$ must also converge

to zero. Now \tilde{B} is a matrix of the form described by Proposition 2.45, so \tilde{B} is invertible, and thus $\{T_l(t_n)y_{2,n}\}_{n \in \mathbb{N}}$ converges to zero. As $\{T_l^{-1}(t_n)\}_{n \in \mathbb{N}}$ is clearly bounded, it must be that $\{y_{2,n}\}_{n \in \mathbb{N}}$ converges to zero. But by construction $\{y_{2,n}\}_{n \in \mathbb{N}}$ converges to y_2 , so by uniqueness of limits $y_2 = 0$, and consequently $x \in K_{l+1}$. Then $\text{WC}(\varphi, X) \subseteq K_{l+1}$ as x was arbitrary.

Next it will be shown that $K_{l+1} \subseteq \text{WC}(\varphi, X)$. Suppose $x \in K_{l+1}$. Then x is of the form $(y, 0)$ for some $y \in \mathbb{C}^{l+1}$ with respect to the chosen basis. Fix a sequence $\{t_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} |t_n| = +\infty$, and define a sequence $\{x_n\}_{n \in \mathbb{N}}$ in X given by $x_n = (y, 0)$ for $n < N$ and $x_n = (y, -t_n^{-(l+1)}T_l^{-1}(t_n)\tilde{B}^{-1}\tilde{A}T_{l+1}(t_n)y)$ otherwise with respect to the chosen basis, where $N \in \mathbb{N}$ is chosen sufficiently large that $t_n \neq 0$ for all $n \geq N$. Note that \tilde{B}^{-1} exists as was seen previously. Also note that, for all $n \geq N$, each entry of the $l \times (l+1)$ matrix $t_n^{-(l+1)}T_l^{-1}(t_n)\tilde{B}^{-1}\tilde{A}T_{l+1}(t_n)$ is the product of some constant real number (independent of n) with t_n^{-j} for some $j \in \mathbb{N}$. Consequently $\{-t_n^{-(l+1)}T_l^{-1}(t_n)\tilde{B}^{-1}\tilde{A}T_{l+1}(t_n)y\}_{n \geq N}$ converges to zero, and thus $\{x_n\}_{n \in \mathbb{N}}$ converges to $(y, 0) = x$. But similarly to the previous case

$$\varphi(t_n, x_n) = \left[\frac{0}{e^{t_n \lambda}(\tilde{O}y - t_n^{-l}T_{l+1}^{-1}(t_n)\tilde{C}\tilde{B}^{-1}\tilde{A}T_{l+1}(t_n)y)} \right]$$

by construction for all $n \geq N$. While $\{-t_n^{-l}T_{l+1}^{-1}(t_n)\tilde{C}\tilde{B}^{-1}\tilde{A}T_{l+1}(t_n)y\}_{n \geq N}$ does not converge to zero, it does at least converge as the highest power of t_n that appears is t_n^0 , so it is in particular bounded. As $\tilde{O}y$ is constant, $\{\tilde{O}y - t_n^{-l}T_{l+1}^{-1}(t_n)\tilde{C}\tilde{B}^{-1}\tilde{A}T_{l+1}(t_n)y\}_{n \geq N}$ is bounded, and so is $\{e^{t_n \lambda}(\tilde{O}y - t_n^{-l}T_{l+1}^{-1}(t_n)\tilde{C}\tilde{B}^{-1}\tilde{A}T_{l+1}(t_n)y)\}_{n \geq N}$. It follows that the sequence $\{\varphi(t_n, x_n)\}_{n \in \mathbb{N}}$ is bounded, and thus $x \in \text{WC}(\varphi, X)$ as $\{t_n\}_{n \in \mathbb{N}}$ was arbitrary. Then $K_{l+1} \subseteq \text{WC}(\varphi, X)$ as x was arbitrary.

By the preceding arguments, $\text{WC}(\varphi, X) \subseteq K_{l+1}$ and $K_{l+1} \subseteq \text{WC}(\varphi, X)$. Consequently $\text{WC}(\varphi, X) = K_{l+1}$. As $l+1 = \lceil \frac{d}{2} \rceil$, this completes the proof of property (ii) for odd d .

Finally it will be shown that $\text{SC}(\varphi, X) = K_{\lfloor \frac{d}{2} \rfloor} = K_l$ when d is odd. As per above, assume that $d \geq 3$ so that $l \geq 1$. In this case the matrix form of φ_t with respect to the chosen basis is

$$\varphi_t = e^{t\lambda} \left[\begin{array}{c|c} T_{l+1}^{-1}(t)\hat{A}T_l(t) & t^l T_{l+1}^{-1}(t)\hat{B}T_{l+1}(t) \\ \hline O & t^{-1}T_l^{-1}(t)\hat{C}T_{l+1}(t) \end{array} \right]$$

where λ is the eigenvalue of φ and \hat{A} , \hat{B} , and \hat{C} are real $(l+1) \times l$, $(l+1) \times (l+1)$, and

$l \times (l + 1)$ matrices respectively satisfying

$$\left[\begin{array}{c|c} \hat{A} & \hat{B} \\ \hline O & \hat{C} \end{array} \right] = \left[\begin{array}{cc|cccc|cc|cc} 1 & 1 & \frac{1}{2} & \frac{1}{(l-2)!} & \frac{1}{(l-1)!} & \frac{1}{l!} & \frac{1}{(l+1)!} & \frac{1}{(l+2)!} & \frac{1}{(d-2)!} & \frac{1}{(d-1)!} \\ 0 & 1 & 1 & \frac{1}{(l-3)!} & \frac{1}{(l-2)!} & \frac{1}{(l-1)!} & \frac{1}{l!} & \frac{1}{(l+1)!} & \frac{1}{(d-3)!} & \frac{1}{(d-2)!} \\ 0 & 0 & 1 & \frac{1}{(l-4)!} & \frac{1}{(l-3)!} & \frac{1}{(l-2)!} & \frac{1}{(l-1)!} & \frac{1}{l!} & \frac{1}{(d-4)!} & \frac{1}{(d-3)!} \\ & & \ddots & & & & & & & \\ 0 & 0 & 0 & 0 & 1 & 1 & \frac{1}{2} & \frac{1}{6} & \frac{1}{l!} & \frac{1}{(l+1)!} \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & \frac{1}{2} & \frac{1}{(l-1)!} & \frac{1}{l!} \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & \frac{1}{(l-2)!} & \frac{1}{(l-1)!} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & \frac{1}{(l-3)!} & \frac{1}{(l-2)!} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{(l-4)!} & \frac{1}{(l-3)!} \\ & & \ddots & & & & & & & \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array} \right].$$

First it will be shown that $\text{SC}(\varphi, X) \subseteq K_l$. Suppose $x \in \text{SC}(\varphi, X)$. With respect to the chosen basis x is of the form (y_1, y_2) for some $y_1 \in \mathbb{C}^l$ and $y_2 \in \mathbb{C}^{l+1}$. Fix a sequence $\{t_n\}_{n \in \mathbb{N}}$ in \mathbb{R} such that $\lim_{n \rightarrow \infty} |t_n| = +\infty$. It may be assumed without loss of generality that $t_n \neq 0$ for all $n \in \mathbb{N}$. As $x \in \text{SC}(\varphi, X)$ there exist sequences $\{y_{1,n}\}_{n \in \mathbb{N}}$ in \mathbb{C}^l and $\{y_{2,n}\}_{n \in \mathbb{N}}$ in \mathbb{C}^{l+1} converging to y_1 and y_2 respectively such that the induced sequence $\{e^{t_n \lambda} (T_{l+1}^{-1}(t_n) \hat{A} T_l(t_n) y_{1,n} + t_n^l T_{l+1}^{-1}(t_n) \hat{B} T_{l+1}(t_n) y_{2,n})\}_{n \in \mathbb{N}}$ converges to zero, as does $\{T_{l+1}^{-1}(t_n) \hat{A} T_l(t_n) y_{1,n} + t_n^l T_{l+1}^{-1}(t_n) \hat{B} T_{l+1}(t_n) y_{2,n}\}_{n \in \mathbb{N}}$. Now $\{t_n^{-l} T_{l+1}(t_n)\}_{n \in \mathbb{N}}$ is clearly bounded, so it follows that $\{t_n^{-l} \hat{A} T_l(t_n) y_{1,n} + \hat{B} T_{l+1}(t_n) y_{2,n}\}_{n \in \mathbb{N}}$ converges to zero. Consider $\{t_n^{-l} \hat{A} T_l(t_n) y_{1,n}\}_{n \in \mathbb{N}}$. As the sequence $\{y_{1,n}\}_{n \in \mathbb{N}}$ converges to y_1 , it is in particular bounded. As $\{t_n^{-l} T_l(t_n)\}_{n \in \mathbb{N}}$ converges to zero, and as \hat{A} is independent of n , it must be that $\{t_n^{-l} \hat{A} T_l(t_n) y_{1,n}\}_{n \in \mathbb{N}}$ converges to zero, so $\{\hat{B} T_{l+1}(t_n) y_{2,n}\}_{n \in \mathbb{N}}$ must also converge to zero. Now \hat{B} is a matrix of the form described by Proposition 2.45, so \hat{B} is invertible, and thus $\{T_{l+1}(t_n) y_{2,n}\}_{n \in \mathbb{N}}$ converges to zero. As $\{T_{l+1}^{-1}(t_n)\}_{n \in \mathbb{N}}$ is clearly bounded, it must be that $\{y_{2,n}\}_{n \in \mathbb{N}}$ converges to zero. But by construction $\{y_{2,n}\}_{n \in \mathbb{N}}$ converges to y_2 , so by uniqueness of limits $y_2 = 0$, and consequently $x \in K_l$. Then $\text{SC}(\varphi, X) \subseteq K_l$ as x was arbitrary.

Next it will be shown that $K_l \subseteq \text{SC}(\varphi, X)$. Suppose $x \in K_l$. Then x is of the form

$(y, 0)$ for some $y \in \mathbb{C}^l$ with respect to the chosen basis. Fix a sequence $\{t_n\}_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} |t_n| = +\infty$, and define a sequence $\{x_n\}_{n \in \mathbb{N}}$ in X given by $x_n = (y, 0)$ for $n < N$ and $x_n = (y, -t_n^{-l} T_{l+1}^{-1}(t_n) \hat{B}^{-1} \hat{A} T_l(t_n) y)$ otherwise with respect to the chosen basis, where $N \in \mathbb{N}$ is chosen sufficiently large that $t_n \neq 0$ for all $n \geq N$. Note that \hat{B}^{-1} exists as was seen previously. Also note that, for all $n \geq N$, each entry of the $(l+1) \times l$ matrix $t_n^{-l} T_{l+1}^{-1}(t_n) \hat{B}^{-1} \hat{A} T_l(t_n)$ is the product of some constant real number (independent of n) with t_n^{-j} for some $j \in \mathbb{N}$. Consequently $\{-t_n^{-l} T_{l+1}^{-1}(t_n) \hat{B}^{-1} \hat{A} T_l(t_n) y\}_{n \geq N}$ converges to zero, and thus $\{x_n\}_{n \in \mathbb{N}}$ converges to $(y, 0) = x$. But by construction

$$\varphi(t_n, x_n) = \varphi_{t_n} \left[\frac{y}{-t_n^{-l} T_{l+1}^{-1}(t_n) \hat{B}^{-1} \hat{A} T_l(t_n) y} \right] = \left[\frac{0}{-e^{t_n \lambda} t_n^{-(l+1)} T_l^{-1}(t_n) \hat{C} \hat{B}^{-1} \hat{A} T_l(t_n) y} \right]$$

for all $n \geq N$. Similarly to a previous argument $\{-t_n^{-(l+1)} T_l^{-1}(t_n) \hat{C} \hat{B}^{-1} \hat{A} T_l(t_n) y\}_{n \geq N}$ converges to zero, as does $\{-e^{t_n \lambda} t_n^{-(l+1)} T_l^{-1}(t_n) \hat{C} \hat{B}^{-1} \hat{A} T_l(t_n) y\}_{n \geq N}$. It follows that the sequence $\{\varphi(t_n, x_n)\}_{n \in \mathbb{N}}$ converges to zero, and thus $x \in \text{SC}(\varphi, X)$ as $\{t_n\}_{n \in \mathbb{N}}$ was arbitrary. Then $K_l \subseteq \text{SC}(\varphi, X)$ as x was arbitrary.

By the preceding arguments, $\text{SC}(\varphi, X) \subseteq K_l$ and $K_l \subseteq \text{SC}(\varphi, X)$. Consequently $\text{SC}(\varphi, X) = K_l$. As $l = \lfloor \frac{d}{2} \rfloor$, this completes the case for the strong centre with d odd. Furthermore, this completes the proof as all cases have now been considered. \square

Based on their definitions alone, one might view strong and weak centres as fairly complicated constructions. Of course, defined as they are, it is a straightforward if tedious exercise to verify that they are preserved by homeomorphic equivalence and decomposition, as per Propositions 4.8 and 4.7. But the structure of the strong and weak centres of an irreducible central linear flow on an arbitrary subspace is non-obvious from their definitions. With that said, the preceding theorem demonstrates that the structure of strong and weak centres on subspaces of the form K_j for some $j \leq d$ is very simple.

4.3 Block Counting

Iterated centres consist of various combinations of strong and weak centres. In light of Proposition 4.9, an iterated centre would start with K_d , reduce that to $K_{\lfloor \frac{d}{2} \rfloor}$ or $K_{\lceil \frac{d}{2} \rceil}$, reduce that by approximately half, and so on, at least for irreducible central linear flows.

It follows that every iterated centre $\text{IC}_n(\varphi)$ of such a flow φ is a subspace of the form K_j for some $j \in \mathbb{N}$; moreover, since every iterated centre (except the initial one) ends with a strong centre, all iterated centres $\text{IC}_n(\varphi)$ for a sufficiently high n will be trivial as $K_0 = \{0\}$. The following proposition demonstrates that this sufficiently high n is in fact d , so that all iterated centres for $n \geq d$ are trivial. As a consequence of the intersection of iterated centres and $\ker(\varphi_t - I)$, it turns out that this is all that is needed to completely describe $\text{BC}_{n,t}(\varphi)$.

Proposition 4.10 *Let φ be an irreducible central linear flow on X with $\mathbb{K} = \mathbb{C}$. Then $\text{IC}_n(\varphi) = \{0\}$ if $d \leq n$ and $\text{IC}_n(\varphi) \supseteq K_1$ if $d > n$ for all $n \in \mathbb{N}_0$.*

Proof. Since Theorem 4.9 states that the strong and weak centres of an irreducible central linear flow on subspaces of the form K_j for some $j \in \mathbb{N}$ are subspaces of the form $K_{\tilde{j}}$ for some $\tilde{j} \in \mathbb{N}$, and since iterated centres consist solely of various combinations of strong and weak centres starting with $K_d = X$, it follows that $\text{IC}_n(\varphi) = K_j$ for some $j \in \mathbb{N}$ for all $n \in \mathbb{N}$. Note that $K_0 = \{0\}$. Also note that the case where $d = 0$ is trivial.

