UNIVERSITY OF CALIFORNIA
SANTA CRUZ

ON AN EXTENSION OF THE MEAN INDEX TO THE LAGRANGIAN GRASSMANNIAN

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

MATHEMATICS

by

Matthew I. Grace

August 2020

The Dissertation of Matthew I. Grace
is approved:

Professor Viktor Ginzburg, Chair

Professor Başak Gürel

Professor Richard Montgomery

Dean Quentin Williams
Interim Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

Abstract

On An Extension of the Mean Index to the Lagrangian Grassmannian

by

Matthew I. Grace

For a symplectic vector space (V, ω), recall the identification of $\mathrm{Sp}(V, \omega)$ with an open and dense subset $\mathrm{Im}(\mathrm{Gr}) \subset \Lambda_{2n} := \mathrm{LagGr}(V \times V, \mathrm{Pr}_1^* \omega - \mathrm{Pr}_2^* \omega)$ of the Lagrangian Grassmannian, where $\mathrm{Gr}(A)$ denotes the graph of any $A \in \mathrm{Sp}(V, \omega)$. Our central result is in extending the mean index via this embedding from the continuous paths in $\mathrm{Sp}(V, \omega) \cong \mathrm{Im}(\mathrm{Gr})$ to those contained in a subset $\mathcal{L}_{2n} \subset \Lambda_{2n}$ with $\mathrm{codim}(\Lambda_{2n} \setminus \mathcal{L}_{2n}) = 2$. Then, using a composition operation on Λ_{2n} inherited from the theory of linear relations which extends the usual group structure of $\mathrm{Sp}(V, \omega)$ to that of a monoid on $\mathcal{L}_{2n}$, we show that the extension retains some of the desirable properties of the (symplectic) mean index if restricted to an open and dense subset $\mathcal{P}_{reg}(\mathcal{L}_{2n}) \subset C^1([0, 1], \mathcal{L}_{2n})$ of 'stratum-regular' paths, which we equip with an equivalence relation $\sim_{comp}$. We show that the point-wise composition of any two $\sim_{comp}$ equivalent stratum-regular paths is piece-wise differentiable and prove that the extended index, when restricted to $\mathcal{P}_{reg}(\mathcal{L}_{2n})$, is homogeneous (for non-negative integers) and satisfies a quasimorphism-type bound for any $\sim_{comp}$ equivalent pair of paths.

Mais qu'est-ce que ça veut dire, la peste?

C'est la vie, et voilà tout.

Albert Camus

Part I

Preliminaries

I.1 Introduction

Throughout this dissertation (V, ω) will be a real symplectic vector space of dimension $2n$ and $\Lambda_{2n} := \mathrm{LagGr}(V \times V, \mathrm{Pr}_1^* \omega - \mathrm{Pr}_2^* \omega)$ will denote[1] the Lagrangian Grassmannian of the twisted symplectic product of (V, ω). Our (Lagrangian) mean index $\hat{\Delta}$ extends the classical mean index Δ from the set of continuous paths residing in the symplectic group $\gamma\colon [0,1] \to \mathrm{Sp}(V)$, to those continuous paths in Λ_{2n} whose images reside in the set of 'admissible' Lagrangian subspaces (see definition I.3.4), an open and dense subset $\mathcal{L}_{2n} \subset \Lambda_{2n}$. The extension of Δ is achieved through the map $\mathrm{Gr}\colon \mathrm{Sp}(V) \to \Lambda_{2n}$ (see definition I.3.3) which identifies $\mathrm{Sp}(V)$ in Λ_{2n} with the open and dense embedded submanifold

$$\mathrm{Im}(\mathrm{Gr}) = \{L \in \Lambda_{2n} \,|\, \exists A \in \mathrm{Sp}(V),\, \mathrm{Gr}(A) = L\} \cong \mathrm{Sp}(V),$$

contained within the set of admissible Lagrangians $\mathcal{L}_{2n}$. We denote the complement of the admissible Lagrangians by $H := \Lambda_{2n} \setminus \mathcal{L}_{2n}$ and will occasionally refer to the elements within as 'exceptional' Lagrangian subspaces. From first inspection the set of admissible Lagrangians $\mathcal{L}_{2n}$ may appear a fairly insignificant addition to $\mathrm{Im}(\mathrm{Gr}) \cong \mathrm{Sp}(V)$, as both $\mathrm{Im}(\mathrm{Gr})$ and $\mathcal{L}_{2n}$ are open and dense in Λ_{2n}, yet their difference becomes more apparent when we observe the complements of each; the complement $\Lambda_{2n} \setminus \mathrm{Im}(\mathrm{Gr})$ is a hypersurface in Λ_{2n} (see proposition II.1.6) whereas we show in theorem I.4.1 that $\mathrm{codim}(H) = 2$.

[1] Since the entirety of the dissertation is concerning a fixed symplectic vector space we will suppress the symplectic form, or the vector space altogether, in our notation (as seen in the shorthand, $\mathrm{Sp}(V)$ and Λ_{2n}).

The mean index is commonly seen in one of two contexts: as a real valued map over all continuous paths[2] in $\mathrm{Sp}(V)$ or, as it is found in the bulk of the literature (e.g. [4, 78, 74, 57]), restricted to those paths γ originating at the identity (i.e. $\gamma(0) = \mathrm{Id}$, we will often refer to such paths as 'identity-based'). Besides the fact that most of the algebraic properties of Δ only exists in the context of the latter definition, it is particularly useful due to the isomorphism between fixed-endpoint homotopy classes of paths originating at some base point in $\mathrm{Sp}(V)$ (the identity in this case) and the universal cover of the symplectic group $\widetilde{\mathrm{Sp}}(V)$; the natural space in which to model linearizations of a Hamiltonian flow. Since this identification, given below in lemma I.3.8, is purely topological in nature it is apparent that the preference for a fixed base-point in the literature (at least that related to symplectic geometry) likely has more to do with the Lie group structure[3] on $\widetilde{\mathrm{Sp}}(V)$, over which the mean index Δ may be characterized (as in [4]) axiomatically as the unique quasimorphism on $\widetilde{\mathrm{Sp}}(V)$ which satisfies certain conditions (see definition I.2.2 and lemma I.3.11). This thesis will adopt the former notion, motivated in part by the loss of group structure in passing to the extended domain's universal cover $\widetilde{\mathcal{L}}_{2n}$ (more discussion on this may be found in remark I.3.13, with the necessary context beginning at definition I.3.7). Of course since the latter definition is simply the former restricted to certain paths, this decision does not preclude later restricting $\hat{\Delta}$ for the purposes of identifying certain

[2]Continuity of paths will be an assumption implicitly maintained throughout the dissertation, if a path is discontinuous it will be specified as such. Differentiability will also be explicitly invoked throughout, save for those instances in which doing so would be unnecessarily redundant.

[3]The group structure of $\widetilde{\mathrm{Sp}}(V)$ is fundamentally related to the mean index, see lemma I.3.9.

algebraic properties of Δ and adapting them, with some alterations, to the extension $\hat{\Delta}$.

As mentioned above, one may define Δ as a real-valued map defined over the homotopy classes of paths as in definition I.3.14 or equivalently[4], as a real valued map on the universal cover $\widetilde{\mathrm{Sp}}(V)$ (given in definition I.3.10). Our interpretation of the mean index may be established using either approach as the technique used in both cases (generally credited to Milnor [65]) to construct Δ involves one of two similar lifting procedures in terms of a certain continuous map $\rho : \mathrm{Sp}(V) \to S^1$, which is defined axiomatically for certain Lie groups in definition I.3.7 (in this context it is often called a circle map) and given explicitly for the symplectic group in definition IV.1.1. A consequence of this is that both the existence and continuity of Δ amount to formal consequences of the continuity of ρ. This allows us to avoid extending the function Δ (defined over the free path space of $\mathrm{Sp}(V)$ in greatest generality) and instead deal with the far more straightforward question of extending ρ (in fact we will need to extend ρ^2, see remark I.3.15) to $\mathcal{L}_{2n}$.

The result at the center of the dissertation is theorem I.4.2 (in which we define the extended mean index and show it to be continuous) which is an immediate corollary of theorem V.3.1 through an application of the aforementioned lifting procedure. Specifically, theorem V.3.1 produces a continuous extension $\hat{\rho} : \mathcal{L}_{2n} \to S^1$ of ρ^2 by identifying $\mathrm{Sp}(V) \cong \mathrm{Im}(\mathrm{Gr})$ so that by simply repeating Milnor's procedure on

[4]These two definitions remain equivalent when applied to $\mathcal{L}_{2n}$, see lemma I.3.8.