Recall that for iterated centres n was considered in its binary form $\delta_l 2^{l-1} + \dots + \delta_1 2^0$ where the $\delta_j \in \{0, 1\}$ with $\delta_l = 1$ unless $n = 0$. As in the case of Propositions 4.7 and 4.8, this proposition is proved via induction on l .

Suppose first that $l = 1$ in which case $n = 0$ or $n = 1$. If $n = 0$, then $n < d$ for all $d \in \mathbb{N}$. This leads to the desired conclusion, since $\text{IC}_0(\varphi) = \text{WC}(\varphi, X) = K_{\lceil \frac{d}{2} \rceil}$ and $\lceil \frac{d}{2} \rceil \geq 1$ for all $d \geq 1$. Similarly, if $n = 1$ and $d > 1$, then since $\text{IC}_1(\varphi) = \text{SC}(\varphi, X) = K_{\lfloor \frac{d}{2} \rfloor}$ and $\lfloor \frac{d}{2} \rfloor \geq 1$ for all $d > 1$, it follows that $\text{IC}_1(\varphi) \supseteq K_1$. Finally, if $n = 1$ and $d = 1$, then $\text{IC}_1(\varphi) = \text{SC}(\varphi, X) = K_{\lfloor \frac{d}{2} \rfloor} = K_0$. Thus the claim is correct for $l = 1$.

Fix $L \in \mathbb{N}$ and suppose that the proposition holds for all n with $l \leq L$. Suppose further that $n \in \mathbb{N}$ satisfies $l = L + 1$. There are two main cases to be considered based on whether n is even or odd. Then within each of these main cases are a number of sub-cases based on whether d is even or odd and whether or not $n \geq d$.

Suppose first that n is even. Then $\delta_1 = 0$, so $n = 2^{l-1} + \delta_{l-1} 2^{l-2} + \dots + \delta_2 2$. It follows that $\frac{n}{2} = 2^{l-2} + \delta_{l-1} 2^{l-3} + \dots + \delta_2$ so $\frac{n}{2}$ has length $l - 1 = L$. Moreover, it holds that $\text{IC}_n(\varphi) = \text{SC}(\varphi, \tilde{\delta}_{l-1}(\varphi, \dots, \tilde{\delta}_2(\varphi, K_{\lceil \frac{d}{2} \rceil}))) = \text{IC}_{\frac{n}{2}}(\varphi|_{K_{\lceil \frac{d}{2} \rceil}})$. Thus the proposition holds for even n assuming it can be shown that $n \geq d$ if and only if $\frac{n}{2} \geq \lceil \frac{d}{2} \rceil$. To see that this indeed is the case, note that if d is even, then trivially $n \geq d$ if and only if $\frac{n}{2} \geq \frac{d}{2} = \lceil \frac{d}{2} \rceil$. If d is

odd and $n \geq d$, then necessarily $n \geq d + 1$, so $\frac{n}{2} \geq \frac{d+1}{2} = \lceil \frac{d}{2} \rceil$. On the other hand, if d is odd and $n < d$, then trivially $n < d + 1$, so $\frac{n}{2} < \frac{d+1}{2} = \lceil \frac{d}{2} \rceil$. This completes the case where n is even.

Suppose now that n is odd. Then $n - 1$ is even, and it follows similarly to above that $\frac{n-1}{2} = 2^{l-2} + \delta_{l-1}2^{l-3} + \dots + \delta_2$ and is of length $l - 1 = L$. As $\delta_1 = 1$ in this case, it follows that $\text{IC}_n(\varphi) = \text{IC}_{\frac{n-1}{2}}(\varphi|_{K_{\lfloor \frac{d}{2} \rfloor}})$. Thus the proposition holds for odd n assuming it can be shown that $n \geq d$ if and only if $\frac{n-1}{2} \geq \lfloor \frac{d}{2} \rfloor$. Now if d is also odd then trivially $n \geq d$ if and only if $\frac{n-1}{2} \geq \frac{d-1}{2} = \lfloor \frac{d}{2} \rfloor$. Otherwise if d is even and $n \geq d$, then necessarily $n - 1 \geq d$, so $\frac{n-1}{2} \geq \frac{d}{2} = \lfloor \frac{d}{2} \rfloor$. Finally, if d is even and $n < d$, then trivially $n - 1 < d$, so $\frac{n-1}{2} < \frac{d}{2} = \lfloor \frac{d}{2} \rfloor$. This completes the case when n is odd, and the proposition follows by induction on l . \square

The other major ingredients in the definition of the family $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}, t \in \mathbb{R}^+}$ are the subspaces $\ker(\varphi_t - I)$ for $t \in \mathbb{R}^+$. These spaces are much simpler than strong, weak, and iterated centres. They are often trivial; in particular, they are trivial for any hyperbolic linear flow. It is for this reason that strong, weak, and iterated centres have only been considered for (irreducible) central flows — the iterated centres of hyperbolic flows will always reduce to $\{0\}$ upon intersection with this kernel.

Proposition 4.11 *Let φ be an irreducible central linear flow on X with $\mathbb{K} = \mathbb{C}$ and eigenvalue λ . Then $\ker(\varphi_t - I) = \{0\}$ except when either $\lambda = 0$ or $\lambda = \frac{z2\pi i}{t}$ for some nonzero $z \in \mathbb{Z}$, in which case $\ker(\varphi_t - I) = K_1$.*

Proof. Fix a basis for X so that φ is of the form $e^{t\lambda}e^{tJ_d(0)}x$ as per Proposition 2.41. Consider that $e^{tJ_d(0)} = I + \sum_{j=1}^{d-1} \frac{1}{j!}t^j J_d(0)^j$. Now $\sum_{j=1}^{d-1} \frac{1}{j!}t^j J_d(0)^j$ is an upper diagonal matrix with zeros along the diagonal, and this remains true if the matrix is multiplied by $e^{t\lambda}$. Thus, the diagonal of $e^{t\lambda}e^{tJ_d(0)} - I$ is the same as the diagonal of $e^{t\lambda}I - I$, and so $\ker(\varphi_t - I) = \{0\}$ whenever $e^{t\lambda} \neq 1$. Now if $\lambda \notin i\mathbb{R}$, then $e^{t\lambda} \neq 1$ for all $t \in \mathbb{R}^+$. Even if $\lambda \in i\mathbb{R}$, then $e^{t\lambda} \neq 1$ unless $\lambda = \frac{z2\pi i}{t}$ for some nonzero $z \in \mathbb{Z}$.

Suppose now that either $\lambda = 0$ or $\lambda = \frac{z2\pi i}{t}$ for some nonzero $z \in \mathbb{Z}$ so that $e^{t\lambda} = 1$. Since $e^{t\lambda}e^{tJ_d(0)} - I = \sum_{j=1}^{d-1} \frac{1}{j!}t^j J_d(0)^j$, it is clear that $K_1 \subseteq \ker(\varphi_t - I)$. Suppose in turn that $x \in \ker(\varphi_t - I)$. Then $\sum_{j=1}^{d-1} \frac{1}{j!}t^j J_d(0)^j x = 0$. Since t is nonzero and $J_d(0)^d = 0$, multiplying this equation by $J_d(0)^{d-2}$ demonstrates that $J_d(0)^{d-1}x = 0$.

Thus $\sum_{j=1}^{d-2} \frac{1}{j!} t^j J_d(0)^j x = 0$. But multiplying this equation by $J_d(0)^{d-3}$ demonstrates that $J_d(0)^{d-2} x = 0$, so $\sum_{j=1}^{d-3} \frac{1}{j!} t^j J_d(0)^j x = 0$. Continuing inductively one concludes that $J_d(0)x = 0$. As $\ker J_d(0) = K_1$, it follows that $x \in K_1$. Thus $\ker(\varphi_t - I) = K_1$. \square

What does this mean for the intersection between iterated centres and these kernels for irreducible linear flows? If the flow is hyperbolic, then the intersection is always trivial due to the kernel. If the flow is central, then the intersection is still trivial due to the kernel, except in the case where the flow's eigenvalue lines up appropriately with t . At this point the iterated centre comes into play. Even with a central flow with appropriate eigenvalue, the intersection is still trivial if $d \leq n$. Only in the case where the flow is central with an appropriate eigenvalue and $n < d$ will the intersection be nontrivial. In such a situation the kernel is K_1 and the iterated centre contains K_1 so the intersection is K_1 . This is what gives the family $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}, t \in \mathbb{R}^+}$ its final property.

Theorem 4.12 *Let φ be a linear flow on X with $\mathbb{K} = \mathbb{C}$. There exists a family of subspaces of X , denoted $\{\text{BC}_{n,t}(\varphi)\}_{n \in \mathbb{N}_0, t \in \mathbb{R}^+}$, that satisfy the following two properties for all $n \in \mathbb{N}_0$ and $t \in \mathbb{R}^+$:*

- (i) *If h is a homeomorphic equivalence, then $h(\text{BC}_{n,t}(\varphi)) = \text{BC}_{n,t}(h \circ \varphi)$;*
- (ii) *The number of irreducible components of φ of dimension greater than n with eigenvalue either 0 or $\frac{z2\pi i}{t}$ for some nonzero $z \in \mathbb{Z}$ equals $\dim \text{BC}_{n,t}(\varphi)$.*

Proof. The existence of this family of subspaces is clear from its construction. Property (i) is just Proposition 4.8 - (v). It remains to prove (ii).

Fix $n \in \mathbb{N}_0$ and $t \in \mathbb{R}^+$. Let $\bigoplus_{k=1}^m \varphi_k$ be the irreducible decomposition of φ guaranteed by Theorem 2.37. Fix a basis \mathcal{B} of X by combining in order the elements of bases \mathcal{B}_k of X_k for each $k \leq m$ constructed so that each $\varphi_k(t, x) = e^{t\lambda_k} e^{tJ_{d_k}(0)} x$ with respect to \mathcal{B}_k . This can be done by Proposition 2.41. Now by Proposition 4.7 $\text{BC}_{n,t}(\varphi) = \bigoplus_{k=1}^m \text{BC}_{n,t}(\varphi_k)$. But $\dim \text{BC}_{n,t}(\varphi) = \dim \bigoplus_{k=1}^m \text{BC}_{n,t}(\varphi_k) = \sum_{k=1}^m \dim \text{BC}_{n,t}(\varphi_k)$ from the basic properties of dimensions. Now if φ_k is of dimension $d_k > n$ with eigenvalue λ_k either zero or of the form $\frac{z2\pi i}{t}$ for some nonzero $z \in \mathbb{N}$, then $\text{BC}_{n,t}(\varphi_k) = \text{IC}_n(\varphi_k) \cap \ker(\varphi_{kt} - I) = K_1$ by Propositions 4.10 and 4.11, so that $\dim \text{BC}_{n,t}(\varphi_k) = 1$ in this case. Otherwise it follows from

these propositions that $\text{BC}_{n,t}(\varphi_k) = \{0\}$, so $\dim \text{BC}_{n,t}(\varphi_k) = 0$. Since each φ_k corresponds to an irreducible component of φ , it follows that $\dim \text{BC}_{n,t}(\varphi)$ is the number of irreducible components of φ of dimension $d_k > n$ with eigenvalue either 0 or of the form $\frac{z2\pi i}{t}$ for some nonzero $z \in \mathbb{Z}$. Since n and t were arbitrary, this completes the proof. \square

This theorem is all that is necessary to complete the characterization of homeomorphic equivalence discussed in Section 3.3.

5 Examples and Other Considerations

With the theoretical work now complete, it is natural to turn toward more practical problems such as explicitly classifying particular flows. On normed vector spaces of small dimension, it is a straightforward exercise to provide a natural representative for all equivalence classes of linear flows. It is also worthwhile to consider additional avenues of investigation beyond this thesis, and this is discussed in the final section.

5.1 Examples

The diffeomorphic classification of flows on \mathbb{K}^d is particularly simple as a consequence of Theorem 3.2. Each distinct $d \times d$ Jordan canonical form matrix with coefficients in \mathbb{K} generates a unique equivalence class with respect to diffeomorphic equivalence, since distinct Jordan canonical forms are nonsimilar, at least up to reordering the individual Jordan blocks. On the other hand, the existence of a Jordan canonical form guarantees that every $d \times d$ matrix with coefficients in \mathbb{K} is diffeomorphically equivalent to a matrix in Jordan canonical form. Consequently the distinct (up to reordering) $d \times d$ Jordan canonical form matrices with coefficients in \mathbb{K} act as a natural set of representatives for all of the equivalence classes of linear flows with respect to diffeomorphic equivalence.

The situation is more interesting when it comes to the homeomorphic classification of linear flows on \mathbb{K}^d . Since diffeomorphic equivalence implies homeomorphic equivalence, it is still worthwhile to start by considering $d \times d$ Jordan canonical form matrices with entries in \mathbb{K} as representatives of equivalence classes. However, in this case multiple distinct Jordan canonical form matrices may generate homeomorphically equivalent flows. This can be seen clearly when it comes to hyperbolic flows — in fact it will quickly become apparent that the number of equivalence classes of hyperbolic linear flows is finite.