$\mathcal{L}_{2n}$ with respect to $\hat{\rho}$, we obtain both the existence and continuity of $\hat{\Delta}$ (as in the symplectic case) as formal consequences of the continuity of $\hat{\rho}$. We note that the universal cover of $\mathcal{L}_{2n}$ (and therefore the space of homotopy classes of paths originating at the identity) is likely to be much larger[5] than that of Λ_{2n} (discussed further in remark I.3.13). This, along with the loss of group structure when passing from $\mathrm{Sp}(V)$ to $\mathcal{L}_{2n}$, contributes to our extended index exhibiting some fairly serious algebraic deficiencies when freshly constructed as compared to Δ over $\widetilde{\mathrm{Sp}}(2n)$. Given that the Lie group structure and resulting characterization of Δ as a quasimorphism are used in almost all of the applications of the mean index (a few examples may be found in [4, 27, 29, 10]) we desire to recover some of these properties, but before discussing this we address certain similarities between this dissertation and the thesis [32].

Several of the key objectives of this dissertation, including the central concept of extending the mean index using definition I.3.3 (as well as the role of the exceptional set H in proving it) were directly motivated by the results given in the dissertation [32], authored by Yusuf Gören. In particular, our theorem V.1.1 matches theorem **2.2.6** in [32], though are proved using distinct methods. Our theorem III.2.5 is a slight refinement (in that it shows uniqueness) of Gören's theorem **2.2.1**, with proofs that do bear some similarities where they overlap. The most significant overlap occurs with our main theorem I.4.2, which corresponds to theorem **2.2.5** in [32].

Given the parallels outlined above, this dissertation necessarily exhibits significant

[5] This is shown explicitly in example VII.1.2, wherein $\pi_1(\mathcal{L}_2)$ is computed and an ad hoc solution is shown in which we identify an intermediate covering space in which lifts of the non-contractible loops with zero mean index remain non-contractible in the cover.

distinguishing features, the first of which is our choice of definition when constructing $\hat{\Delta}$, central in the recovery of the algebraic properties found in section VI. We have already mentioned how we will construct $\hat{\Delta}$ (formally in terms of a continuous extension $\hat{\rho}$ of the map ρ), whereas in [32] this is circumvented rather cleverly by choosing a representative $\gamma : I \to \mathcal{L}_{2n}$ for each identity-based fixed-endpoint homotopy class $[\gamma]$ such that γ may be written as the concatenation of a pair of paths τ and η. Specifically, the author requires that $\tau(t) \in \mathcal{L}_{2n}$ be an identity-based path which is the graph of a symplectic map for all $t \in [0, 1)$ with $\tau(1) = L \in \mathcal{L}_{2n} \setminus \mathrm{Im}(Gr)$, while η is some loop based at $\tau(1)$ such that $[\eta * \tau] = [\gamma]$, thereby reducing the extension problem to only those paths which leave Im(Gr) at the very last moment. This means that our continuous extension of ρ^2 in theorem V.3.1 should be considered the true 'main result' of this dissertation as it marks a significant departure from Gören's methods in proving theorem I.4.2. One immediate benefit of this definition (following that given in [4]) is that it better suits establishing variants of the algebraic properties enjoyed by the mean index (over the open and dense subset of 'stratum-regular'[6] paths, see proposition II.4.13), defined in the natural manner by replacing the group operation with Lagrangian composition (see definition I.3.2). These begin with the necessary prerequisite that the composite path of any 'compatible' pair (see definition VI.1.1) of stratum-regular paths is piece-wise differentiable and include proofs that the Lagrangian mean index is both homogeneous and satisfies a quasimorphism-type bound

[6]In this case a stratum-regular path is one satisfying a transversality condition (see definition I.4.4) with respect to a certain stratification of Λ_{2n} as given in [46].

for the aforementioned compatible pairs. Despite this, the similar goals that [32] and this dissertation share ultimately proved invaluable in the writing of this thesis, evident in the fact that the hypothesis of this work's central theorem may be credited to Gören's dissertation.

A more precise description of the results mentioned above, namely those regarding properties of the symplectic mean index partially retained in the Lagrangian mean index, begins with the immediate formal consequence of fixed-endpoint homotopy invariance (due to the definition). More importantly, we prove for each stratum-regular path $\gamma : I \to \mathcal{L}_{2n}$ that the set-theoretic composition γ^l is a piece-wise differentiable path in $\mathcal{L}_{2n}$ for all $l \in \mathbb{N}$ (a consequence of theorem I.4.5 mentioned below) which allows us to show in corollary I.4.6 that the mean index is homogeneous with respect to this composition when $l \geq 0$ (we lose the negative integers here due to complications which arise in defining the notion of an inverse within the category of linear canonical relations, see remark II.2.3). We next give an equivalence relation on the set of stratum-regular paths and see in theorem I.4.5 that all compatible paths γ, τ sharing an equivalence class have piece-wise differentiable composite paths $\gamma \circ \tau, \tau \circ \gamma$. The finale comes in section VI.2 with theorem I.4.7 in showing that any such pair satisfies a quasimorphism-type bound (namely, the inequality in definition I.2.2) with respect to $\hat{\Delta}$. A word of warning should be said concerning the instability this equivalence relation exhibits under reparameterization; given any non-identity C^1 map β from the unit interval to itself which fixes endpoints, the resulting reparameterized path $\gamma \circ \beta$

will in general not be a γ-compatible path regardless of how close β is to the identity (unless one only considers very restricted families of reparameterizations which would necessarily depend on each equivalence class). Regardless, as unstable as these results may be under reparameterization, they are evidence that certain algebraic properties of Δ have at least partial analogues in the Lagrangian case.

We make use of the stratification detailed in section II.1 of the Lagrangian Grassmannian of a $4n$ dimensional symplectic vector space as in [39, 46]. These $n+1$ strata, parameterized by $0 \leq k \leq n$, each form a fiber bundle over the space of isotropic pairs of dimension k (see definition II.1.1) and are shown in proposition II.1.6 to have codimension k^2 in Λ_{2n}. The purpose of distinguishing the stratum-regular γ in the manner they are is that such paths induce a finite partition $\{t_i\}_{i=1}^{M}$ of the unit interval which mark each departure from the generic stratum (the image of $\mathrm{Sp}(2n)$ under the graph map given in definition I.3.3). This allows one to write γ as the concatenation of a finite number of (open) paths, each lying in the image of the graph map (and therefore the symplectic group). This decomposition (shown in lemma VI.1.5) is purely technical, the utility lies in the fact that each restriction of γ may be identified with a symplectic path, thereby providing a gateway through which the algebraic properties (lost in extending Δ to the Lagrangian Grassmannian) may be recovered. We obtain the above results by using the index theory summarized below in section I.3 on each restriction and afterward show that the collection of paths may be continuously stitched back together with the given property intact.

The structure of this dissertation roughly parallels three central theorems which each contribute to the proof of theorem I.4.2, contained in parts II,III and V (for the precise theorem statements we refer the reader to section I.4). Part II introduces isotropic pairs and the fiber structure of the strata of Λ_{2n}, in addition to elaborating on the set of stratum-regular paths. In part III we prove theorem I.4.1 which establishes that $\operatorname{codim}(H) = 2$, as well as prove theorem III.2.5 which gives a standard procedure for decomposing $L \in \mathcal{L}_{2n}$. Following this, part IV consists of mostly formal results regarding the construction of the map ρ and its relationship with the Conley-Zehnder index (in addition to a brief discussion regarding which properties are retained by $\hat{\rho}$) while part V involves some of the heavier technical details needed to prove our central theorem V.3.1, which yields theorem I.4.2 as a formal consequence (though we show this explicitly in section V.3). The penultimate part VI is reserved for the proofs of the algebraic results for stratum-regular paths while part VII closes the dissertation by giving a few examples and speculative remarks.

Acknowledgements. *The author would like to thank his advisor Viktor Ginzburg for his unwavering patience, understanding and helpful feedback throughout the last four years.*

I.2 Historical Context

In regards to the developments discussed below (and most ideas in general), coming to a definitive and complete answer as to when certain steps were reached, who did so and with which references is almost certainly an impossible task. Along similar lines it must also be made explicit that the following account of symplectic and Lagrangian index theory is necessarily incomplete and subject to major omissions and misconceptions. This is partially due to the complexity inherent in any thorough account of a mathematical topic which is unlikely to be known by any one person in its entirety, nor capable of being thoroughly expressed in so few pages. Philosophy aside, this deficiency is likely a simple consequence of the author's naivetè on the matter so that many vital contributions, and therefore contributors, to the field are likely to be missing from the following account.

Now that we have completed our disclaimer, before we proceed in outlining some of the advances in Morse-type index theory (and the particular branch of subtopics grown out of the Maslov index theory) we will briefly discuss a central concept which has driven the advances in Morse-type index theory since it began, that of the calculus of variations. The notion of a variational principle is generally understood to have originated in the works of Maupertuis and Euler during the mid 18^{th} century [48] and has consistently appeared in various physical and mathematical contexts through to the modern day. In the first half of the 19^{th} century following the work of Lagrange's reformulation of Newton's laws of motion, William Hamilton established what is now known as Hamilton's principle [37], leading to his own formulation of classical mechanics which ushered in the field of Hamiltonian dynamics. Despite this, the intuition behind variational principles has proven surprisingly ancient, with written examples dating back roughly two millennia. Included among these is Archimedes' *Law of the Lever* [81], now a standard introductory example for the concept of virtual work[7], as well as a Hellenic treatise on optics[8] [48]. These examples are almost certainly predated by works lost to time or by ideas that were never written down, evidenced by the appearance of a variational principle being used to solve an isoperimetric area maximization problem within a Phoenician myth [48] by making use of the free boundary of the sea shore; an ancient foreshadowing of the near ubiquitous introductory calculus exercises involving the construction of rectangular animal pens alongside rivers.

More recently in the first half of the 20^{th} century the establishment of Morse theory [66] ushered in the field deemed by Morse as *variational calculus at large* (encompassing topics such as Lusternik-Schnirelmann theory and later motivating Bott's periodicity theorems). As defined by Morse, *variational calculus at large* is the study of the qualitative behavior of variational problems and their relationship with global topological properties, demonstrating that the utility of variational methods extend beyond that of simply identifying local extrema. As seen in Morse's work on function

[7]The principle of virtual work has seen modern uses in applying symplectic techniques to physics, e.g. the generalized Hamiltonian dynamics in elastic continuum mechanics as given in [7].

[8]Heron's Catoptrica (*Theory of Mirrors*) argues that the trajectory of light determined by Euclid's law of reflection is the shortest (reflected) path possible from the source to the observer.

spaces [66], early Morse theory was generally confined to finite dimensional manifolds, though this was often used in combination with various tricks to extend the theory to infinite dimensional manifolds like loop spaces or spaces of geodesics. This was later adapted to include well behaved functionals on separable Banach and Hilbert spaces [14], albeit with the possibility of infinite Morse (co)indices. Some of the main problems encountered during this time period were those of strongly indefinite functionals (A strongly indefinite functional is unbounded in both directions and retains this property modulo any finite dimensional subspace) and infinite (co)indices (both of which occur when applying variational techniques to most Hamiltonian systems). This suggests that any Morse-type results under these conditions (and on Hamiltonian systems in particular) will require more than purely topological data. One example motivating this is given in [41], in which the fact that the topology of many infinite dimensional spaces tends to be excessively fine to the effect that the topological invariance of (co)homology tends to force the (co)homology groups to be trivial. This obviously should not be considered a commandment given the existence of work such as Palais' extension of Lusternik-Schnirelmann theory to Banach manifolds [70], although admittedly Finsler structures are used to do so. We mention these examples with the express purpose of exhibiting the difficulties encountered when applying variational techniques to Hamiltonian systems.

The prospect of using variational methods established before the late 1970s to identify the periodic orbits of all but the nicest Hamiltonian systems was an intractable one for several reasons (including but certainly not limited to the few listed above). The 1978 paper [72] authored by Rabinowitz is occasionally marked as the beginning of a decades-long refinement in using variational techniques to study Hamiltonian systems, a refinement which is far from complete and which continues to this day. When seeking periodic orbits of a Lagrangian dynamical system Rabinowitz had the novel realization that, in some cases, one may eschew classical Lagrangian methods over the phase space in favor of a Lagrangian system defined on the loop space of the manifold, thereby identifying periodic orbits (modulo parameterization) by seeking the extrema of the new action functional[9]. The main obstacle Rabinowitz and his contemporaries encountered was that the action functional could take different values for the various parameterizations of geometrically indistinguishable orbits, e.g. iterates. Two major breakthroughs (among many others) during the 1980s would help to revolutionize the field of Hamiltonian dynamics (in addition to countless other topics in symplectic topology) and most relevant to our discussion, would contribute to certain methods for identifying the distinct geometrically indistinguishable orbits that Rabinowitz encountered.

The first of these prominent steps would come in 1984 after Conley and Zehnder authored the paper [13], commonly cited for introducing what is now known as the Conley-Zehnder index. The authors define a Morse-type index of a (non-degenerate) periodic orbit for certain linear Hamiltonian systems on $(\mathbb{R}^{2n}, \omega_{std})$ when

[9]This description is a paraphrasing of Helmut Hofer during his 60^{th} birthday conference. In the video he is speaking of Rabinowitz's 1978 paper and it's influence on his decision to specialize in symplectic/contact geometry [68].

$n \geq 2$, provided a number of conditions are satisfied (including the non-degeneracy of the trivial orbit and an asymptotic linearity condition on the Hamiltonian). In other words, when $n \geq 2$ the index assigns an integer to every identity-based path $\gamma : [0,1] \to \mathrm{Sp}(2n)$ arising from a linearized Hamiltonian orbit which terminates at a non-degenerate symplectic map. This allowed a prime periodic orbit to be distinguished from its iterates, thereby solving Rabinowitz's problem, albeit only for certain linear Hamiltonian systems (at the time leaving much of the behavior of Hamiltonian systems on symplectic manifolds an open question). In the second half of the 1980's, Andreas Floer would begin developing several seminal contributions to the field of symplectic dynamics which would later influence a remarkable series of results spanning nearly every field of symplectic mathematics (and well beyond). In the interest of brevity, Floer developed several novel infinite dimensional Morse theories [22, 23] which built on his earlier results [20, 21] regarding the symplectic action functional and its associated gradient flow. The ultimate consequence was a novel Morse-type homology, what is now known as Floer homology[10]. This work would quickly rise in prominence and ultimately motivate many of the rapid advances in symplectic mathematics which proceeded it. In particular, Floer's contributions would result in the rapid growth of both the generalization and application of Conley and Zehnder's results including an early example by Floer [19] and continued on by other authors

[10]The descriptor 'Floer homology' is rather vague, even if one ignores analogues developed later. Floer's work exhibited homology theories associated to Lagrangian intersections [22] as well as non-degenerate symplectic endomorphisms [20, 23] (and this is entirely omitting his work on three-manifolds).

in the papers [54, 82, 55], although it should be said that these few citations naturally only account for an incredibly small portion of the relevant contributions and contributors.

Two decades before the Conley-Zehnder index and Floer's breakthrough, the Maslov index for paths of Lagrangian subspaces provided an important motivation for developing an index theory for $\mathrm{Sp}(2n)$. The Maslov index as defined by Arnol'd [2] is a characteristic class for Lagrangian submanifolds. More concretely, for any Lagrangian subspace $L \subset (M, \omega)$ the Maslov index of a (generic) path $\gamma : [0, 1] \to L$ may be written as a signed intersection count of the associated path $\gamma_T : [0, 1] \to \Lambda_n$ (induced by the tangent map) with a co-oriented hypersurface of Λ_n whose cohomology class coincides with the Poincaré dual of the Maslov class (e.g. the set of Lagrangian subspaces non-transversal to some fixed Lagrangian, often called a Maslov cycle)[11]. After a few years much of the material defined in [2] would become relatively standard after a chapter authored by Arnol'd was featured in the textbook [63]. The manner in which the Maslov index was formed would foreshadow the methods later used in [78] to define the Conley-Zehnder index.