The simplest case is of course when $d = 1$. The linear flows on \mathbb{R} are exactly the maps $(t, x) \mapsto e^{rt}x$ for some $r \in \mathbb{R}$. If r is negative, the flow is stable and thus equivalent to $e^{-t}x$. Similarly, if r is positive, then the flow is equivalent to $e^t x$. Finally, if $r = 0$, then the flow cannot be equivalent to either $e^{-t}x$ or $e^t x$, so it generates its own single element equivalence

class. It follows that the distinct equivalence classes of linear flows on \mathbb{R} are exactly

$$[-I_1], [I_1], \text{ and } [O_1]$$

where, to simplify notation, $[L]$ represents the equivalence class of the flow generated by some $L \in \mathcal{L}(X)$. The situation is similar for linear flows on \mathbb{C} . In this case the flows are exactly the maps $(t, x) \mapsto e^{ct}x$ for some $c \in \mathbb{C}$. In this case all flows with c in the left open half-plane and the right open half-plane belong to the equivalence classes $[-I_1]$ and $[I_1]$ respectively. All that remains is to deal with the situation when $c \in i\mathbb{R}$. While the two equivalence classes $[-I_1]$ and $[I_1]$ are sufficient to completely represent hyperbolic flows, an uncountable collection of equivalence classes is required to deal with the remaining nonhyperbolic flows. Each flow of the form $e^{irt}x$ is homeomorphically equivalent to $e^{-irt}x$ for all $r \in \mathbb{R}$ but is in turn not homeomorphically equivalent to $e^{i\tilde{r}t}x$ for any other $\tilde{r} \in \mathbb{R}$. Consequently each $r \in \mathbb{R}^+$ generates a distinct equivalence class, so that $\{[irI_1]\}_{r \in \mathbb{R}^+}$ is a family of distinct equivalence classes. Note that it is not necessary to consider $r \in \mathbb{R}^-$ as $[irI_1] = [\overline{irI_1}] = [-irI_1]$ for all $r \in \mathbb{R}$. It follows that the distinct equivalence classes of linear flows on \mathbb{C} are exactly

$$[-I_1], [I_1], [O_1], \text{ and the family } \{[irI_1]\}_{r \in \mathbb{R}^+}.$$

Handling the case when $d = 2$ comes with a small but nonnegligible increase in difficulty. As it is now possible for a real matrix to have a variety of purely imaginary eigenvalues, the set of equivalence classes of linear flows on \mathbb{R}^2 is no longer finite. With that said, the situation for hyperbolic flows is still straightforward. A hyperbolic flow on \mathbb{R}^2 is either completely stable, completely unstable, or a mix of stable and unstable, and the associated equivalence classes are respectively $[-I_2]$, $[I_2]$, and $[\text{diag}(-I_1, I_1)]$. The situation for non-hyperbolic flows with no zero eigenvalues is similar to the nonhyperbolic case on \mathbb{C} . In this case, by defining

$$\hat{I} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix},$$

each $r \in \mathbb{R}^+$ generates a distinct equivalence class $[r\hat{I}]$. More finesse is required to handle flows with a zero eigenvalue — two equivalence classes are required to handle flows which

only have the eigenvalue 0, along with two distinct equivalence classes to handle flows with the zero eigenvalue and a stable or unstable subspace. These equivalence classes are as follows: $[O_2]$, $[J_2(0)]$, $[\text{diag}(O_1, -I_1)]$, and $[\text{diag}(O_1, I_1)]$. Consequently, for linear flows on \mathbb{R}^2 , there are exactly seven specific equivalence classes, namely

$$[O_2], [J_2(0)], [\text{diag}(O_1, -I_1)], [\text{diag}(O_1, I_1)], [-I_2], [\text{diag}(-I_1, I_1)], \text{ and } [I_2],$$

together with the family $\{[r\hat{I}]\}_{r \in \mathbb{R}^+}$.

The first seven preceding matrices all also represent distinct equivalence classes when it comes to linear flows over \mathbb{C}^2 and the reasoning is similar. Since it is now possible for a flow to have a single purely imaginary eigenvalue (counting algebraic multiplicity), it is necessary to include some additional families of equivalences to handle flows with a zero and a purely imaginary eigenvalue, as well as to handle flows with a purely imaginary eigenvalue and a stable or unstable subspace. The required families are $\{[\text{diag}(O_1, irI)]\}_{r \in \mathbb{R}^+}$, $\{[\text{diag}(-I_1, irI)]\}_{r \in \mathbb{R}^+}$, and $\{[\text{diag}(I_1, irI)]\}_{r \in \mathbb{R}^+}$. Again note that it is only necessary to consider $r \in \mathbb{R}^+$ for the same reason as in the $d = 1$ case, and this will continue to be true throughout these examples. The situation is also more complex than in the \mathbb{R}^2 case for flows with two purely imaginary eigenvalues. It is necessary to construct a family of representatives for when a flow consists of a single Jordan Block with a purely complex eigenvalue. It is also necessary to construct a family (now ranging over two variables) to handle flows consisting of two Jordan Blocks with purely imaginary eigenvalues. These two families are $\{[J_2(ir)]\}_{r \in \mathbb{R}^+}$ and $\{[\text{diag}(ir_1 I_1, ir_2 I_1)]\}_{r_1, r_2 \in \mathbb{R}^+}$. In combination with the previously discussed three families, these five families along with the seven matrices discussed previously combine to form a complete set of distinct representatives of the equivalence classes of linear flows on \mathbb{C}^2 with respect to homeomorphic equivalence.

The explicit representation of the homeomorphic classification of linear flows on \mathbb{R}^3 is only just barely reasonable to write up fully, given how many equivalence classes and families of equivalences that are required. Now fifteen fixed matrices are necessary to represent the distinct equivalence classes for flows consisting of various combinations of stable and unstable subspaces, along with subspaces associated with the zero eigenvalue. A full list is as follows:

$$\begin{aligned}
& [O_3], [\text{diag}(O_1, J_2(0))], [J_3(0)], [\text{diag}(O_2, -I_1)], [\text{diag}(O_2, I_1)], [\text{diag}(J_2(0), -I_1)], \\
& [\text{diag}(J_2(0), I_1)], [\text{diag}(O_1, -I_2)], [\text{diag}(O_1, -I_1, I_1)], [\text{diag}(O_1, I_2)], [-I_3], [\text{diag}(-I_2, I_1)], \\
& [\text{diag}(-I_1, I_2)], \text{ and } [I_3].
\end{aligned}$$

It is also now necessary to have three separate families, namely

$$\{[\text{diag}(O_1, r\hat{I})]\}_{r \in \mathbb{R}^+}, \{[\text{diag}(-I_1, r\hat{I})]\}_{r \in \mathbb{R}^+}, \text{ and } \{[\text{diag}(I_1, r\hat{I})]\}_{r \in \mathbb{R}^+},$$

to handle the combination of a central, stable, or unstable subspace with a subspace generated by a purely imaginary eigenvalue pair respectively. This is similar to how it was necessary on \mathbb{C}^2 to have three cases to handle the combination of a purely imaginary eigenvalue with a stable and unstable space and with subspaces associated with a zero eigenvalue. The case \mathbb{C}^3 is even worse than \mathbb{R}^3 . The same fixed matrices as in the \mathbb{R}^3 case are required, along with a fairly substantial number of families. In particular, seven families consisting of a single purely imaginary eigenvalue combined with the various fixed matrices of the \mathbb{C}^2 case are required; for example, the family $\{[\text{diag}(O_2, irI_1)]\}_{r \in \mathbb{R}^+}$ is needed. Also, for both families associated with only purely imaginary eigenvalues (the single Jordan block and dual Jordan block families) in the \mathbb{C}^2 case, there are now three associated families in the \mathbb{C}^3 case. For instance $\{[\text{diag}(O_1, J_2(ir))]\}_{r \in \mathbb{R}^+}$ and $\{[\text{diag}(-I_1, ir_1I_1, ir_2I_1)]\}_{r_1, r_2 \in \mathbb{R}^+}$ are both required to completely represent the equivalence classes of linear flows on \mathbb{C}^3 . And these families are not yet sufficient; it is still necessary to handle flows on \mathbb{C}^3 with only purely imaginary eigenvalues via the families $\{[J_3(ir)]\}_{r \in \mathbb{R}^+}$, $\{[\text{diag}(ir_1I_1, J_2(ir_2))]\}_{r_1, r_2 \in \mathbb{R}^+}$, and $\{[\text{diag}(ir_1I_1, ir_2I_1, ir_3I_1)]\}_{r_1, r_2, r_3 \in \mathbb{R}^+}$.

It is perhaps becoming clear that homeomorphic classification for a given $d = D$ can be described by building upon the classifications for $d < D$. The pattern is sufficiently complex that it is not worth stating explicitly. It should be clear that explicitly expressing representatives for the equivalence classes of linear flows with d higher than 3 quickly becomes tedious.

5.2 Other Considerations

This thesis has focused on the classification of linear flows on finite-dimensional normed spaces based primarily on the specific notions of homeomorphic and diffeomorphic equiva-

lence. These notions of equivalence are fairly natural choices to work with, but one could certainly consider other notions as well. For instance, one could construct a notion of higher order derivatives. If $Dh : X \rightarrow \mathcal{L}(X)$ is the derivative of a diffeomorphic equivalence h , then there may exist a map $D^2h : X \rightarrow \mathcal{L}(X, \mathcal{L}(X))$ that is the derivative of Dh , and so on. One could then discuss $\mathcal{C}^n(X)$ equivalent flows for all $n \in \mathbb{N}_0$ where homeomorphic and diffeomorphic equivalence are simply the $n = 0$ and $n = 1$ cases respectively. One could even define a notion of smooth equivalence.

With that said, notions of equivalence based on higher order differentiability are not particularly interesting. Consider that a linear equivalence $x \mapsto hx$ is always differentiable with derivative $x \mapsto h$. This derivative is a constant map, so higher order derivatives will all exist and be the zero map from X to some appropriate space. Consequently, linear equivalence implies smooth equivalence, which in turn implies $\mathcal{C}^n(X)$ equivalence for any $n \in \mathbb{N}$, finally implying diffeomorphic equivalence in particular. It would then follow from Lemma 3.1 that all these notions of equivalence are identical to diffeomorphic equivalence.

Notions of equivalence weaker than diffeomorphic are potentially more interesting. For instance, one can examine how diffeomorphic classification morphs into homeomorphic classification by considering α -Hölder continuity for various $\alpha \in (0, 1]$. This approach is discussed in [10]. One may also consider notions of equivalence weaker than homeomorphic; for example, one could consider flow equivalences that are bimeasurable rather than bi-continuous. Even further, one could discuss basic flow equivalence, without any additional structure on the equivalence. Now less structure results in fewer tools to work with to analyze those equivalences, but there are still things that can be said. The two flows in Example 3.5 are not flow equivalent even in the most basic sense, for instance, as bicontinuity was not used in that example.

As mentioned when flow equivalence was first introduced, one may also consider equivalences for which the time variable can vary; that is, given flows φ and ψ on X one may consider bijections on X with the property that $h(\varphi(t, x)) = \psi(\tilde{h}(t, x), h(x))$ where $\tilde{h} : \mathbb{R} \times X \rightarrow \mathbb{R}$. Such constructions are sometimes referred to as conjugacies rather than equivalences. This topic is discussed more in [9], but it is not too hard to see what happens when $\tilde{h}(t, x) = rt$ for all $(t, x) \in \mathbb{R} \times X$ where r is some fixed real number. When r

is positive, diffeomorphic equivalence is made a little more flexible in that the underlying operators of two flows now need only be similar up to some uniform rescaling. In essence this introduces a single degree of freedom for diffeomorphic conjugacy that does not exist for diffeomorphic equivalence. In the latter case the underlying operators of two flows must have identical Jordan block structure once a basis is fixed. In the former case the Jordan block structure must be similar, but the eigenvalues no longer need to be exactly the same — it is now only necessary that the ratios between eigenvalues are the same. For example, consider the two flows on \mathbb{C}^2 generated by matrices L and M with eigenvalues $\{i, 2i\}$ and $\{2i, 4i\}$ respectively. These two flows clearly cannot be diffeomorphically equivalent, but they are diffeomorphically conjugate, as the Jordan canonical form of the second matrix is twice the Jordan canonical form of the first.

The situation is similar for homeomorphic conjugacy with a fixed positive r . Upon rescaling the underlying operator by r , eigenvalues with positive real part still have positive real part, eigenvalues with negative real part still have negative real part, and eigenvalues with zero real part still have zero real part. Consequently, homeomorphic conjugacy behaves identically to homeomorphic equivalence when it comes to the stable and unstable parts of a flow, and it behaves similar to diffeomorphic conjugacy when it comes to the central part of a flow, as homeomorphic equivalence functions are linearly equivalent for that part. The situation is also similar for conjugacies with negative r . Negative r values not only rescale time but also reverse the direction of paths. With that said, negative r values still only allow for one degree of freedom — either all paths are reversed or none are.

Aside from considering forms of equivalence outside those discussed in this thesis, one may also consider other types of flows. Less can be said about such situations here, as this thesis relies heavily on the innate structural properties of linear flows on finite-dimensional normed spaces. One could for example consider the case of nonlinear flows, but virtually all of this thesis is predicated on the fact that all linear flows are of the form $e^{tL}x$ for some linear operator, and this clearly is not true for nonlinear flows. One could also consider flows when the dimension is infinite. Properties of linear operators on finite-dimensional normed spaces are used both explicitly and implicitly. While it may be possible to avoid using these properties via more elegant arguments, this would be a substantial endeavour,

would likely increase the complexity of proofs, and may not be possible everywhere. One could even consider a notion of flows defined for discrete time (using say $t \in \mathbb{Z}$ rather than $t \in \mathbb{R}$), though the behaviour of discrete flows differs greatly from that of continuous flows.

Bibliography

- [1] H. Amann, *Ordinary Differential Equations: An Introduction to Nonlinear Analysis*, de Gruyter Studies in Mathematics, Walter de Gruyter & Co., Berlin · New York, 1990.
- [2] V. Arnold, *Ordinary Differential Equations*, Third Edition, Springer, Verlag · Berlin · Heidelberg · New York, 2006.
- [3] K. Conrad, Complexification, <http://www.math.uconn.edu/~kconrad/blurbs/linmultialg/complexification.pdf> (accessed December 4, 2015).
- [4] J. Dugundji, *Topology*, Allyn and Bacon Inc., Boston, 1966.
- [5] D. Dummit, R. Foote, *Abstract Algebra*, Third Edition, John Wiley and Sons Inc., 2004
- [6] J. Guckenheimer, P. Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, Springer, Verlag · New York · Berlin · Heidelberg · Tokyo, 1983.
- [7] L. Hogben *Handbook of Linear Algebra*, Second Edition, Discrete Mathematics and Its Applications, Taylor and Francis Group LLC, 2014.
- [8] N. Kuiper, The Topology of the Solutions of a Linear Differential Equation on \mathbf{R}^n , *Manifolds - Tokyo 1973: Proceedings of the International Conference on Manifolds and Related Topics in Topology*, 195–203, University of Tokyo Press, Tokyo 1973.
- [9] N. Ladis, Topological Equivalence of Linear Flows, *Differential Equations*, 9(7), 938–947, 1975.

- [10] P. McSwiggen, K. Meyer, Conjugate Phase Portraits of Linear Systems, *American Mathematical Monthly*, 115(7), 596–614, 2008.
- [11] J. Meiss, *Differential Dynamical Systems*, SIAM Series on Mathematical Modelling and Computation, Society for Industrial and Applied Mathematics, Philadelphia, 2007.
- [12] Riemann Integration, https://en.wikibooks.org/wiki/Real_Analysis/Riemann_integration (accessed December 4, 2015).

A Operator Exponential

The exponential of an operator is defined by naively applying the Taylor series expansion of the real exponential to linear operators. For this definition to make sense, it is essential that such a construction is always well-defined.

Proposition A.1 *Let $L \in \mathcal{L}(X)$. The sequence of partial sums $\{\sum_{j=0}^n \frac{1}{j!} L^j\}_{n \in \mathbb{N}}$ converges in $\mathcal{L}(X)$.*

Proof. As X is assumed to be a finite-dimensional normed space, X is in fact a Banach space, and thus it is sufficient to prove that the sequence is absolutely convergent. But $L \in \mathcal{L}(X)$ implies that $\|L\| < \infty$, so $\lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} \|L^j\| \leq e^{\|L\|} < \infty$. \square

The following definition is now justified.

Definition A.2 Let $L \in \mathcal{L}(X)$. The exponential of L is $e^L := \sum_{j=0}^{\infty} \frac{1}{j!} L^j \in \mathcal{L}(X)$.