In slightly more detail, both the Maslov and Conley-Zehnder indices may be constructed starting from the determinant map on the unitary group $U(n)$ to produce a continuous S^1-valued function ρ on the relevant space (Λ_n and $\mathrm{Sp}(2n)$ respectively),

[11]The idea of using a signed intersection count and some co-oriented hypersurface in $\mathrm{Sp}(2n)$ (as in [78]) to define an index for linear-symplectic paths was nearly contemporaneous, as can be seen in figure I.1 which depicts the analogous 'Maslov cycle' in $\mathrm{Sp}(2)$, consisting of those symplectomorphism with an eigenvalue equal to 1.

each of which inducing an isomorphism of the relevant fundamental group[12] with $\pi_1(S^1) \cong \mathbb{Z}$ (see definition I.3.7). More precisely, in the Lagrangian case [2] Arnol'd uses the square of the determinant modulo its kernel, identifying the domain through the diffeomorphism $\Lambda_n \cong U(n)/O(n)$ to produce the map ρ whereas Salamon and Zehnder's paper [78] uses a method depending on the eigenvalues of a symplectic map, which has no analogous method in terms of the Maslov index on Λ_n (though they do recognize that ρ corresponds with the square of the determinant map on $U(n)$ as a subgroup of $\mathrm{Sp}(2n)$).

From a newcomer's perspective, it's commonly understood that equivalent definitions for the Conley-Zehnder index (and in turn, the mean index) which closely follow Arnol'd's methods in obtaining ρ on Λ_n would only appear many years after Arnol'd's publications; the first in 1984 [13] already briefly mentioned, included a procedure using the squared determinant and the usual polar decomposition to define ρ in terms of the determinant on $U(n)$. The second definition came in 1989 in the paper [75] in which ρ is given as the normalized squared determinant of the holomorphic part of a linear symplectomorphism (seen as an element of $GL(n,\mathbb{C})$). We will discuss those in their appropriate temporal context shortly, but first we must take a step back to trace the roots of the Conley-Zehnder index, as it should be noted that Conley and Zehnder's 1984 paper [13] does not mark the first appearance of the mean and Conley-Zehnder indices. The mean index is foreshadowed as far back as

[12]In our case though, the map $\hat{\rho}$ will fail in satisfying this as the induced map $\hat{\rho}^* : \pi_1(\mathcal{L}_{2n}) \to \pi_1(S^1)$ is highly non-injective, see remark I.3.13).

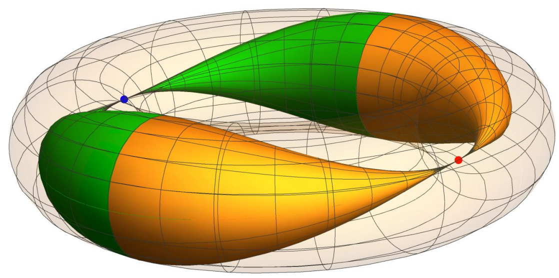

© Mathematica 10

Figure I.1: The above is a depiction of $\mathrm{Sp}(2,\mathbb{R})$ as a solid (open) torus wherein the surface of symplectomorphisms with eigenvalues $+1,-1$ are drawn in orange/green respectively. The figure is of interest as an analogous figure (not shown here to avoid violating copyright law) appeared in the 1955 paper [25] published by Gel'fand and Lidskii, in which the authors exhibit (in the $n=1$ case) the co-oriented (hyper)surface $Sp^*(2n,\mathbb{R})$ consisting of the symplectomorphisms with eigenvalue equal to 1. This appearance notably precedes (by several decades) the procedure given in [78] for defining the Conley-Zehnder index in terms of a signed intersection count (à la the Maslov class by Arnol'd [2]).

1885 in Poincaré's rotation number [71] later to be associated with Morse theory by Hedlund [39] in 1932, a paper which notably features the two dimensional case of inequality (I.2.3). The more well known precursor was published by Bott in 1956 [8], in which he defines and computes analogues of the mean index (via equation (I.2.1) below), the Conley-Zehnder index as well as a nullity index for iterates of closed geodesics.

Along these lines we highlight a textbook and paper (see figure I.1) which appear as citations (the former more often than the latter) within many of the early developments mentioned below ([16, 15, 78, 55] to name a few) which shows significant developments in the application of the above index theory to Hamiltonian systems in the 1980s nearly 30 years beforehand. One particularly surprising example is the alternate construction for the Conley-Zehnder index in [75] mentioned above, which appears in Yakubovich and Starzhinskii's text [87]. Impressively, the crucial step of concatenating a symplectic path so as to obtain an integer from the lift of ρ had also already been given in the 1975 textbook, citing a 1955 paper published by Gel'fand and Lidskii [25]. This pair of documents mark early examples of a significant portion of the fundamentals of Hamiltonian index theory, including the stability of periodic orbits of Hamiltonian systems, in turn giving a (very intuitive) description of the dynamics of eigenvalue quadruples belonging to a real family of symplectic linearizations of a Hamiltonian orbit.

Jumping ahead again to 1983, despite the paper [13] still being a work in progress (or perhaps a preprint), Conley and Zehnder hastily put their index to use proving Arnol'd's conjecture for tori in [12]. The techniques used within would spark a cascade of new research, soon to become foundational topics in symplectic geometry. One example is Eliashberg's proof [17] that the image $\mathrm{Symp}(M, \omega) \hookrightarrow \mathrm{Diff}(M)$ is C^0-closed for symplectic (M, ω), thereby establishing symplectic topology as a topic in its own right. Soon after, Gromov used traditional techniques from enumerative algebraic geometry to identify symplectic topological invariants through the analysis of pseudoholomorphic curves [33], marking their origin within symplectic and contact geometry. The paper also established the first symplectic non-squeezing results giving the first example (what is known today as the Gromov width) of a symplectic capacity, a topic which has since contributed to many non-intuitive results in the field of symplectic embeddings [64, 44, 45, 79].

Two papers of note which utilized Morse theory for Hamiltonian systems (pre-Floer homology) were published by Ekeland [16], and Ekeland and Hofer [15], in 1984 and 1987 respectively. The former paper is of particular interest for our purposes since it marks an early appearance of the mean index for Hamiltonian systems, defined as in equation (I.2.1), for which inequality (I.2.3) is shown. Additionally the $\{1, 2, \ldots, 2n\}$-valued nullity index is defined, foreshadowing the role it would play in the paper [55] in extending the Conley-Zehnder index to degenerate maps. Both of the papers naturally assumed fairly restrictive conditions on the Hamiltonian yet they mark the first signs of the following decade's flood of results, of which Floer's contributions at the end of the 1980's could be marked as when the flood passed a critical threshold (of the contribution made by Floer, the series of papers [18, 19, 20, 21, 22, 23, 24] represents only a portion of the relevant papers). We note that Floer was likely motivated in part by the introduction of Conley and Zehnder's index theory (as evidenced by [19]) and undoubtedly by Gromov's introduction of pseudoholomorphic curves (see the paper [18] published in vol. 25 of *Travaux en Cours*, or *Works in Progress*). Floer's contributions have since proven to have been very effective kindling for the following years' advances in symplectic topology, as fueled by this development much progress was made in refining the relationship between Conley-Zehnder index theory and the various symplectic homology theories. In particular, Hamiltonian Floer Homology yielded many novel results in Hamiltonian dynamics [76, 9], being just one of many fields during this time which experienced rapid growth through the application of various Floer-type homology theories established in the early 90's.

Returning again to the various constructions of the Conley-Zehnder index (and in most cases, the mean index as well), we recall that in [13] the map ρ (and in turn the Conley-Zehnder index) is defined for certain linear Hamiltonian systems on $\mathbb{R}^{2n}$ with $n \geq 2$ using a continuous map given by some polar projection sending the non-degenerate linear symplectic maps to $U(n)$, followed by the application of the squared determinant. In 1990 a pair of papers [58, 54] were published each of which were (co)authored by Yiming Long, following four years of his work in analyzing the peculiar dynamics of certain forced Hamiltonian systems (see [52, 53]). In the first, published with Zehnder, they defined the index for non-degenerate linear Hamiltonian systems in two dimensions [58] filling in the gap left in [13], while the second paper finally extended the index to degenerate fundamental solutions of linear Hamiltonian systems (provided they have symmetric and continuous coefficients) [54]. Remarkably, a nearly identical result in [82] was published that same year by Viterbo. In 1992, Salamon and Zehnder published their hallmark paper [78] which moved the theory beyond linear system, extending the Conley-Zehnder index to periodic orbits of non-degenerate Hamiltonian systems on any compact symplectic manifold of dimension $2n$. They also established an axiomatic formulation[13] of the Conley-Zehnder index

[13] An analogous list of axioms for the mean index was also published by Barge and Ghys in the same paper [4], although our index unfortunately does not satisfy these axioms (see remark I.3.13).

under which the index is unique, which coincidentally was also shown the same year by Barge and Ghys in [4].

A paper [74] authored by Robbin and Salamon in 1993 used the Maslov index and the graph map in definition I.3.3 to construct a novel index for linear symplectic paths (generally called the Robbin-Salamon index today), in addition to giving a pair of alternate definitions for the Conley-Zehnder index. These included writing the index as a signed intersection count with a co-oriented hypersurface (as in Maslov's initial work) in addition to embedding $\mathrm{Sp}(2n)$ into $GL_n(\mathbb{C})$ (a technique mentioned above for appearing nearly two decades earlier in the text [87]). Much of the material covered until this point (and much more) would later be compiled and published in 1999 as notes [77] from a lecture delivered by Salamon in 1997, in a sense standardizing much of the aforementioned topics. The field was growing quickly though as in just two years those notes showed their age, as 1997 also marks when Long extended the Conley-Zehnder index [55] to every path in $\mathrm{Sp}(2n)$ (which had only been shown in dimension two by Long and Zehnder until that point). The two papers [78, 55] and their offspring consequently opened the door to a new avenue of research regarding degenerate periodic Hamiltonian orbits, or in many cases simply the non-degenerate orbits of a degenerate Hamiltonian system on a compact symplectic manifold. Since then many advances have been made in answering the early conjectures of the field; Floer established Arnol‘d's conjecture for closed symplectically aspherical manifolds in [24] followed two decades later by the Conley conjecture (under identical assumption), which was shown to hold in [28, 30] (in addition to several alternate hypotheses).

Given the amount of time spent addressing the Conley-Zehnder index, it is appropriate to describe how the mean and Conley-Zehnder indices are related beyond their method of construction, with perhaps the most obvious example being that the former can be expressed as a 'weighted average' of the latter (see equation (I.2.1) below). As mentioned above, both indices have corresponding axiomatic definitions in terms of the group structure on $\widetilde{Sp}(2n)$ (see [4, 5] for the former and [78, 4] for the latter) but in our case the desired extension's domain $\widetilde{\mathcal{L}}_{2n}$ (the universal cover of the admissible Lagrangians $\mathcal{L}_{2n}$) is only a monoid, which limits us to a definition which does not presume the existence of a group structure. Fortunately such definitions exist, like that used in [78, 4] (among many others), wherein the mean index is defined using the same map $\rho : \mathrm{Sp}(2n) \to S^1$ used to construct the Conley-Zehnder index. There is another definition of the mean index in terms of the Conley-Zehnder index, albeit it too will be of little use to us precisely because of its dependence on the Conley-Zehnder index. Regardless this should not preclude the mention of it, particularly since it exhibits the fundamental relationship between the two indices described at the beginning of the paragraph[14]. In this definition the mean index is given as the continuous real valued map Δ defined over identity-based fixed-endpoint homotopy classes of paths $\gamma : [0, 1] \to \mathrm{Sp}(2n)$ which satisfies the following equation

[14]In fact, the etymology of the term 'mean' index is a reference to this 'weighted average', expressed in equation I.2.1.

for all such γ,

$$\Delta(\gamma) = \lim_{k\to\infty} \frac{\mu_{cz}(\gamma^k)}{k}. \tag{I.2.1}$$

Remark I.2.1. *Here μ_{cz} denotes the Conley-Zehnder index. We refer the reader to lemmas I.3.8 and I.3.9 for an explicit definition of the iterate γ^k of an identity based path (not necessarily a loop) γ, and more generally how the composition of any two identity based paths in* $\mathrm{Sp}(2n)$ *is defined.*

Fortunately the two indices' relationship extends beyond the above definition (or equation), namely the two are governed by the following inequality which holds for both degenerate and non-degenerate paths in $\mathrm{Sp}(2n)$ originating at the identity

$$|\Delta(\gamma) - \mu_{cz}(\gamma)| \leq n, \tag{I.2.2}$$

a generalization of an earlier development for non-degenerate γ

$$|\Delta(\gamma) - \mu_{cz}(\gamma)| < n, \tag{I.2.3}$$

which may be found in [16, 15, 78] holding for whichever 'admissible' paths γ each paper is considering. In 1997, Yiming Long established (I.2.2) for all paths [55], constructing a $\mathbb{C}\backslash\{0\}$ family of Maslov-type indices, $(\mu_\omega, \eta_\omega) \in \mathbb{Z}\times\{0, 1, \ldots, 2n\}$, defined for all (including ω-degenerate) paths in $\mathrm{Sp}(2n)$ such that each ω-index has a corresponding 'ω-mean index'. In particular when $\omega = 1$ this Morse-type index coincides with the classic μ_{cz} thereby forming a 'generalized' Conley-Zehnder index in that all paths γ with non-degenerate endpoints satisfy $(\mu_{\omega=1}(\gamma), \eta_{\omega=1}(\gamma)) = (\mu_{cz}(\gamma), 0)$. It is also shown by Long that equation (I.2.2) holds for all paths with equality only if the path is degenerate. That same year, a stronger inequality (centered about μ_{cz} with asymmetric upper/lower bounds) was given and shown to be optimal, varying with the nullity index of a given iterate [51].

In addition to satisfying (I.2.2), the mean index possesses other algebraic properties, including homogeneity and the following quasimorphism property.

Definition I.2.2. *Given a group G, a map $\Delta : G \to \mathbb{R}$ is called a quasimorphism if there exists some $c \in \mathbb{R}$ for which all $\phi, \theta \in G$ satisfy the following inequality,*

$$|\Delta(\phi\theta) - \Delta(\phi) - \Delta(\theta)| \leq c. \tag{I.2.4}$$

As mentioned above, it was the 1992 paper [4] which established the crucial fact that Δ may be characterized axiomatically as the unique continuous quasimorphism $\Delta : \widetilde{Sp}(2n) \to \mathbb{R}$ which is both homogeneous and continuous; see the 2008 paper [5] by Ben Simon and Salamon for more details. We observe here that Δ may be expressed (using some more advanced machinery) as the continuous and homogeneous quasimorphism with $[\partial^1\Delta] \neq 0$ as a cochain in the (continuous) bounded Lie group cohomology $H^2_{bc}(\widetilde{Sp}(2n), \mathbb{R})$ (as defined in [69]). With some of the rich structure enjoyed by the mean and Conley-Zehnder indices now introduced, we remind the reader that a hefty price must be paid in extending the mean index from the symplectic group to $\mathcal{L}_{2n}$, which is perhaps most plainly demonstrated by the lack of a group structure over the linear canonical relations of a fixed symplectic vector space when equipped with set-theoretic composition (definition I.3.2) so that in particular, the universal cover is no longer a group and consequently, does not admit quasimorphisms (at least in the sense of the above definition). We will discuss this in more detail in the following section I.3.

In regards to the utility of Δ, it is shown in [10] using equation (I.2.1) that the mean index and its associated spectrum might be use to extract information, even when about a degenerate periodic point, regarding the grading shift isomorphisms relating the various local Floer homologies of the iterates of a given Hamiltonian, which the author then shows can be used to give a novel proof for the symplectically aspherical Conley conjecture. One may also find an example of inequality (I.2.2) above being applied within Hamiltonian dynamics in [27] wherein the authors prove a local variant of the Conley conjecture about an isolated periodic point of a Hamiltonian on a closed and symplectically aspherical manifold, utilizing the mean index by defining a filtration on the local Floer homology to supplement the usual action filtration. More than just helping to validate the Conley conjecture, the exceptional cases wherein only finitely many periodic points exist present interesting applications for the mean index in their own right. In particular, the paper [29] establishes various conditions (some regarding the mean indices of the periodic points) that a Hamiltonian 'pseudo-rotation'[15] of complex projective space must satisfy. The results reached in [11], this time regarding any symplectic manifold admitting a pseudo-rotation, again utilize the

[15]We adopt for the above paragraph the nomenclature that any Hamiltonian with finitely many periodic points is called a pseudo-rotation

mean indices of the periodic points to extract symplectic topological information. It is shown for any manifold admitting a pseudo-rotation (along with some additional conditions which vary with each of the following) that there exists an upper bound on the minimal Chern number, a lower bound on the quantum cohomology cup length, and the existence of some non-zero Gromov-Witten invariants. In particular the cup length lower bound is determined by the mean indices in a manner resembling certain 'non-resonance' conditions for stability as found in KAM theory [1, 67, 49]. The underlying index theory common to all of these references may be found in [31, 57] where the former focuses on Lusternik-Schnirelmann theoretical aspects and the later is a comprehensive text consisting of a detailed exposition on Maslov-type index theory.