As one might expect, many properties of exponentials of real numbers carry over to exponentials of operators. For example, $|e^t| \leq e^{|t|}$ for all $t \in \mathbb{R}$, and it is easily seen that a similar result holds for the operator exponential.

Proposition A.3 *Let $L \in \mathcal{L}(X)$. Then $\|e^L\| \leq e^{\|L\|}$.*

Proof. It follows from the triangle inequality that $\|\sum_{j=0}^n \frac{1}{j!} L^j\| \leq \sum_{j=0}^n \frac{1}{j!} \|L\|^j$ for all $n \in \mathbb{N}$. The desired result is obtained by taking n to ∞ . \square

Another basic property of exponentials of real numbers is that $e^{s+t} = e^s e^t = e^t e^s$ whenever $s, t \in \mathbb{R}$. This property also extends to exponentials of operators but with a catch. If $L, M \in \mathcal{L}(X)$ do not commute, then it is not necessarily true that $e^{L+M} = e^L e^M = e^M e^L$. This issue does not arise in the real case as commutativity is automatic.

Proposition A.4 *Let $L, M \in \mathcal{L}(X)$ commute. Then $e^{L+M} = e^L e^M = e^M e^L$.*

Proof. For notational simplicity, set $L_n = \sum_{j=0}^n \frac{1}{j!} L^j$ and $M_n = \sum_{j=0}^n \frac{1}{j!} M^j$. Now e^{L+M} is simply $\sum_{j=0}^{\infty} \frac{1}{j!} (L+M)^j$. By uniqueness of limits it is sufficient to show that for every $\epsilon \in \mathbb{R}^+$ there exists an $N \in \mathbb{N}$ so that $\|\sum_{j=0}^n \frac{1}{j!} (L+M)^j - e^L e^M\| < \epsilon$ for all $n \geq N$, and similarly for $e^M e^L$.

Fix $\epsilon \in \mathbb{R}^+$. As a consequence of the binomial formula,

$$\begin{aligned} \sum_{j=0}^n \frac{1}{j!} (L + M)^j &= \sum_{j=0}^n \sum_{k=0}^j \frac{1}{(j-k)!} L^{j-k} \frac{1}{k!} M^k \\ &= \sum_{j=0}^n \left(\sum_{k=0}^j \frac{1}{k!} L^k \right) \frac{1}{(n-j)!} M^{n-j} \\ &= \sum_{j=0}^n L_j \frac{1}{(n-j)!} M^{n-j} \end{aligned}$$

after an appropriate rearrangement and relabelling of the terms. Since, again by reordering, $M_n = \sum_{j=0}^n \frac{1}{(n-j)!} M^{n-j}$, it follows that

$$\sum_{j=0}^n \frac{1}{j!} (L + M)^j - e^L e^M = \sum_{j=0}^n (L_j - e^L) \frac{1}{(n-j)!} M^{n-j} + e^L (M_n - e^M).$$

Now there exists an $N_M \in \mathbb{N}$ such that $\|M_n - e^M\| < \frac{\epsilon}{3\|e^L\|}$ whenever $n \geq N_M$. Similarly, there exists an $N_L \in \mathbb{N}$ so that $\|L_n - e^L\| < \frac{\epsilon}{3e\|M\|}$ for all $n \geq N_L$. Finally, $\|\frac{1}{j!} M^j\|$ must converge to zero for the series $\sum_{j=0}^n \frac{1}{j!} M^j$ to be absolutely convergent, and consequently there exists an $\tilde{N}_M \in \mathbb{N}$ so that $\|\frac{1}{j!} M^j\| < \frac{\epsilon}{3(N_L+1)\max_{0 \leq l \leq N_L} \|L_l - e^L\|}$ whenever $j \geq \tilde{N}_M$.

Let $N = \max\{N_L + \tilde{N}_M, N_M\}$, and let $n \geq N$. Then for all $j \leq N_L$ it must be the case that $n - j \geq \tilde{N}_M$, so $\sum_{j=0}^{N_L} \frac{1}{(n-j)!} M^{n-j} \|L_j - e^L\| < \frac{1}{3}\epsilon$ by the construction of \tilde{N}_M . Similarly, that $\sum_{j=N_L+1}^n \frac{1}{(n-j)!} M^{n-j} \|L_j - e^L\| < \frac{1}{3}\epsilon$ follows from the construction of N_L , and that $\|e^L\| \|M_n - e^M\| < \frac{1}{3}\epsilon$ follows from the construction of N_M . Now

$$\begin{aligned} \left\| \sum_{j=0}^n \frac{1}{j!} (L + M)^j - e^L e^M \right\| &\leq \sum_{j=0}^{N_L} \frac{1}{(n-j)!} M^{n-j} \|L_j - e^L\| \\ &\quad + \sum_{j=N_L+1}^n \frac{1}{(n-j)!} M^{n-j} \|L_j - e^L\| \\ &\quad + \|e^L\| \|M_n - e^M\|, \end{aligned}$$

and consequently $\|\sum_{j=0}^n \frac{1}{j!} (L + M)^j - e^L e^M\| < \epsilon$. The desired result follows as n and ϵ were arbitrary. The $e^M e^L$ case is proved similarly. \square

This thesis is primarily interested in exponentials of the form e^{tL} for $t \in \mathbb{R}$ and $L \in \mathcal{L}(X)$. In this case the situation is more straightforward.

Corollary A.5 *Let $L \in \mathcal{L}(X)$. Then $e^{(s+t)L} = e^{sL} e^{tL}$ for all $s, t \in \mathbb{R}$.*

Proof. As sL and tL necessarily commute for all $s, t \in \mathbb{R}$, the result follows immediately from the preceding proposition. \square

Another property of real exponentials that extends to operator exponentials is differentiability — a property that plays a crucial role in this thesis.

Proposition A.6 *Let $L \in \mathcal{L}(X)$. Then e^{tL} is differentiable with respect to t , and the derivative of e^{tL} is given by Le^{tL} for each fixed $t \in \mathbb{R}$.*

Proof. Fix $\epsilon \in \mathbb{R}^+$, and let $\delta = \min\{1, \frac{\epsilon}{e^{\|L\|}\|L\| \|e^{tL}\| + 1}\} \in \mathbb{R}^+$. For all $h \in \mathbb{R} \setminus \{0\}$ such that $|h| < \delta$ it holds that

$$\begin{aligned}
\left\| \frac{1}{h}(e^{(t+h)L} - e^{tL}) - Le^{tL} \right\| &\leq \left\| \frac{1}{h}(e^{hL} - I) - L \right\| \|e^{tL}\| \\
&= \lim_{n \rightarrow \infty} \left\| \frac{1}{h} \left(\sum_{j=0}^n \frac{1}{j!} h^j L^j - I \right) - L \right\| \|e^{tL}\| \\
&= \lim_{n \rightarrow \infty} \left\| \sum_{j=1}^n \frac{1}{j!} h^{j-1} L^j - L \right\| \|e^{tL}\| \\
&\leq \lim_{n \rightarrow \infty} \left\| \sum_{j=1}^n \frac{1}{j!} h^{j-1} L^{j-1} - I \right\| \|L\| \|e^{tL}\| \\
&= \lim_{n \rightarrow \infty} \left\| \sum_{j=1}^{n-1} \frac{1}{j+1} \frac{1}{j!} h^j L^j \right\| \|L\| \|e^{tL}\| \\
&\leq \lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} |h|^j \|L\|^j \|L\| \|e^{tL}\| \\
&< \lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} \delta^j \|L\|^j \|L\| \|e^{tL}\| \\
&\leq \epsilon
\end{aligned}$$

by the definition of δ , noting that $\delta \leq 1$ guarantees $\delta^j \leq \delta$ for all $j \in \mathbb{N}$. This completes the proof as ϵ and t were arbitrary. \square

There are some properties of the operator exponential that are not considered properties of the real exponential, as they are trivial in the real case. The problem again is that commutativity is automatic for real numbers. For example, the real exponential equivalent of the following property is trivial.

Proposition A.7 *Let $L, M \in \mathcal{L}(X)$ commute. Then $e^L M = M e^L$.*

Proof. It is easily verified by induction on j that $L^j M = M L^j$ for all $j \in \mathbb{N}_0$ as a consequence of the commutativity of L and M (the $j = 0$ case is trivial) so $\sum_{j=0}^n \frac{1}{j!} L^j M = \sum_{j=0}^n \frac{1}{j!} M L^j$ for all $n \in \mathbb{N}$. The desired result then follows directly by taking limits on both sides, as the M can be pulled out of the sums and limits. \square

The following proposition is another such example. Here the proposition does not even require commutativity; rather the result is completely uninteresting with commutativity, so it is never worth considering for real exponentials.

Proposition A.8 *Let $L, M \in \mathcal{L}(X)$ with M invertible. Then $e^{MLM^{-1}} = Me^LM^{-1}$.*

Proof. It is easily verified by induction on j that $(MLM^{-1})^j = ML^jM^{-1}$ for all $j \in \mathbb{N}_0$ (the $j = 0$ case is again trivial) and it is an immediate consequence of this fact that $e^{MLM^{-1}} = \sum_{j=0}^{\infty} \frac{1}{j!} (MLM^{-1})^j = \sum_{j=0}^{\infty} \frac{1}{j!} ML^jM^{-1} = M(\sum_{j=0}^{\infty} \frac{1}{j!} L^j)M^{-1} = Me^LM^{-1}$. \square

It is sometimes necessary in this thesis to calculate the exponential of an operator with respect to some fixed basis. As a consequence of the previous proposition, since similarity transformations are exactly changes of basis, it is sufficient to calculate the exponential of the operator's Jordan canonical form. But any operator in Jordan canonical form is a block diagonal matrix, say $t\text{diag}(J_{d_1}(\lambda_1), \dots, J_{d_m}(\lambda_m))$, and it is clear that $\sum_{j=0}^n \frac{1}{j!} t\text{diag}(J_{d_1}(\lambda_1), \dots, J_{d_m}(\lambda_m))^j = \text{diag}(\sum_{j=0}^n \frac{1}{j!} t^j J_{d_1}(\lambda_1)^j, \dots, \sum_{j=0}^n \frac{1}{j!} t^j J_{d_m}(\lambda_m)^j)$ for all $n \in \mathbb{N}_0$ by induction on n (the $j = 0$ case is trivial as usual). Consequently, it follows that $e^{t\text{diag}(J_{d_1}(\lambda_1), \dots, J_{d_m}(\lambda_m))} = \text{diag}(e^{tJ_{d_1}(\lambda_1)}, \dots, e^{tJ_{d_m}(\lambda_m)})$, so it is really only necessary to calculate the exponential of a single Jordan block.

When $\mathbb{K} = \mathbb{C}$, a Jordan block $J_d(\lambda)$ is of the form

$$\begin{bmatrix} \lambda & 1 & 0 & & 0 & 0 \\ 0 & \lambda & 1 & & 0 & 0 \\ 0 & 0 & \lambda & & 0 & 0 \\ & & & \ddots & & \\ 0 & 0 & 0 & & \lambda & 1 \\ 0 & 0 & 0 & & 0 & \lambda \end{bmatrix}$$

for some $\lambda \in \mathbb{C}$. But then $e^{tJ_d(\lambda)} = e^{t\lambda I + tJ_d(0)} = e^{t\lambda I} e^{tJ_d(0)}$ as $t\lambda I$ and $tJ_d(0)$ necessarily commute for all $t \in \mathbb{R}$. Since $I^n = I$ for all $n \in \mathbb{N}_0$, it consequently must be the case that $\sum_{j=0}^{\infty} \frac{1}{j!} (t\lambda I)^j = \sum_{j=0}^{\infty} \frac{1}{j!} t^j \lambda^j I = (\sum_{j=0}^{\infty} \frac{1}{j!} t^j \lambda^j) I = e^{t\lambda} I$, so $e^{tJ_d(\lambda)} = e^{t\lambda} e^{J_d(0)}$ for all $t \in \mathbb{R}$. Now it is easily verified by induction on n that $J_d(0)^n$ is a matrix with ones down the n^{th} super diagonal and zeros everywhere else whenever $n < d$; moreover $J_d(0)^n = O$ for all

$n \geq d$. It follows that $e^{tJ_d(\lambda)} = e^{t\lambda} \sum_{j=0}^{d-1} \frac{1}{j!} t^j J_d(0)^j$ for all $t \in \mathbb{R}$, and so is a real $d \times d$ matrix of the form

$$e^{t\lambda} \begin{bmatrix} 1 & t & \frac{1}{2}t^2 & \frac{1}{(d-2)!}t^{d-2} & \frac{1}{(d-1)!}t^{d-1} \\ 0 & 1 & t & \frac{1}{(d-3)!}t^{d-3} & \frac{1}{(d-2)!}t^{d-2} \\ 0 & 0 & 1 & \frac{1}{(d-4)!}t^{d-4} & \frac{1}{(d-3)!}t^{d-3} \\ & & & \ddots & \\ 0 & 0 & 0 & 1 & t \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

If $\mathbb{K} = \mathbb{R}$, the situation is the same for Jordan blocks with $\lambda \in \mathbb{R}$. The situation for Jordan blocks with $\lambda \in \mathbb{C} \setminus \mathbb{R}$ is similar as well, but more complex. In this case the Jordan block is of the form

$$\begin{bmatrix} aI_2 + b\hat{I} & I_2 & 0 & 0 & 0 \\ 0 & aI_2 + b\hat{I} & I_2 & 0 & 0 \\ 0 & 0 & aI_2 + b\hat{I} & 0 & 0 \\ & & & \ddots & \\ 0 & 0 & 0 & aI_2 + b\hat{I} & I_2 \\ 0 & 0 & 0 & 0 & aI_2 + b\hat{I} \end{bmatrix} \quad \text{with } \hat{I} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

where $a = \Re\lambda$ and $b = \Im\lambda$. Applying similar reasoning as above, one concludes that $e^{t\tilde{J}_d(\lambda)} = e^{taI + t\text{diag}(b\hat{I}, \dots, \hat{I}) + tJ_d(0)^2} = e^{ta} \text{diag}(e^{tb\hat{I}}, \dots, e^{tb\hat{I}}) \sum_{j=0}^{\frac{d}{2}} \frac{1}{j!} t^j J_d(0)^{2j}$ for all $t \in \mathbb{R}$, and thus it is a real $2d \times 2d$ matrix of the form

$$e^{ta} \begin{bmatrix} E(t, b) & tE(t, b) & \frac{1}{2}t^2E(t, b) & \frac{1}{(d-2)!}t^{d-2}E(t, b) & \frac{1}{(d-1)!}t^{d-1}E(t, b) \\ 0 & E(t, b) & tE(t, b) & \frac{1}{(d-3)!}t^{d-3}E(t, b) & \frac{1}{(d-2)!}t^{d-2}E(t, b) \\ 0 & 0 & E(t, b) & \frac{1}{(d-4)!}t^{d-4}E(t, b) & \frac{1}{(d-3)!}t^{d-3}E(t, b) \\ & & & \ddots & \\ 0 & 0 & 0 & E(t, b) & tE(t, b) \\ 0 & 0 & 0 & 0 & E(t, b) \end{bmatrix}$$

where

$$E(t, b) := e^{tb\hat{I}} = \begin{bmatrix} \cos(tb) & \sin(tb) \\ -\sin(tb) & \cos(tb) \end{bmatrix}.$$

B Operator Integral

Operator integrals are defined and behave almost identically to the standard Riemann integral of real-valued functions. As such, it is necessary to start by defining notions such as the partition of an interval.

Definition B.1 A *partition* of $[a, b] \subset \mathbb{R}$ is a finite sequence $\{t_n\}_{n=0}^N$ of real numbers such that $\{t_n\}$ is strictly increasing (so $t_{n-1} < t_n$ for all $n \leq N$) and $t_0 = a$ and $t_N = b$. A *tagged partition* of $[a, b]$ is a partition of $[a, b]$ paired with a second finite sequence $\{\hat{t}_n\}_{n=1}^N$ of real numbers such that $\hat{t}_n \in [t_{n-1}, t_n]$ for all $n \leq N$. The *size* of a partition \mathcal{P} is given by $|\mathcal{P}| := \max_{n \leq N} |t_n - t_{n-1}|$. The *size* of a tagged partition $\hat{\mathcal{P}}$ is the size of its underlying partition \mathcal{P} .

Riemann sums of operators are also defined exactly as in the case of real-valued functions.

Definition B.2 Let $f : \mathbb{R} \rightarrow \mathcal{L}(X)$, and let $\hat{\mathcal{P}}$ be a tagged partition of $[a, b] \subset \mathbb{R}$. The *Riemann sum* of f over $\hat{\mathcal{P}}$ is given by $S(f, \hat{\mathcal{P}}) := \sum_{j=1}^N f(\hat{t}_j)(t_j - t_{j-1})$.

The operator integral is then constructed using Riemann sums.

Definition B.3 Let $f : \mathbb{R} \rightarrow \mathcal{L}(X)$. Then f is *integrable* on $[a, b] \subset \mathbb{R}$ if there exists an $M \in \mathcal{L}(X)$ such that, for every $\epsilon \in \mathbb{R}^+$, one may choose a $\delta \in \mathbb{R}^+$ so that $\|S(f, \hat{\mathcal{P}}) - M\| < \epsilon$ for every tagged partition $\hat{\mathcal{P}}$ satisfying $|\hat{\mathcal{P}}| < \delta$. The operator M is referred to as an (*operator*) *integral* of f on $[a, b]$ and is denoted $\int_a^b f(t) dt$. If f is integrable on $[a, b]$ for every $[a, b] \subset \mathbb{R}$, then f is *integrable*.

It is unsurprising that an operator integral is unique (for a given $[a, b]$) if it exists, and the proof of this fact is similar to the standard proof of uniqueness of limits.

Proposition B.4 Let $f : \mathbb{R} \rightarrow \mathcal{L}(X)$ be integrable on $[a, b] \subset \mathbb{R}$. If $M, \tilde{M} \in \mathcal{L}(X)$ are two operator integrals of f on $[a, b]$, then $M = \tilde{M}$.

Proof. Note that $\|M - \tilde{M}\| \leq \|M - S(f, \hat{\mathcal{P}})\| + \|S(f, \hat{\mathcal{P}}) - \tilde{M}\|$ for any tagged partition $\hat{\mathcal{P}}$ of $[a, b]$. In particular, given a fixed $\epsilon \in \mathbb{R}^+$, by definition there exist $\delta, \tilde{\delta} \in \mathbb{R}^+$ such that $|\hat{\mathcal{P}}| < \min\{\delta, \tilde{\delta}\}$ implies that $\|S(f, \hat{\mathcal{P}}) - M\| < \frac{\epsilon}{2}$ and $\|S(f, \hat{\mathcal{P}}) - \tilde{M}\| < \frac{\epsilon}{2}$, and consequently $\|M - \tilde{M}\| < \epsilon$ via an appropriate choice of tagged partition. Since ϵ was arbitrary, it follows that $M = \tilde{M}$. \square

For the purposes of this thesis it is essential that every linear flow be integrable. This is indeed the case as demonstrated in the next proposition, but it is first necessary to prove the following lemma.

Lemma B.5 *Let $f : \mathbb{R} \rightarrow \mathcal{L}(X)$, and let $[a, b] \subset \mathbb{R}$. Then f is integrable on $[a, b]$ if and only if for all $\epsilon \in \mathbb{R}^+$ there exists a $\delta \in \mathbb{R}^+$ such that $\|S(f, \hat{\mathcal{P}}) - S(f, \hat{\mathcal{Q}})\| < \epsilon$ whenever $\hat{\mathcal{P}}$ and $\hat{\mathcal{Q}}$ are tagged partitions of $[a, b]$ with $|\hat{\mathcal{P}}|, |\hat{\mathcal{Q}}| < \delta$.*

Proof. Suppose first that f is integrable on $[a, b]$, and fix an $\epsilon \in \mathbb{R}^+$. By definition there exists a $\delta \in \mathbb{R}^+$ such that any tagged partition $\hat{\mathcal{P}}$ of $[a, b]$ with $|\hat{\mathcal{P}}| < \delta$ must satisfy $\|S(f, \hat{\mathcal{P}}) - \int_a^b f(t) dt\| < \frac{\epsilon}{2}$. But then for two partitions $\hat{\mathcal{P}}$ and $\hat{\mathcal{Q}}$ such that $|\hat{\mathcal{P}}|, |\hat{\mathcal{Q}}| < \delta$ it must be that $\|S(f, \hat{\mathcal{P}}) - S(f, \hat{\mathcal{Q}})\| \leq \|S(f, \hat{\mathcal{P}}) - \int_a^b f(t) dt\| + \|S(f, \hat{\mathcal{Q}}) - \int_a^b f(t) dt\| < \epsilon$. This is sufficient to prove the ‘only if’ case as ϵ was arbitrary.

Suppose now that for all $\epsilon \in \mathbb{R}^+$ there exists a $\delta \in \mathbb{R}^+$ such that $\|S(f, \hat{\mathcal{P}}) - S(f, \hat{\mathcal{Q}})\| < \epsilon$ whenever $\hat{\mathcal{P}}$ and $\hat{\mathcal{Q}}$ are tagged partitions of $[a, b]$ with $|\hat{\mathcal{P}}|, |\hat{\mathcal{Q}}| < \delta$. Then in particular for any $n \in \mathbb{N}$ there exists a $\delta_n \in \mathbb{R}^+$ such that $\|S(f, \hat{\mathcal{P}}) - S(f, \hat{\mathcal{Q}})\| < \frac{1}{n}$ whenever $\hat{\mathcal{P}}$ and $\hat{\mathcal{Q}}$ are tagged partitions of $[a, b]$ with $|\hat{\mathcal{P}}|, |\hat{\mathcal{Q}}| < \delta_n$. It may be assumed without loss of generality that $\{\delta_n\}_{n \in \mathbb{N}}$ converges to 0. Now construct a sequence of tagged partitions $\{\hat{\mathcal{P}}_n\}_{n \in \mathbb{N}}$ of $[a, b]$ such that $|\hat{\mathcal{P}}_n| < \delta_n$ for all $n \in \mathbb{N}$. The induced sequence of Riemann sums $\{S(f, \hat{\mathcal{P}}_n)\}_{n \in \mathbb{N}}$ is clearly Cauchy by above, and since $\mathcal{L}(X)$ is a Banach space it follows that $\{S(f, \hat{\mathcal{P}}_n)\}_{n \in \mathbb{N}}$ converges to some limit $M \in \mathcal{L}(X)$.

Fix an $\epsilon \in \mathbb{R}^+$. As $\{S(f, \hat{\mathcal{P}}_n)\}_{n \in \mathbb{N}}$ converges to M , there exists a $\delta_{n_1} \in \mathbb{R}^+$ such that $\|S(f, \hat{\mathcal{P}}_n) - M\| < \frac{\epsilon}{2}$ whenever $|\hat{\mathcal{P}}_n| < \delta_{n_1}$. As $\{\delta_n\}_{n \in \mathbb{N}}$ converges to 0 there also exists a $\delta_{n_2} \in \mathbb{R}^+$ such that $\|S(f, \hat{\mathcal{Q}}) - S(f, \hat{\mathcal{P}}_n)\| < \frac{\epsilon}{2}$ whenever $|\hat{\mathcal{Q}}|, |\hat{\mathcal{P}}_n| < \delta_{n_2}$. Set $\delta = \min\{\delta_{n_1}, \delta_{n_2}\}$. Then for all tagged partitions $\hat{\mathcal{Q}}$ of $[a, b]$ with $|\hat{\mathcal{Q}}| < \delta$ it must be that $\|S(f, \hat{\mathcal{Q}}) - M\| \leq \|S(f, \hat{\mathcal{Q}}) - S(f, \hat{\mathcal{P}}_n)\| + \|S(f, \hat{\mathcal{P}}_n) - M\|$ for any $\hat{\mathcal{P}}_n$. As this holds for all $\hat{\mathcal{P}}_n$, this is in particular true for $\hat{\mathcal{P}}_n$ such that $|\hat{\mathcal{P}}_n| < \delta$, in which case $\|S(f, \hat{\mathcal{Q}}) - M\| < \epsilon$. As ϵ was arbitrary, this proves the ‘if’ case. \square

The standard approach to proving that every continuous real-valued map is integral uses upper and lower sums. Such constructs do not make sense in the context of Banach-valued maps, so an alternative approach is required. The preceding lemma demonstrates

that integrability on an interval is equivalent to requiring that all Riemann sums satisfy a Cauchy-type property. Using the preceding lemma, it is now possible to demonstrate that every continuous Banach-valued map is integrable, simply by showing that every such map satisfies this Cauchy-type property for every $[a, b] \subset \mathbb{R}$.

Proposition B.6 *Let $f : \mathbb{R} \rightarrow \mathcal{L}(X)$ be continuous. Then f is integrable.*

Proof. Fix $\epsilon \in \mathbb{R}^+$, and fix $[a, b]$. Since f is continuous, and since $[a, b]$ is compact, there exists a $\delta \in \mathbb{R}^+$ such that $\|f(t) - f(s)\| < \frac{\epsilon}{2(b-a)}$ whenever $|t - s| < \delta$. Let $\hat{\mathcal{P}}$ and $\hat{\mathcal{Q}}$ be two partitions of $[a, b]$ such that $|\hat{\mathcal{P}}|, |\hat{\mathcal{Q}}| < \delta$. Define the partition \mathcal{R} of $[a, b]$ to be the finite sequence $\{r_n\}_{n=0}^{N_R}$ generated by the distinct elements of $\{t_n\}_{n=0}^{N_P}$ and $\{s_n\}_{n=0}^{N_Q}$ taken in increasing order. Then form a tagged partition $\hat{\mathcal{R}}$ using \mathcal{R} in combination with any sequence $\{\hat{r}_n\}_{n=1}^{N_R}$. Note that necessarily $|\hat{\mathcal{R}}| < \delta$. It is sufficient to show that $\|S(f, \hat{\mathcal{P}}) - S(f, \hat{\mathcal{R}})\| < \frac{\epsilon}{2}$, as $\|S(f, \hat{\mathcal{P}}) - S(f, \hat{\mathcal{Q}})\| \leq \|S(f, \hat{\mathcal{P}}) - S(f, \hat{\mathcal{R}})\| + \|S(f, \hat{\mathcal{R}}) - S(f, \hat{\mathcal{Q}})\|$.

The idea now is to take advantage of the fact that the sequence $\{t_n\}_{n=0}^{N_P}$ is contained in $\{r_n\}_{n=0}^{N_R}$. Construct a sequence $\{\hat{t}'_n\}_{n=1}^{N_R}$ by setting $\hat{t}'_n = \hat{t}_j$ whenever $r_n \in (t_{j-1}, t_j]$ for all $j \leq N_P$ and $n \leq N_R$. Then $S(f, \hat{\mathcal{P}}) = \sum_{j=0}^{N_P} f(\hat{t}'_j)(r_n - r_{n-1})$, and it follows that

$$\begin{aligned} \|S(f, \hat{\mathcal{P}}) - S(f, \hat{\mathcal{R}})\| &= \left\| \sum_{n=1}^{N_R} f(\hat{t}'_n)(r_n - r_{n-1}) - \sum_{n=1}^{N_R} f(\hat{r}_n)(r_n - r_{n-1}) \right\| \\ &= \left\| \sum_{n=1}^{N_R} (f(\hat{t}'_n) - f(\hat{r}_n))(r_n - r_{n-1}) \right\| \\ &\leq \sum_{n=1}^{N_R} \|f(\hat{t}'_n) - f(\hat{r}_n)\| (r_n - r_{n-1}) \\ &< \sum_{n=1}^{N_R} \frac{\epsilon}{2(b-a)} (r_n - r_{n-1}) \\ &= \frac{\epsilon}{2} \end{aligned}$$

since $\sum_{n=1}^{N_R} (r_n - r_{n-1}) = b - a$. As ϵ was arbitrary, it follows from Lemma B.5 that f is integrable on $[a, b]$. This completes the proof, as $[a, b]$ was arbitrary. \square

Not only are continuous maps all integrable as in the case of real integrals, but many of the basic properties of operator integrals are similar to properties of real integrals. For example, operator integrals behave well under operator addition and scalar multiplication.

Proposition B.7 *Let $f, g : \mathbb{R} \rightarrow \mathcal{L}(X)$ be continuous, let $[a, b] \subset \mathbb{R}$, and let $x, y \in \mathbb{K}$. Then $\int_a^b xf(t) + yg(t) dt = x \int_a^b f(t) dt + y \int_a^b g(t) dt$.*

Proof. Let $\hat{\mathcal{P}}$ be a tagged partition of $[a, b]$. It is clear that $S(xf + yg, \hat{\mathcal{P}}) = xS(f, \hat{\mathcal{P}}) + yS(g, \hat{\mathcal{P}})$ holds, so $\|\int_a^b xf(t) + yg(t) dt - (x \int_a^b f(t) dt + y \int_a^b g(t) dt)\|$ is less than or equal to $\|\int_a^b xf(t) + yg(t) dt - S(xf + yg, \hat{\mathcal{P}})\| + |x| \|S(f, \hat{\mathcal{P}}) - \int_a^b f(t) dt\| + |y| \|S(g, \hat{\mathcal{P}}) - \int_a^b g(t) dt\|$ by the triangle inequality. It follows that $\|\int_a^b xf(t) + yg(t) dt - (x \int_a^b f(t) dt + y \int_a^b g(t) dt)\|$ may be made arbitrarily small by choosing a tagged partition $\hat{\mathcal{P}}$ of $[a, b]$ with $|\hat{\mathcal{P}}|$ sufficiently small. \square

Proposition B.8 *Let $f : \mathbb{R} \rightarrow \mathcal{L}(X)$ be continuous, let $[a, b] \subset \mathbb{R}$, and let $L \in \mathcal{L}(X)$. Then $L \int_a^b f(t) dt = \int_a^b Lf(t) dt$ and $\int_a^b f(t) dt L = \int_a^b f(t) L dt$.*

Proof. Note that for any tagged partition $\hat{\mathcal{P}}$ of $[a, b]$ it must hold that $S(Lf, \hat{\mathcal{P}}) = LS(f, \hat{\mathcal{P}})$ and $S(fL, \hat{\mathcal{P}}) = S(f, \hat{\mathcal{P}})L$. The rest of the proof follows similarly to the proof of the previous proposition. \square

It is often useful to manipulate the intervals of a real integral, and these techniques extend to the operator integral as well.