I.3 Definitions and Conventions

We will be working over a real symplectic vector space (V^{2n}, ω) and adopt the following shorthand notation to denote V's twisted symplectic product,

$$V \times \overline{V} := (V \times V, \tilde{\omega} = \pi_1^*\omega - \pi_2^*\omega).$$

We will distinguish set theoretic from linear subspace inclusion by using '$\leq$' for the latter. A Lagrangian subspace $L \leq V \times \overline{V}$, also referred to as a linear canonical relation, will be said to have source and target V and $\overline{V}$ respectively[16] and we intro-

[16] Many authors define the target and source in the opposite manner to better complement composition.

duce the following notation, used in [80] for linear relations with the exception being the colorful notation halo(L) (introduced in [60] yet lately supplanted by the more conventional indet(L) as in [50]).

Definition I.3.1. *Recall for a symplectic vector space (V^{2n}, ω) that the Lagrangian Grassmannian* $\mathrm{LagGr}(V, \omega)$ *is defined as follows (where L^ω denotes the symplectic orthogonal),*

$$\mathrm{LagGr}(V, \omega) := \{L \leq V \mid L^\omega = L\} \subset \mathrm{Gr}_n(2n),$$

where $\mathrm{Gr}_n(2n)$ is the standard Grassmannian manifold of n-planes in a $2n$-dimensional vector space.

Given a linear canonical relation $L \in \Lambda_{2n} := \mathrm{LagGr}(V \times V, \tilde{\omega})$, we denote the following distinguished subspaces of V (where as above, π_1, π_2 are the projections from $V \times \overline{V}$ to the first and second coordinate),

- $\mathrm{dom}(L) := \{v \in V \mid \exists w \in V\ ,\ (v, w) \in L\} = \pi_1(L)$
- $\mathrm{ran}(L) := \{v \in V \mid \exists w \in V\ ,\ (w, v) \in L\} = \pi_2(L)$
- $\ker(L) := \{v \in V \mid (v, 0) \in L\} = \mathrm{dom}(L)^\omega$
- $\mathrm{halo}(L) := \{v \in V \mid (0, v) \in L\} = \mathrm{ran}(L)^\omega.$

It is true for any linear canonical relation $L \subseteq V \times \overline{V}$ (being a consequence of the choice of symplectic form) that both subspaces $\mathrm{dom}(L)$ and $\mathrm{ran}(L)$ are coisotropic (equivalently, $\ker(L), \mathrm{halo}(L)$ are isotropic) and that $\dim(\mathrm{dom}(L)) = \dim(\mathrm{ran}(L))$.

Definition I.3.2. *Given $L, L' \in \Lambda_{2n}$, the set theoretic composition for linear canonical relations is defined as follows,*

$$L \circ K = \{(v,z) \in V \times V \,|\, \exists w \in V \text{ s.t. } (v,w) \in K,\, (w,z) \in L\}.$$

We also denote

$$L^l = \underbrace{L \circ L \circ \cdots \circ L}_{l \text{ times}} \tag{I.3.1}$$

for any $l \geq 0$.

We let $L^0 := \{(v,v) \in V \times V \,|\, v \in V\}$ denote the diagonal $\triangle_V$ which is the identity in the monoid of linear Lagrangian relations [84]; for any $L \in \Lambda_{2n}$, $\triangle_V \circ L = L = L \circ \triangle_V$.

Definition I.3.3. *Define the following smooth map sending each $A \in \mathrm{Sp}(V)$ to its graph, a Lagrangian subspace of $V \times \overline{V}$,*

$$\begin{array}{rcl} \mathrm{Sp}(V) & \underset{\mathrm{Gr}}{\hookrightarrow} & \Lambda_{2n} \\ A & \mapsto & \left\{(v, Av) \in V \times \overline{V} \,|\, v \in V\right\}. \end{array}$$

As shown in [39], the above map has an open and dense image in Λ_{2n} (in fact, this embedding is analytic), so that in particular Λ_{2n} is a compactification of $\mathrm{Sp}(2n)$. Next we partition Λ_{2n} into $n+1$ pair-wise disjoint sets defined for each $0 \leq k \leq n$ (note that $\Lambda_{2n}^0 = \mathrm{Im}(\mathrm{Gr}) \cong \mathrm{Sp}(V)$),

$$\Lambda_{2n}^k := \{L \in \Lambda_{2n} \,|\, \dim(\ker(L)) = \dim(\mathrm{halo}(L)) = k\}.$$

As shown in [46], these $n+1$ sets are in fact smooth submanifolds that form a stratification[17] of Λ_{2n},

$$\Lambda_{2n}^{n} \subset \Lambda_{2n}^{\geq n-1} \subset \cdots \subset \Lambda_{2n}^{\geq 1} \subset \Lambda_{2n}^{\geq 0} = \Lambda_{2n},$$

where $\Lambda_{2n}^{\geq r} := \bigcup_{k=r}^{n} \Lambda_{2n}^{k}$. Note that since $\Lambda_{2n}^{\geq 1} = \Lambda_{2n} \setminus \mathrm{Im}(\mathrm{Gr})$ is the complement of an open set, we see that $\Lambda_{2n}^{\geq 1}$ is a closed stratified space. We show in theorem I.4.1 that each smooth submanifold Λ_{2n}^{k} is of codimension k^2 in Λ_{2n} for $0 \leq k \leq n$, and will use this stratification to define our stratum-regular paths in section II.4.

Definition I.3.4. *For all $n \geq 1$ we define the set of exceptional Lagrangian subspaces,*

$$H := \{L \in \Lambda_{2n} \mid \ker(L) \cap \mathrm{ran}(L) \neq \{0\}\},$$

and let $\mathcal{L}_{2n} := \Lambda_{2n} \setminus H$.

Remark I.3.5. *We explore the possibility of choosing a smaller exceptional set $\hat{H}$ in section VII.3, taking advantage of a tuple of invariants* (II.3.1)*, shown in Lorand's paper* [59] *to completely characterize the conjugacy classes of (co)isotropic pairs (see definition II.1.1 for more details on isotropic pairs).*

As first determined in [43] following [62], within the context of microlocal analysis (which tends to be the field in which one finds early examples of Lagrangian relations), given three symplectic vector spaces[18] X, Y and Z and a pair of linear canonical

[17] The author gives their apologies for breaking with the traditional indexing of a stratification, which is reversed in our notation, i.e. the lowest index denotes the highest (dimension) stratum.

[18] The category whose objects are symplectic vector spaces and morphisms are linear canonical relations is often denoted **SLREL**, and (to the best of the author's knowledge) was first formally constructed in [6]

relations $L \leq X \times Y$ and $L' \leq Y \times Z$, the set-theoretic composition $L' \circ L$ presents some serious issues, even in the linear case, unless one at least imposes the transversality condition $\mathrm{dom}(L') \oplus \mathrm{ran}(L) = Y$ or equivalently, $(L' \times L) \cap (\{0_X\} \times \triangle_Y \times \{0_Z\}) = \{0\}$.

Remark I.3.6. *Note that our transversality condition $L \notin H$ is stronger than that introduced above. By taking the symplectic orthogonal of Hörmander's transversality condition we get* $\ker(L') \cap \mathrm{halo}(L) = \{0\}$ *whereas we require (under a more restrictive assumption) that* $\ker(L') \cap \mathrm{ran}(L) = \{0\}$*, which implies the former.*

The manner in which composition is defective when L, L' fail to satisfy the above transversality condition is easiest to observe if we consider the composition operation as a function; $* \circ * : \mathrm{LagGr}(X \times \overline{Y}) \times \mathrm{LagGr}(Y \times \overline{Z}) \to \mathrm{LagGr}(X \times \overline{Z})$. This function, while well defined, fails to be continuous unless one restricts the domain to those pairs (L, L') satisfying Hörmander's transversality condition (see example II.2.1 for a classic example of this failure of continuity). Several techniques have been established to circumnavigate this issue, and as one might expect, even more significant issues arise when translating this operation to non-linear objects (e.g. smooth canonical relations). One early solution (again, in the field of microlocal analysis) may be found in [34] in which the authors augment their Lagrangian relations with half densities. We will postpone a brief discussion to section VII.2 regarding some of the many issues that one might encounter in adapting our extended mean index to smooth objects, in which we also speculate on what the most promising categorical 'extension' of linear canonical relations might be for our purposes (a technique first developed in [83] and later refined in [50]). That said, our extension faces some more immediate issues.