Proposition B.9 *Let $f : \mathbb{R} \rightarrow \mathcal{L}(X)$ be continuous. Then $\int_a^b f(t) dt = \int_a^c f(t) dt + \int_c^b f(t) dt$ for any $c \in (a, b)$.*

Proof. Let $\hat{\mathcal{P}}$ be any tagged partition of $[a, b]$ so that c appears in the sequence $\{t_n\}_{n=0}^N$. Then $\hat{\mathcal{P}}$ splits into two partitions $\hat{\mathcal{P}}_1$ and $\hat{\mathcal{P}}_2$ of $[a, c]$ and $[c, b]$ respectively, and necessarily $S(f, \hat{\mathcal{P}}) = S(f, \hat{\mathcal{P}}_1) + S(f, \hat{\mathcal{P}}_2)$. Now $\|\int_a^b f(t) dt - (\int_a^c f(t) dt + \int_c^b f(t) dt)\|$ is less than or equal to $\|\int_a^b f(t) dt - S(f, \hat{\mathcal{P}})\| + \|S(f, \hat{\mathcal{P}}_1) - \int_a^c f(t) dt\| + \|S(f, \hat{\mathcal{P}}_2) - \int_c^b f(t) dt\|$. It follows that $\|\int_a^b f(t) dt - (\int_a^c f(t) dt + \int_c^b f(t) dt)\|$ can be made arbitrarily small by choosing $\hat{\mathcal{P}}$ as above with $|\hat{\mathcal{P}}|$ sufficiently small. \square

Proposition B.10 *Let $f : \mathbb{R} \rightarrow \mathcal{L}(X)$ be continuous. Then $\int_a^{a+b} f(t) dt = \int_0^b f(t+a) dt$ for all $a \in \mathbb{R}$ and $b \in \mathbb{R}^+$.*

Proof. Consider that any tagged partition $\hat{\mathcal{P}}$ of $[0, b]$ generates a tagged partition $\hat{\mathcal{Q}}$ of $[a, a+b]$ by setting $s_n = t_n + a$ for all $0 \leq n \leq N$ and $\hat{s}_n = \hat{t}_n + a$ for all $1 \leq n \leq N$; moreover $|\hat{\mathcal{P}}| = |\hat{\mathcal{Q}}|$ and $S(f, \hat{\mathcal{Q}}) = S(\tilde{f}, \hat{\mathcal{P}})$ where $\tilde{f}(t) := f(t+a)$. It must hold that $\|\int_a^{a+b} f(t) dt - \int_0^b f(t+a) dt\| \leq \|S(f, \hat{\mathcal{Q}}) - \int_a^{a+b} f(t) dt\| + \|S(\tilde{f}, \hat{\mathcal{P}}) - \int_0^b f(t+a) dt\|$,

so $\|\int_a^{a+b} f(t) dt - \int_0^b f(t+a) dt\|$ can be made arbitrarily small by choosing $\hat{\mathcal{P}}$ with $|\hat{\mathcal{P}}|$ sufficiently small. \square

For the purposes of this thesis, it is essential to prove the following proposition.

Proposition B.11 *Let $f : \mathbb{R} \rightarrow \mathcal{L}(X)$ be continuous. Then $\lim_{s \rightarrow 0} \frac{1}{s} \int_a^{a+s} f(t) dt = f(a)$ for all $a \in \mathbb{R}$.*

Proof. Fix $\epsilon \in \mathbb{R}^+$. By continuity there exists a $\delta \in \mathbb{R}^+$ such that $\|f(t) - f(a)\| < \frac{\epsilon}{2}$ whenever $|t - a| < \delta$. Let $s < \delta$. Now for any tagged partition $\hat{\mathcal{P}}$ of $[a, a+s]$ it must hold that

$$\begin{aligned}
\left\| \frac{1}{s} \int_a^{a+s} f(t) dt - f(a) \right\| &\leq \left\| \int_a^{a+s} \frac{1}{s} f(t) dt - S\left(\frac{1}{s}f, \hat{\mathcal{P}}\right) \right\| + \left\| S\left(\frac{1}{s}f, \hat{\mathcal{P}}\right) - \frac{1}{s}f(a)s \right\| \\
&= \left\| \int_a^{a+s} \frac{1}{s} f(t) dt - S\left(\frac{1}{s}f, \hat{\mathcal{P}}\right) \right\| \\
&\quad + \left\| \sum_{n=1}^N \frac{1}{s} f(\hat{t}_n)(t_n - t_{n-1}) - \sum_{n=1}^N \frac{1}{s} f(a)(t_n - t_{n-1}) \right\| \\
&= \left\| \int_a^{a+s} \frac{1}{s} f(t) dt - S\left(\frac{1}{s}f, \hat{\mathcal{P}}\right) \right\| \\
&\quad + \frac{1}{s} \left\| \sum_{n=1}^N (f(\hat{t}_n) - f(a))(t_n - t_{n-1}) \right\| \\
&\leq \left\| \int_a^{a+s} \frac{1}{s} f(t) dt - S\left(\frac{1}{s}f, \hat{\mathcal{P}}\right) \right\| \\
&\quad + \frac{1}{s} \sum_{n=1}^N \|f(\hat{t}_n) - f(a)\| (t_n - t_{n-1}) \\
&< \left\| \int_a^{a+s} \frac{1}{s} f(t) dt - S\left(\frac{1}{s}f, \hat{\mathcal{P}}\right) \right\| + \frac{\epsilon}{2}.
\end{aligned}$$

In particular this is true for any tagged partition with $|\hat{\mathcal{P}}|$ sufficiently small that $\hat{\mathcal{P}}$ satisfies $\|\int_a^{a+s} \frac{1}{s} f(t) dt - S(\frac{1}{s}f, \hat{\mathcal{P}})\| < \frac{\epsilon}{2}$, so $\|\frac{1}{s} \int_a^{a+s} f(t) dt - f(a)\| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$. As ϵ was arbitrary, the result follows as a and s were also arbitrary. \square

C Complexification and Realification

It is often the case that complex normed spaces are easier to work with than real normed spaces. As such, when working with real normed spaces, it is desirable to be able to in some sense convert the spaces to complex normed spaces prior to analyzing them. This conversion

can be done in a canonical fashion called complexification, wherein a real normed space is viewed as being embedded in a complex normed space.

Definition C.1 The *complexification* of X with $\mathbb{K} = \mathbb{R}$, denoted $X_{\mathbb{C}}$, is given by the set $\{(x_1, x_2) : x_1, x_2 \in X\}$ equipped with notions of addition and (complex) scalar multiplication, where addition is given by $(x_1, x_2) + (y_1, y_2) = (x_1 + y_1, x_2 + y_2)$ for all $(x_1, x_2), (y_1, y_2) \in X_{\mathbb{C}}$ and scalar multiplication is given by $(a + ib)(x_1, x_2) := (ax_1 - bx_2, ax_2 + bx_1)$ for all $a + ib \in \mathbb{C}$ and $(x_1, x_2) \in X_{\mathbb{C}}$.

For simplicity $X_{\mathbb{C}} := X$ for X with $\mathbb{K} = \mathbb{C}$. As such, it is only interesting to consider X with $\mathbb{K} = \mathbb{R}$, so this will be assumed to be the case until stated otherwise. In either case $X_{\mathbb{C}}$ is a finite-dimensional normed space over \mathbb{C} .

Proposition C.2 $X_{\mathbb{C}}$ is a normed space over \mathbb{C} with $\dim_{\mathbb{C}} X_{\mathbb{C}} = \dim_{\mathbb{R}} X$.

Proof. First it is shown that $X_{\mathbb{C}}$ is a vector space over \mathbb{C} . It is clear from the definition that $x + y \in X_{\mathbb{C}}$ for all $x, y \in X_{\mathbb{C}}$, and similarly that $cx \in X_{\mathbb{C}}$ for all $c \in \mathbb{C}$ and $x \in X_{\mathbb{C}}$. The associativity and commutativity of addition follow from the definition of addition and the fact that X is a vector space; moreover, it is clear that $0_{\mathbb{C}}$ is given by $(0, 0)$ and $-x$ is given by $(-x_1, -x_2)$ for all $x = (x_1, x_2) \in X_{\mathbb{C}}$. To see that scalar multiplication also behaves properly, first note that clearly $1x = x$ for all $x \in X_{\mathbb{C}}$. Scalar multiplication is associative, since

$$\begin{aligned} c_1(c_2x) &= (a_1 + ib_1)(a_2x_1 - b_2x_2, a_2x_2 + b_2x_1) \\ &= (a_1a_2x_1 - a_1b_2x_2 - b_1a_2x_2 - b_1b_2x_1, a_1a_2x_2 + a_1b_2x_1 + b_1a_2x_1 - b_1b_2x_2) \\ &= ((a_1a_2 - b_1b_2) + i(a_1b_2 + b_1a_2))(x_1, x_2) \\ &= (c_1c_2)x \end{aligned}$$

for all $x = (x_1, x_2) \in X_{\mathbb{C}}$ and $c_1 = a_1 + ib_1, c_2 = a_2 + ib_2 \in \mathbb{C}$. Finally,

$$\begin{aligned} c(x + y) &= (a + ib)(x_1 + y_1, x_2 + y_2) \\ &= (ax_1 + ay_1 - bx_2 - by_2, ax_2 + ay_2 + bx_1 + by_1) \\ &= (ax_1 - bx_2, ax_2 + bx_1) + (ay_1 - by_2, ay_2 + by_1) \\ &= cx + cy \end{aligned}$$

for all $x = (x_1, x_2), y = (y_1, y_2) \in X_{\mathbb{C}}$ and $c = a + ib \in \mathbb{C}$, and

$$\begin{aligned}
(c_1 + c_2)x &= ((a_1 + a_2) + i(b_1 + b_2))(x_1, x_2) \\
&= (a_1x_1 + a_2x_1 - b_1x_2 - b_2x_2, a_1x_2 + a_2x_2 + b_1x_1 + b_2x_1) \\
&= (a_1x_1 - b_1x_2, a_1x_2 + b_1x_1) + (a_2x_1 - b_2x_2, a_2x_2 + b_2x_1) \\
&= c_1x + c_2x
\end{aligned}$$

for all $x = (x_1, x_2) \in X_{\mathbb{C}}$ and $c_1 = a_1 + ib_1, c_2 = a_2 + ib_2 \in \mathbb{C}$, so addition and scalar multiplication distribute appropriately.

To see that $X_{\mathbb{C}}$ is a normed space, consider the map $\|\cdot\|_{\mathbb{C}} : X_{\mathbb{C}} \rightarrow \mathbb{R}$ defined by setting $\|x\|_{\mathbb{C}} = \sup_{\theta \in [0, 2\pi]} \|\cos(\theta)x_1 + \sin(\theta)x_2\|$ for all $x = (x_1, x_2) \in X_{\mathbb{C}}$. Clearly $\|x\|_{\mathbb{C}} \geq 0$ for all $x \in X_{\mathbb{C}}$ with $\|0\|_{\mathbb{C}} = 0$. If $\|x\|_{\mathbb{C}} = 0$ for some $x = (x_1, x_2) \in X_{\mathbb{C}}$, then it must be that $\cos(\theta)x_1 + \sin(\theta)x_2 = 0$ for all $\theta \in [0, 2\pi]$. In particular, this must be true for $\theta = 0$ and $\theta = \frac{\pi}{2}$, so $x_1 = 0$ and $x_2 = 0$, and consequently $x = 0$. Now

$$\begin{aligned}
\|x + y\|_{\mathbb{C}} &= \|(x_1 + y_1, x_2 + y_2)\|_{\mathbb{C}} \\
&= \sup_{\theta \in [0, 2\pi]} \|\cos(\theta)x_1 + \cos(\theta)y_1 + \sin(\theta)x_2 + \sin(\theta)y_2\| \\
&\leq \sup_{\theta \in [0, 2\pi]} (\|\cos(\theta)x_1 + \sin(\theta)x_2\| + \|\cos(\theta)y_1 + \sin(\theta)y_2\|) \\
&\leq \sup_{\theta \in [0, 2\pi]} \|\cos(\theta)x_1 + \sin(\theta)x_2\| + \sup_{\theta \in [0, 2\pi]} \|\cos(\theta)y_1 + \sin(\theta)y_2\| \\
&= \|x\|_{\mathbb{C}} + \|y\|_{\mathbb{C}}
\end{aligned}$$

for all $x = (x_1, x_2), y = (y_1, y_2) \in X_{\mathbb{C}}$. Finally,

$$\begin{aligned}
\|cx\|_{\mathbb{C}} &= \|(r_c \cos(\theta_c)x_1 - r_c \sin(\theta_c)x_2, r_c \cos(\theta_c)x_2 + r_c \sin(\theta_c)x_1)\|_{\mathbb{C}} \\
&= \sup_{\theta \in [0, 2\pi]} \|r_c \cos(\theta) \cos(\theta_c)x_1 - r_c \cos(\theta) \sin(\theta_c)x_2 \\
&\quad + r_c \sin(\theta) \cos(\theta_c)x_2 + r_c \sin(\theta) \sin(\theta_c)x_1\| \\
&= r_c \sup_{\theta \in [0, 2\pi]} \|\cos(\theta - \theta_c)x_1 + \sin(\theta - \theta_c)x_2\| \\
&= |r_c(\cos(\theta_c) + i \sin(\theta_c))| \sup_{\theta \in [0, 2\pi]} \|\cos(\theta)x_1 + \sin(\theta)x_2\| \\
&= |c| \|x\|_{\mathbb{C}}
\end{aligned}$$

for all $x = (x_1, x_2) \in X_{\mathbb{C}}$ and $c = r_c(\cos(\theta_c) + i \sin(\theta_c)) \in \mathbb{C}$. It follows that $\|\cdot\|_{\mathbb{C}}$ is a norm on $X_{\mathbb{C}}$.