The following definitions and lemmas are only applicable to Lie groups and as we have warned, they will be of no real use to us going forward. Regardless, they are important inclusions; in stating them we elucidate the algebraic properties instrumental in applying the mean index as described above [27, 10, 29, 11] so as to fully grasp what our extended index is missing. We express them in general terms with this aim in mind and begin with the algebraic characterization of the map ρ.

Definition I.3.7. *[73] For a connected Lie group G with $\pi_1(G) \cong \mathbb{Z}$ we call any smooth $\rho : G \to S^1$ for which $\rho_* : \pi_1(G) \xrightarrow{\cong} \mathbb{Z}$ is an isomorphism a circle map. Additionally, if the following two properties hold, we call it a normalized circle map;*

1. $\rho(\phi^{-1}) = \rho(\phi)^{-1}$

2. $\rho(Id) = 1$.

As shown in [73] our map ρ is a normalized circle map on $\mathrm{Sp}(2n)$.

Lemma I.3.8. *[38] Given a pointed topological space (X, x_0) for which a universal cover $\widetilde{X}$ exists then the set of fixed-endpoint homotopy classes of paths in X which originate at the point x_0 may be identified with the universal cover $\widetilde{X}$.*

As is becoming clear, the vast majority of applications for the mean index rely on the fact that $\pi_1(\mathrm{Sp}(2n)) \cong \mathbb{Z} \cong \pi_1(S^1)$ and that ρ is a circle map. We proceed with a characterization from Rawlins (based on Milnor's [65]) of the universal cover for such Lie groups.

Lemma I.3.9. *[73] Given any connected Lie group G with $\pi_1(G) \cong \mathbb{Z}$ equipped with a normalized circle map ρ, the universal cover of G may be written as*

$$\tilde{G} = \left\{ (g, c) \in G \times \mathbb{R} \,|\, \rho(g) = e^{ic} \right\}.$$

Where the group action for $\tilde{G}$ is given by,

$$(g_1, c_1) \cdot (g_2, c_2) = (g_1 g_2, c_1 + c_2).$$

The similarity between the 2^{nd} real coordinate in the above lemma and the mean index is no coincidence as the following definition exhibits.

Definition I.3.10. *[65] If we let $\tilde{\rho} : \tilde{G} \to S^1$ denote the composition of the universal cover projection $j : \tilde{G} \to G$ with ρ we may write an alternative definition of the mean index $\Delta : \tilde{G} \to \mathbb{R}$ as the lift of $\tilde{\rho}$:*

$$\begin{tikzcd} \tilde{G} \arrow[r, dashed, "\Delta"] \arrow[dr, "\tilde{\rho}"'] & \mathbb{R} \arrow[d, "e^{it}"] \\ & S^1 \end{tikzcd}$$

We leave the explicit reconciliation of definition I.3.10 and definition I.3.14 below to [73], but it's not a stretch to see from lemma I.3.9 that any path $\gamma : I \to G$ with $\gamma(0) = Id$, when lifted: $\tilde{\gamma} : I \to \tilde{G}$, terminates at some point $(\gamma(1), c)$ so that $\Delta((g, c)) = c$.

Now we take a deeper look into the role the quasimorphism property plays in this algebraic context.

Lemma I.3.11. *[73] Given any connected Lie group G of dimension n with $\pi_1(G) \cong \mathbb{Z}$ equipped with a normalized circle map then there exists a unique map $\eta : G \times G \to \mathbb{R}$ for which $\eta(Id, Id) = 0$ and*

$$\rho(g_1 g_2) = \rho(g_1)\rho(g_2)e^{i\eta(g_1,g_2)}.$$

Where $\eta(g_1, g_2) := \Delta(\tilde{g}_1\tilde{g}_2) - \Delta(\tilde{g}_1) - \Delta(\tilde{g}_2)$ is independent of the choice of lifts $\tilde{g}_1$ and $\tilde{g}_2$ and $|\eta(g_1, g_2)| < \frac{n\pi}{2}$.

Additionally η satisfies the cocycle condition,

$$\eta(g_1, g_2) + \eta(g_1 g_2, g_3) = \eta(g_1, g_2 g_3) + \eta(g_2, g_3).$$

Remark I.3.12. *Since our definition normalizes multiples of 2π to the integers the above inequality $|\eta(g_1, g_2)| < \frac{n\pi}{2}$ is scaled incorrectly with respect to definition I.3.14.*

As warned, the constructions for Δ and η above are heavily dependent on the group structure of G so it should come as no surprise that proceeding with a monoid is not a feasible strategy, in particular since the fundamental group is not likely to be isomorphic to $\mathbb{Z}$.

Remark I.3.13. *Since the non-singular portions of H are codimension two, it is a reasonable supposition that $\pi_1(\mathcal{L}_{2n})$ will not be isomorphic to $\mathbb{Z}$. If so, and even if one ignores the lack of group structure and simply proceeds with definition I.3.10 in lifting $\tilde{\rho} : \widetilde{\mathcal{L}_{2n}} \to S^1$ one will end up grappling with, among other issues, a pre-image $\Delta^{-1}(0)$ (paths up to homotopy) which is far too large, an effect of the induced map*

$\rho_* : \pi_1(\mathcal{L}_{2n}) \to \mathbb{Z}$ no longer being injective. In example VII.1.2 we see this is indeed the case when $n = 1$; there exist many distinct and non-contractible (homotopy classes of) loops in $\mathcal{L}_2$ which have zero mean index. We also exhibit a rather ad-hoc strategy of replacing the universal cover with an intermediate covering space over which $\hat{\Delta}$ is injective on homotopy classes of paths.

We conclude the section by giving a second definition for the mean index which we will follow in our own construction of $\hat{\Delta}$ (although using definition I.3.10 would work just as well for identity-based paths).

Definition I.3.14. *Given any path $\gamma\colon [0,1]1 \to \mathrm{Sp}(V)$ there exists a unique, continuous lift $\theta\colon [0,1] \to \mathbb{R}$ such that $(\rho \circ \gamma)(t) = e^{i\theta(t)}$ and $\theta(0) \in [-\pi, \pi)$. Then the mean index for the path γ is defined as*

$$\Delta(\gamma) := \frac{\theta(1) - \theta(0)}{2\pi}.$$

Remark I.3.15. *The above definition must be slightly altered before attempting to extend ρ as shown below in example VII.3.4. Specifically, we will be continuously extending ρ^2 and because of this the Lagrangian mean index $\hat{\Delta}$ will differ from Δ by a factor of two for paths $\gamma : [0,1] \to \mathrm{Im}(\mathrm{Gr}) \cong \mathrm{Sp}(V)$.*

I.4 Dissertation Outline

I.4.1 Outline of Results

Theorem I.4.1. *The set H given above in definition I.3.4 has codimension two in Λ_{2n}.*

Recall that H is the exceptional set on which the circle map ρ^2 may not be continuously extended. It manifests in the two dimensional case as a circle bridging the two connected components of the parabolic transformations at a projective 'line at infinity' outside the image $\mathrm{Gr}(\mathrm{Sp}(2)) \subset \Lambda_2$ (see figure VII.1.2). We show later in example VII.3.2 that in higher dimensions (namely $n \geq 3$) there exist $L \in H$ to which ρ^2 may be continuously extended. For more details as to how these conditions might be relaxed see proposition VII.3.1 for partial results which depend on the invariants associated with isotropic pair conjugacy classes as defined in section II.3.

Theorem I.4.2. *There exists a unique real valued continuous function $\hat{\Delta}$ defined on fixed endpoint homotopy classes of paths in $\mathcal{L}_{2n}$ such that for any path $\gamma \in \mathrm{Sp}(2n)$ we have $\hat{\Delta}(\mathrm{Gr}(\gamma)) = 2\Delta(\gamma)$.*

This claim requires a more intricate proof than the others, although that privilege rightly belongs to theorem V.3.1 wherein we produce the continuous extension $\hat{\rho}$ of ρ^2. Indeed, provided a continuous extension $\hat{\rho}$ exists we may apply definition I.3.14 and see that any path $\gamma\colon I \to \mathcal{L}_{2n}$, when composed with $\hat{\rho}$, lifts to a unique continuous

$\hat{\theta}\colon I \to \mathbb{R}$ such that $(\hat{\rho} \circ \gamma)(t) = e^{i\hat{\theta}(t)}$ and $\hat{\theta}(0) \in [-\pi, \pi)$. Then the extended mean index for the path γ may be defined just as in definition I.3.14 as $\hat{\Delta}(\gamma) := \frac{\hat{\theta}(1)-\hat{\theta}(0)}{2\pi}$.