Now let $\{b_n\}_{n=1}^d$ be a basis for X , and consider $\{(b_n, 0)\}_{n=1}^d$ in $X_{\mathbb{C}}$. Given a point $x = (x_1, x_2) \in X_{\mathbb{C}}$, there exists $a_n, \tilde{a}_n \in \mathbb{R}$ for $n \leq d$ such that $x_1 = \sum_{n=1}^d a_n b_n$ and $x_2 = \sum_{n=1}^d \tilde{a}_n b_n$ since $\{b_n\}_{n=1}^d$ is a basis for X . Now define $\{c_n\}_{n=1}^d$ in \mathbb{C} by $c_n = a_n + i\tilde{a}_n$. Then

$$x = (x_1, x_2) = \left(\sum_{n=1}^d a_n b_n, \sum_{n=1}^d \tilde{a}_n b_n \right) = \sum_{n=1}^d (a_n b_n, \tilde{a}_n b_n) = \sum_{n=1}^d c_n (b_n, 0),$$

and thus $\{(b_n, 0)\}_{n=1}^d$ spans $X_{\mathbb{C}}$ since x was arbitrary. On the other hand, suppose that $\sum_{n=1}^d c_n (b_n, 0) = 0$ for some $c_n = a_n + i\tilde{a}_n \in \mathbb{C}$. Then $(0, 0) = (\sum_{n=1}^d a_n b_n, \sum_{n=1}^d \tilde{a}_n b_n)$. Since $\{b_n\}_{n=1}^d$ is a basis for X , it must be that $a_n = \tilde{a}_n = 0$ for all $n \leq d$. But then $c_n = 0$ for all $n \leq d$. It follows that $\{(b_n, 0)\}_{n=1}^d$ is linearly independent, and consequently $\{(b_n, 0)\}_{n=1}^d$ is a basis for $X_{\mathbb{C}}$. Thus $\dim_{\mathbb{C}} X_{\mathbb{C}} = \dim_{\mathbb{R}} X < \infty$. \square

It makes sense to write $x_1 + ix_2$ in place of $(x_1, x_2) \in X_{\mathbb{C}}$. Identifying X with the subspaces $\{(x, 0) : x \in X\}$ and $\{(0, x) : x \in X\}$, it follows that $X_{\mathbb{C}}$ can be viewed as $X \oplus iX$. While it took some work to construct a norm on $X_{\mathbb{C}}$ built from a norm on X , it is easily seen that the restriction any norm on $X_{\mathbb{C}}$ induces a norm on X . Now that it is clear that $X_{\mathbb{C}}$ is a complex finite-dimensional normed space as desired, one may also define the complexification of a map in a straightforward fashion.

Definition C.3 Let $f : X \rightarrow Y$. The *complexification* of f , denoted $f_{\mathbb{C}} : X_{\mathbb{C}} \rightarrow Y_{\mathbb{C}}$, is given by $f_{\mathbb{C}}(x) = f(x_1) + if(x_2)$, where $x_1, x_2 \in X$ such that $x = x_1 + ix_2$. Let φ be a flow on X . The *complexification* of φ , denoted $\varphi_{\mathbb{C}} : \mathbb{R} \times X_{\mathbb{C}} \rightarrow X_{\mathbb{C}}$, is given by $\varphi_{\mathbb{C}}(t, x) = \varphi(t, x_1) + i\varphi(t, x_2)$, where $x_1, x_2 \in X$ such that $x = x_1 + ix_2$.

Many properties of maps are preserved by their complexifications. For example the complexification of a continuous map is itself continuous.

Proposition C.4 Let $f : X \rightarrow Y$ be continuous. Then $f_{\mathbb{C}}$ is continuous.

Proof. Let $x_0 \in X_{\mathbb{C}}$, and let $\{x_n\}_{n \in \mathbb{N}}$ be a sequence in $X_{\mathbb{C}}$ converging to x_0 . Then there exists $x_{0,1}, x_{0,2} \in X$ and sequences $\{x_{n,1}\}_{n \in \mathbb{N}}$ and $\{x_{n,2}\}_{n \in \mathbb{N}}$ in X such that $x_0 = x_{0,1} + ix_{0,2}$ and $x_n = x_{n,1} + ix_{n,2}$ for all $n \in \mathbb{N}$. It follows from the convergence of $\{x_n\}_{n \in \mathbb{N}}$ that the sequences $\{x_{n,1}\}_{n \in \mathbb{N}}$ and $\{ix_{n,2}\}_{n \in \mathbb{N}}$ converge to $x_{0,1}$ and $ix_{0,2}$ respectively, and thus the

sequence $\{x_{n,2}\}_{n \in \mathbb{N}}$ converges to $x_{0,2}$. Consequently,

$$\begin{aligned} \lim_{n \rightarrow \infty} f_{\mathbb{C}}(x_n) &= \lim_{n \rightarrow \infty} (f(x_{n,1}) + if(x_{n,2})) \\ &= \lim_{n \rightarrow \infty} f(x_{n,1}) + i \lim_{n \rightarrow \infty} f(x_{n,2}) \\ &= f(x_{0,1}) + if(x_{0,2}) \\ &= f_{\mathbb{C}}(x_0), \end{aligned}$$

and $f_{\mathbb{C}}$ is therefore continuous as $\{x_n\}_{n \in \mathbb{N}}$ and x_0 were arbitrary. \square

A similar result holds for linear operators.

Proposition C.5 *Let $L \in \mathcal{L}(X, Y)$. Then $L_{\mathbb{C}} \in \mathcal{L}(X_{\mathbb{C}}, Y_{\mathbb{C}})$.*

Proof. Continuity of $L_{\mathbb{C}}$ follows from the previous proposition, so it suffices to demonstrate linearity. Let $x, y \in X_{\mathbb{C}}$, and let $c_1, c_2 \in \mathbb{C}$. Then there exist $x_1, x_2, y_1, y_2 \in X$ and $a_1, b_1, a_2, b_2 \in \mathbb{R}$ such that $x = x_1 + ix_2, y = y_1 + iy_2, c_1 = a_1 + ib_1$, and $c_2 = a_2 + ib_2$. It follows that

$$\begin{aligned} L_{\mathbb{C}}(c_1x + c_2y) &= L_{\mathbb{C}}((a_1x_1 - b_1x_2) + i(a_1x_2 + b_1x_1) + (a_2y_1 - b_2y_2) + i(a_2y_2 + b_2y_1)) \\ &= L_{\mathbb{C}}((a_1x_1 - b_1x_2 + a_2y_1 - b_2y_2) + i(a_1x_2 + b_1x_1 + a_2y_2 + b_2y_1)) \\ &= L(a_1x_1 - b_1x_2 + a_2y_1 - b_2y_2) + iL(a_1x_2 + b_1x_1 + a_2y_2 + b_2y_1) \\ &= a_1Lx_1 - b_1Lx_2 + a_2Ly_1 - b_2Ly_2 + i(a_1Lx_2 + b_1Lx_1 + a_2Ly_2 + b_2Ly_1) \\ &= (a_1 + ib_1)(Lx_1 + iLx_2) + (a_2 + ib_2)(Ly_1 + iLy_2) \\ &= c_1L_{\mathbb{C}}x + c_2L_{\mathbb{C}}y, \end{aligned}$$

so $L_{\mathbb{C}}$ is linear (over \mathbb{C}) as x, y, c_1 , and c_2 were all arbitrary. \square

Not only does the complexification of a linear operator result in a linear operator, but in fact the complexification of a linear combination of linear operators results in a linear operator that is simply the linear combination of the complexifications of the linear operators in the linear combination.

Proposition C.6 *Let $L, M \in \mathcal{L}(X)$, and let $r, s \in \mathbb{R}$. Then $(rL + sM)_{\mathbb{C}} = rL_{\mathbb{C}} + sM_{\mathbb{C}}$.*

Proof. Let $x \in X_{\mathbb{C}}$. There exist $x_1, x_2 \in X$ such that $x = x_1 + ix_2$. Now

$$\begin{aligned}
(rL + sM)_{\mathbb{C}}x &= (rL + sM)x_1 + i(rL + sM)x_2 \\
&= rLx_1 + sMx_2 + i(rLx_2 + sMx_2) \\
&= r(Lx_1 + iLx_2) + s(Mx_1 + iMx_2) \\
&= rL_{\mathbb{C}}x + sM_{\mathbb{C}}x,
\end{aligned}$$

so $(rL + sM)_{\mathbb{C}} = rL_{\mathbb{C}} + sM_{\mathbb{C}}$ as x was arbitrary. \square

Note that any linear operator L can be viewed as a real matrix once a basis has been fixed. If this matrix is instead viewed as a complex matrix (with entries that happen to be all real) then $iL = Li$, and consequently $L_{\mathbb{C}}x = Lx_1 + iLx_2 = Lx_1 + Li x_2 = L(x_1 + ix_2) = Lx$, so the matrix representation of $L_{\mathbb{C}}$ is the same as the matrix representation of L once a basis is fixed.

Proposition C.7 *Let φ be a flow on X . Then $\varphi_{\mathbb{C}}$ is a flow on $X_{\mathbb{C}}$.*

Proof. Let $x \in X_{\mathbb{C}}$. Then there exists $x_1, x_2 \in X$ such that $x = x_1 + ix_2$. But then $\varphi_{\mathbb{C}}(0, x) = \varphi(0, x_1) + i\varphi(0, x_2) = x_1 + ix_2 = x$. Also

$$\begin{aligned}
\varphi_{\mathbb{C}}(s, \varphi_{\mathbb{C}}(t, x)) &= \varphi_{\mathbb{C}}(s, \varphi(t, x_1) + i\varphi(t, x_2)) \\
&= \varphi(s, \varphi(t, x_1)) + i\varphi(s, \varphi(t, x_2)) \\
&= \varphi(s + t, x_1) + i\varphi(s + t, x_2) \\
&= \varphi_{\mathbb{C}}(s + t, x)
\end{aligned}$$

for all $s, t \in \mathbb{R}$. As x was arbitrary, it remains to show that $\varphi_{\mathbb{C}}$ is continuous. This is proved similarly to the proof of Proposition C.4. \square

It follows from Proposition C.5 that the complexification of a linear flow is itself a linear flow. Furthermore, it is easily seen that the complexification of a linear flow is generated by the complexification of the operator that generates the original flow.

Proposition C.8 *Let φ be a linear flow on X . If $L \in \mathcal{L}(X)$ generates φ , then $\varphi_{\mathbb{C}}$ is a linear flow on $X_{\mathbb{C}}$ generated by $L_{\mathbb{C}}$.*

Proof. Linearity follows from Proposition C.5. To show that $\varphi_{\mathbb{C}}$ is generated by $L_{\mathbb{C}}$, it suffices to show that the complexification of e^{tL} is really just $e^{tL_{\mathbb{C}}}$. Let $x \in X_{\mathbb{C}}$. There exists $x_1, x_2 \in X$ such that $x = x_1 + ix_2$. Note first that $(L_{\mathbb{C}})^j = (L^j)_{\mathbb{C}}$ for all $j \in \mathbb{N}$ by induction on j , since $(L_{\mathbb{C}})^{j+1}x = (L_{\mathbb{C}})^j L_{\mathbb{C}}x = (L^j)_{\mathbb{C}}(Lx_1 + iLx_2) = L^{j+1}x_1 + iL^{j+1}x_2 = (L^{j+1})_{\mathbb{C}}x$ if $(L_{\mathbb{C}})^j = (L^j)_{\mathbb{C}}$, so $(L_{\mathbb{C}})^{j+1} = (L^{j+1})_{\mathbb{C}}$ as x was arbitrary. Let $t \in \mathbb{R}$. It now follows from Proposition C.6 that

$$\begin{aligned}
(e^{tL})_{\mathbb{C}}x &= e^{tL}x_1 + ie^{tL}x_2 \\
&= \lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} t^j L^j x_1 + i \lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} t^j L^j x_2 \\
&= \lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} t^j (L^j x_1 + iL^j x_2) \\
&= \lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} t^j (L^j)_{\mathbb{C}}x \\
&= \lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} t^j (L_{\mathbb{C}})^j x \\
&= e^{tL_{\mathbb{C}}}x
\end{aligned}$$

and thus $(e^{tL})_{\mathbb{C}} = e^{tL_{\mathbb{C}}}$ as x and t were arbitrary. \square

Even flow equivalence is preserved by complexification.

Proposition C.9 *Let h be a flow equivalence between flows φ and ψ on X and Y respectively. Then $h_{\mathbb{C}}$ is a flow equivalence between $\varphi_{\mathbb{C}}$ and $\psi_{\mathbb{C}}$.*

Proof. First it is shown that $h_{\mathbb{C}}$ is a bijection. Let $y \in Y_{\mathbb{C}}$. Then $y = y_1 + iy_2$ for some $y_1, y_2 \in Y$, but then $y = h(x_1) + ih(x_2) = h_{\mathbb{C}}(x_1 + ix_2)$ for some $x_1, x_2 \in X$, and consequently $y = h_{\mathbb{C}}(x)$ for some $x \in X_{\mathbb{C}}$. Thus $h_{\mathbb{C}}$ is surjective. Let $x, \hat{x} \in X_{\mathbb{C}}$. There exists $x_1, x_2, \hat{x}_1, \hat{x}_2 \in X$ such that $x = x_1 + ix_2$ and $\hat{x} = \hat{x}_1 + i\hat{x}_2$. If $h_{\mathbb{C}}(x) = h_{\mathbb{C}}(\hat{x})$, then $h(x_1) + ih(x_2) = h(\hat{x}_1) + ih(\hat{x}_2)$, and thus $h(x_1) = h(\hat{x}_1)$ and $h(x_2) = h(\hat{x}_2)$. It then follows from the injectivity of h that $x_1 = \hat{x}_1$ and $x_2 = \hat{x}_2$, so $x = \hat{x}$. Thus $h_{\mathbb{C}}$ is injective. Finally, let $x \in X_{\mathbb{C}}$. There exists $x_1, x_2 \in X$ such that $x = x_1 + ix_2$. It follows directly from the Definition C.3, along with the fact that h is an equivalence between φ and ψ , that $h_{\mathbb{C}}(\varphi_{\mathbb{C}}(t, x)) = \psi_{\mathbb{C}}(t, h_{\mathbb{C}}(x))$ for all $t \in \mathbb{R}$. This is sufficient to complete the proof as x was arbitrary. \square

It follows from Proposition C.4 that the complexification of a homeomorphic equivalence is itself a homeomorphic equivalence. Similarly, it follows from Proposition C.5 that the complexification of a linear equivalence is itself a linear equivalence.

Not only are many properties of maps preserved by complexification, but decompositions are preserved by complexification as well.

Proposition C.10 *If $X = \bigoplus_{k=1}^m X_k$, then $X_{\mathbb{C}} = \bigoplus_{k=1}^m (X_k)_{\mathbb{C}}$. Let $L \in \mathcal{L}(X)$. If L can be written as $\bigoplus_{k=1}^m L_k$ for some $L_k \in \mathcal{L}(X_k)$, then $L_{\mathbb{C}} = \bigoplus_{k=1}^m (L_k)_{\mathbb{C}}$. Let φ be a linear flow on X . If φ can be written as $\bigoplus_{k=1}^m \varphi_k$ for some linear flows φ_k on X_k , then $\varphi_{\mathbb{C}} = \bigoplus_{k=1}^m (\varphi_k)_{\mathbb{C}}$.*

Proof. It is easily verified that $X_{\mathbb{C}} = \sum_{k=1}^m (X_k)_{\mathbb{C}}$. Now fix $k \leq m$. Given some point $x \in (X_k)_{\mathbb{C}} \cap \sum_{j \neq k} (X_j)_{\mathbb{C}}$, then there must exist some $x_{j,1}, x_{j,2} \in X_j$ for $j \leq m$ such that $x_{k,1} + ix_{k,2} = x_k = \sum_{j \neq k} x_{j,1} + ix_{j,2}$. But then $x_{k,1} = \sum_{j \neq k} x_{j,1}$ and $x_{k,2} = \sum_{j \neq k} x_{j,2}$. As by assumption $X_k \cap \sum_{j \neq k} X_j = \{0\}$, it follows that $x = 0$. Thus $X_{\mathbb{C}} = \sum_{k=1}^m (X_k)_{\mathbb{C}} = \{0\}$, as x was arbitrary. As k was also arbitrary, this holds for all $k \leq m$, and this is sufficient to prove the first statement.

For the other two statements, let $x \in X_{\mathbb{C}}$. There exist $y, z \in X$ such that $x = y + iz$. Now

$$\begin{aligned}
L_{\mathbb{C}}x &= Ly + iLz \\
&= \left(\bigoplus_{k=1}^m L_k \right) y + i \left(\bigoplus_{k=1}^m L_k \right) z \\
&= \sum_{k=1}^m L_k y_k + i \sum_{k=1}^m L_k z_k \\
&= \sum_{k=1}^m (L_k y_k + iL_k z_k) \\
&= \sum_{k=1}^m (L_k)_{\mathbb{C}} x_k \\
&= \left(\bigoplus_{k=1}^m (L_k)_{\mathbb{C}} \right) x
\end{aligned}$$

and a similar argument works for $\varphi_{\mathbb{C}}$. \square

So many things are preserved by complexification that one may wonder what isn't preserved. It turns out that irreducibility is not preserved by complexification. Of course the irreducible decomposition of a flow necessarily generates a decomposition of the complexification by the preceding proposition, but this decomposition need not be irreducible.

Consider an irreducible flow φ of dimension d with eigenvalue λ . As has been demonstrated $\varphi_{\mathbb{C}}$ is generated by $L_{\mathbb{C}}$ where L is the operator that generates φ . Since φ is irreducible, L takes the form of a single real Jordan block $\tilde{J}_d(\lambda)$ under an appropriate choice of basis. But under that basis $L_{\mathbb{C}}$ has the same form. Now if λ is real, then $L_{\mathbb{C}}$ (and consequently $\varphi_{\mathbb{C}}$) is also irreducible as complex Jordan blocks with a real eigenvalue are the same as real Jordan blocks with a real eigenvalue. However, real Jordan blocks with a nonreal eigenvalue are constructed from conjugate pairs of complex Jordan blocks, so with an appropriate change of basis $L_{\mathbb{C}} = \text{diag}(J_{\frac{d}{2}}(\lambda), J_{\frac{d}{2}}(\bar{\lambda}))$. Note that $\frac{d}{2}$ makes sense as d must be even for a real Jordan block with a nonreal eigenvalue. But this means that $\varphi_{\mathbb{C}}$ reduces to the flows generated by $J_{\frac{d}{2}}(\lambda)$ and $J_{\frac{d}{2}}(\bar{\lambda})$. Consequently, given a real flow φ , each irreducible component of φ of dimension d with eigenvalue $\lambda \in \mathbb{R}$ corresponds to an irreducible component of $\varphi_{\mathbb{C}}$ of dimension d with eigenvalue λ , while each irreducible component of φ of dimension d with conjugate pair of eigenvalues $\{\lambda, \bar{\lambda}\}$ corresponds to a pair of irreducible components of $\varphi_{\mathbb{C}}$ of dimension $\frac{d}{2}$ with eigenvalue λ and $\bar{\lambda}$ respectively.

One might wonder if complexification can be reversed; that is, given the complexification of a space, can one reconstruct the original space? Unfortunately this is not possible. The problem is that there is an infinite collection of subspaces of a complex normed space that can be viewed as embedded real normed spaces, and there is no way to tell which one is the original space. With that said, considering Definition C.1, complexification amounts to introducing multiplication by i to a real normed space. One can reverse this idea — in effect one can forget multiplication by i in a complex normed space — to construct a real normed space from a complex normed space. This process is called realification.

Definition C.11 The *realification* of X with $\mathbb{K} = \mathbb{C}$, denoted $X_{\mathbb{R}}$, is given by the set $\{x : x \in X\}$ equipped with notions of addition and (real) scalar multiplication, where addition is simply the addition of X and scalar multiplication is simply the (complex) scalar multiplication of X restricted to the real numbers.

For simplicity $X_{\mathbb{R}} = X$ for X with $\mathbb{K} = \mathbb{R}$. Similar to complexification the case where $\mathbb{K} = \mathbb{R}$ is entirely uninteresting, so it will be assumed from here on that $\mathbb{K} = \mathbb{C}$. In many respects realification is a simpler process than complexification. This can be seen in the proof that $X_{\mathbb{R}}$ is a real finite-dimensional normed space for X with $\mathbb{K} = \mathbb{C}$.

Proposition C.12 $X_{\mathbb{R}}$ is a normed space over \mathbb{R} with $\dim_{\mathbb{R}} X_{\mathbb{R}} = 2 \dim_{\mathbb{C}} X$.

Proof. That $X_{\mathbb{R}}$ is a vector space is a trivial consequence of the fact that X is a vector space and \mathbb{R} is a subfield of \mathbb{C} . It is also clear that, given any norm $\|\cdot\|$ on X , the map $\|\cdot\|_{\mathbb{R}} : X_{\mathbb{R}} \rightarrow \mathbb{R}$ given by $\|x\|_{\mathbb{R}} = \|x\|$ for all $x \in X_{\mathbb{R}}$ is a norm on $X_{\mathbb{R}}$. Let $\{b_n\}_{n=1}^d$ be any basis for X , and consider $\{\hat{b}_n\}_{n=1}^{2d}$ given by $\hat{b}_n = b_n$ whenever $n \leq d$ and $\hat{b}_n = ib_{n-d}$ whenever $d < n \leq 2d$. If $x \in X_{\mathbb{R}}$, then $x = \sum_{n=1}^d c_n b_n = \sum_{n=1}^d (a_n + i\tilde{a}_n) b_n$ for some $c_n = a_n + i\tilde{a}_n \in \mathbb{C}$. But then $x = \sum_{n=1}^d a_n b_n + \sum_{n=d+1}^{2d} i\tilde{a}_{n-d} b_{n-d} = \sum_{n=1}^{2d} r_n \hat{b}_n$ where $r_n = a_n \in \mathbb{R}$ whenever $n \leq d$ and $r_n = \tilde{a}_{n-d} \in \mathbb{R}$ whenever $d < n \leq 2d$. It follows that $\{\hat{b}_n\}_{n=1}^{2d}$ spans $X_{\mathbb{R}}$ since x was arbitrary. On the other hand, suppose $\sum_{n=1}^{2d} r_n \hat{b}_n = 0$ for some $r_n \in \mathbb{R}$, and define $\{c_n\}_{n=1}^d$ by setting $c_n = r_n + ir_{n+d} \in \mathbb{C}$. Then it must be that $\sum_{n=1}^d c_n b_n = \sum_{n=d+1}^{2d} r_n \hat{b}_n = 0$. Since $\{b_n\}_{n=1}^d$ is a basis for X , it follows that $c_n = 0$ for all $n \leq d$, and consequently $r_n = 0$ for all $n \leq 2d$. Thus $\{\hat{b}_n\}_{n=1}^{2d}$ is linearly independent, and it follows that $\{\hat{b}_n\}_{n=1}^{2d}$ is a basis for $X_{\mathbb{R}}$, so $\dim_{\mathbb{R}} X_{\mathbb{R}} = 2 \dim_{\mathbb{C}} X < \infty$. \square .

It turns out that $X_{\mathbb{R}}$ and X can be identified with each other in a very strong sense.

Proposition C.13 The map $\iota_X : X_{\mathbb{R}} \rightarrow X$ given by $\iota_X(x) = x$ is a homeomorphism with ι_X and ι_X^{-1} both \mathbb{R} -linear.

Proof. This is clear from the definition in light of the proof of Proposition C.12

One then defines the realification of a map between complex normed spaces based on this identification.

Definition C.14 Let $f : X \rightarrow Y$ with $\mathbb{K}_X = \mathbb{K}_Y = \mathbb{C}$. The *realification* of f , denoted $f_{\mathbb{R}} : X_{\mathbb{R}} \rightarrow Y_{\mathbb{R}}$, is given by $f_{\mathbb{R}}(x) = \iota_X^{-1}(f(\iota_X(x)))$. Let φ be a flow on X . The *realification* of φ , denoted $\varphi_{\mathbb{R}} : \mathbb{R} \times X_{\mathbb{R}} \rightarrow Y_{\mathbb{R}}$, is given by $\varphi_{\mathbb{R}}(t, x) = \iota_X^{-1}(\varphi(t, \iota_X(x)))$.

Previously it was shown that complexification preserves many properties of maps, and the situation is the same when it comes to realification — in fact, many of the proofs in this case are essentially trivial. For instance, $f_{\mathbb{R}}$ is continuous whenever f is continuous, as $f_{\mathbb{R}}$ is just a composition of continuous functions. The only result that requires care is the following proposition.

Proposition C.15 *Let φ be a linear flow on X . If φ is generated by $L \in \mathcal{L}(X)$, then $\varphi_{\mathbb{R}}$ is a linear flow on $X_{\mathbb{R}}$ generated by $L_{\mathbb{R}}$.*

Proof. Note that $(L_{\mathbb{R}})^k = (L^k)_{\mathbb{R}}$ for all $k \in \mathbb{N}$ by induction on k since if $(L_{\mathbb{R}})^k = (L^k)_{\mathbb{R}}$, then

$$\begin{aligned} (L_{\mathbb{R}})^{k+1}x &= (L_{\mathbb{R}})^k L_{\mathbb{R}}x \\ &= (L^k)_{\mathbb{R}} \iota_X^{-1}(L \iota_X(x)) \\ &= \iota_X^{-1}(L^k \iota_X(\iota_X^{-1}(L \iota_X(x)))) \\ &= \iota_X^{-1}(L^{k+1} \iota_X(x)) \\ &= (L^{k+1})_{\mathbb{R}}x \end{aligned}$$

for all $x \in X_{\mathbb{R}}$, so $(L_{\mathbb{R}})^{k+1} = (L^{k+1})_{\mathbb{R}}$. It follows that

$$\begin{aligned} (e^{tL})_{\mathbb{R}}x &= \iota_X^{-1}(e^{tL} \iota_X(x)) \\ &= \iota_X^{-1}\left(\lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} t^j L^j \iota_X(x)\right) \\ &= \lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} t^j \iota_X^{-1}(L^j \iota_X(x)) \\ &= \lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} t^j (L^j)_{\mathbb{R}}x \\ &= \lim_{n \rightarrow \infty} \sum_{j=0}^n \frac{1}{j!} t^j (L_{\mathbb{R}})^j x \\ &= e^{tL_{\mathbb{R}}}x \end{aligned}$$

for all $x \in X_{\mathbb{R}}$ and $t \in \mathbb{R}$. The desired result then follows. \square

The only other result that needs to be considered is how realification affects decompositions. Unsurprisingly, just as with complexification, the realification of a decomposition decomposes over the realifications of the various components.

Proposition C.16 *If $X = \bigoplus_{k=1}^m X_k$, then $X_{\mathbb{R}} = \bigoplus_{k=1}^m (X_k)_{\mathbb{R}}$. Let $L \in \mathcal{L}(X)$. If L can be written as $\bigoplus_{k=1}^m L_k$ for some $L_k \in \mathcal{L}(X_k)$, then $L_{\mathbb{R}} = \bigoplus_{k=1}^m (L_k)_{\mathbb{R}}$. Let φ be a linear flow on X . If φ can be written as $\bigoplus_{k=1}^m \varphi_k$ for some linear flows φ_k on X_k , then $\varphi_{\mathbb{R}} = \bigoplus_{k=1}^m (\varphi_k)_{\mathbb{R}}$.*

Proof. Let \mathcal{B}_k be a basis for X_k for each $k \leq m$. Then $\mathcal{B} := \bigcup_{k=1}^m \mathcal{B}_k$ is a basis for X , and using this basis to construct $X_{\mathbb{R}}$ the first statement then follows from the definition. It is

also clear using this basis that $\iota_X((X_k)_\mathbb{R}) \subseteq X_k$ and $\iota_X^{-1}(X_k) \subseteq (X_k)_\mathbb{R}$, and the other two statements follow from the definition of $L_\mathbb{R}$ and $\varphi_\mathbb{R}$. \square

As in the case of complexification, it is worthwhile to consider what happens to irreducible components under realification. As previously, this amounts to determining what happens to Jordan Blocks $J_d(\lambda)$ under realification. If φ is generated by $L = J_d(\lambda)$ with respect to a basis \mathcal{B} , then one can determine $L_\mathbb{R}$ by direct calculation of Lx for $x \in X$. If $\lambda \in \mathbb{R}$, then order the basis $\mathcal{B} \cup i\mathcal{B}$ of $X_\mathbb{R}$ by taking the elements of \mathcal{B} in order followed by the elements of $i\mathcal{B}$ again in order. With the basis in this form, direct calculation of $L_\mathbb{R}$ determines that $L_\mathbb{R} = \text{diag}(J_d(\lambda), J_d(\lambda))$ — that is, $\varphi_\mathbb{R}$ is reducible with irreducible components both generated by $J_d(\lambda)$. If $\lambda \in \mathbb{C} \setminus \mathbb{R}$, then order $\mathcal{B} \cup i\mathcal{B}$ by taking the first element of \mathcal{B} followed by the first element of $i\mathcal{B}$ followed by the second element of \mathcal{B} followed by the second element of $i\mathcal{B}$ and so on. Direct calculation of $L_\mathbb{R}$ determines that $L_\mathbb{R} = \tilde{J}_{2d}(\lambda)$, so in this case $\varphi_\mathbb{R}$ is still irreducible. Therefore, given a linear flow φ , each irreducible component of φ of size d with eigenvalue $\lambda \in \mathbb{R}$ corresponds to a pair of irreducible components of $\varphi_\mathbb{R}$ of size d with eigenvalue λ , and each irreducible component of φ of size d with eigenvalue $\lambda \in \mathbb{C} \setminus \mathbb{R}$ corresponds to an irreducible component of $\varphi_\mathbb{R}$ of size $2d$ with conjugate eigenvalue pair $\{\lambda, \bar{\lambda}\}$.