Similarly we may precompose $\hat{\rho}$ with the universal covering map $j : \widetilde{\mathcal{L}}_{2n} \to \mathcal{L}_{2n}$ and lift this to obtain $\hat{\Delta} : \widetilde{\mathcal{L}}_{2n} \to \mathbb{R}$ which agrees with the above definition by letting a homotopy class of paths be mapped to the lifted paths' shared terminal point via lemma I.3.8.

Now we state the theorem at the core of the proof for theorem I.4.2.

Theorem V.3.1. There exists a unique continuous map $\hat{\rho}\colon \mathcal{L}_{2n} \to S^1$ such that for all $\phi \in \mathrm{Sp}(V)$ we have $\hat{\rho}(\mathrm{Gr}(\phi)) = \rho^2(\phi)$.

Already in dimension two the purpose of squaring ρ is clear; there exist sequences $A_i^{\pm}$ for which $Tr(A_i^{\pm}) > 2$ for all $i \in \mathbb{N}$ (equivalent to hyperbolicity) and $\mathrm{Gr}(A_i^{\pm}) \to L \notin H$ yet $\rho(A_i^{\pm}) = \pm 1$ for all $i \in \mathbb{N}$, exhibiting that even in the nicest case there will still be two distinct limiting values for ρ approaching $\Lambda_{2n} \setminus \mathrm{Im}(\mathrm{Gr})$, meaning no continuous extension of ρ exists (see example VII.3.4 for an example applicable to symplectic vector spaces of all dimensions).

Remark I.4.3. *When approaching H via the elliptic transformations significant discontinuities arise even in low dimensions, in [32] an example is given for some fixed $L \in H$ in which an $S^1 \setminus \{\pm 1\}$ family of sequences of symplectic maps $\left\{A_i^{\theta}\right\}_{i=1}^{\infty} \in \mathrm{Sp}(2)$ are constructed such that each has $\lim_{i\to\infty} \mathrm{Gr}(A_i^{\theta}) = L$ yet $\rho(A_i^{\theta}) = \theta$ for all $i \in \mathbb{N}$.*

The following lemma is a partial result towards theorem I.4.5.

Lemma IV.2.1. The extended circle map $\hat{\rho}$ is homogeneous on $\mathcal{L}_{2n}$; given any $L \in \mathcal{L}_{2n}$ then $\hat{\rho}(L) = l \cdot \hat{\rho}(L) \in \mathbb{R}/\mathbb{Z}$ for all $l \in \mathbb{N}$.

The remaining results are regarding the algebraic properties retained in the extended mean index, each of which relies on the technical lemma VI.1.5. Before stating these results, we must give a definition.

Definition I.4.4 (Stratum-Regular Paths)**.** *We denote the set of stratum-regular paths as $\mathcal{P}_{reg}(\mathcal{L}_{2n}) \subset C^1([0,1], \mathcal{L}_{2n})$, defined as the set of all paths $\gamma \in C^1([0,1], \mathcal{L}_{2n})$ which are transversal (in the sense of definition II.4.3) to the stratified space $(\mathcal{L}_{2n}^k)_{k=1}^n$.*

we show in lemma II.4.13 that $P_{reg}(\mathcal{L}_{2n})$ is open and dense in $C^1([0,1], \mathcal{L}_{2n})$. We postpone the complete definition of the equivalence relation $\sim_{comp}$ on $\mathcal{P}_{reg}(\mathcal{L}_{2n})$ to definition VI.1.1, as the fact that it is an equivalence relation will be sufficient to state the following theorem.

Theorem I.4.5. *Given any $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$ and $\tau \in [\gamma]$ (the set of all stratum-regular paths compatible with γ, see proposition VI.1.3 for details on this equivalence relation) then both $\gamma \circ \tau$ and $\tau \circ \gamma$ are well defined, piece-wise differentiable paths in $\mathcal{L}_{2n}$ which are smooth on their intersection with the symplectic group. In particular γ^l is defined for all $l \geq 0$.*

The first property we regain from the original mean index is homogeneity as a corollary of lemmas IV.2.1 and VI.1.5.

Corollary I.4.6. *For any $l \geq 0$,*

$$\hat{\Delta}(\gamma^l) = l \cdot \hat{\Delta}(\gamma).$$

i.e. The mean index $\hat{\Delta}$ is homogeneous over stratum-regular paths.

Theorem I.4.7. *For any $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$ and $\tau \in [\gamma]$ the Lagrangian mean index $\hat{\Delta}$ satisfies the quasimorphism-type bound,*

$$|\hat{\Delta}(\gamma \circ \tau) - \hat{\Delta}(\tau) - \hat{\Delta}(\gamma)| < C,$$

where $C \in \mathbb{R}$ and the bound is uniform over all pairs of paths residing in $[\gamma]$.

Remark I.4.8. *One might extend the aforementioned results to those paths γ satisfying the less restrictive condition that $|\pi_0(\gamma^{-1}(\mathcal{L}_{2n}^k))| < \infty$ for each $0 \leq k \leq n$, thereby allowing paths which intersect arbitrary strata (as well as even non-transverse intersections). The added technical details in showing such a result are not trivial and include the process of collating each fiber component $\phi_{\gamma(t)} \in \mathrm{Sp}(\mathrm{dom}(\gamma(t)) \cap \mathrm{ran}(\gamma(t)))$ (see theorem III.2.5) into a single path of symplectomorphisms when encountering non-transverse stratum intersections with higher strata. The added complexity in proving analogous results over this alternate definition is disproportionate when compared to the generality gained. Because of this we will content ourselves in defining stratum-regularity as given in definition I.4.4 above.*

I.4.2 Outline of Proofs

We first note that part VI contains the proofs regarding the stratum-regular paths including corollary I.4.6 as well as the quasimorphism-type bound stated in theorem I.4.7. In addition to these, the part concludes with proofs for both the technical lemma VI.1.5 and theorem I.4.5, together being critical in proving the preceding results. In regards to the proof of our central theorem V.3.1 (that is, continuously extending ρ^2 as described above) it will come as a fairly straightforward consequence of the following three theorems.

Theorem III.2.5. For a given $L \in \mathcal{L}_{2n}$ there exists a unique symplectic decomposition of $V = V_s \oplus V_g$ and $\phi \in \mathrm{Sp}(V_g)$ such that

$$L = (\ker(L) \times \{0\}) \oplus (\{0\} \times \mathrm{halo}(L)) \oplus \mathrm{Gr}(\phi) \leq (V_s \times \overline{V}_s) \oplus (V_g \times \overline{V}_g) \qquad \text{(I.4.1)}$$

where $\ker(L), \mathrm{halo}(L) \in \mathrm{LagGr}(V_s)$ are transverse.

Theorem V.1.1. Given any sequence $\{A_i\}_{i=1}^{\infty} \subset \mathrm{Sp}(V)$ for which each A_i has distinct eigenvalues such that

$$\mathrm{Gr}(A_i) \to L_{\mathrm{dom}} \times \{0\} \oplus \{0\} \times L_{\mathrm{ran}} \in \mathcal{L}_{2n},$$

(i.e. $L_{\mathrm{dom}}, L_{\mathrm{ran}} \in \Lambda_n$ are transversal), then the A_i will eventually have no eigenvalues $\lambda \in S^1 \setminus \{\pm 1\}$. In particular this shows that $\rho^2(A_i) = 1$ for sufficiently large i.

Theorem V.2.1. Consider any sequence $\{A_i\}_{i=1}^{\infty} \subset \mathrm{Sp}(V)$ where each A_i has distinct eigenvalues and for which

$$\mathrm{Gr}(A_i) \underset{i \to \infty}{\to} L \in \mathcal{L}_{2n},$$

where the graph part $\phi_L \in \mathrm{Sp}(V_g)$ (see remark III.2.3) of L has semisimple eigenvalues.

Then this sequence eventually induces an associated sequence of unique, A_i invariant symplectic decompositions $V = E_s^i \oplus E_g^i$ so that may write $A_i = \alpha_i \oplus \beta_i \in \mathrm{Sp}(E_s^i) \times \mathrm{Sp}(E_g^i)$ such that $\mathrm{Gr}(\alpha_i) \to \ker(L) \times \{0\} \oplus \{0\} \oplus \mathrm{halo}(L)$.

Additionally there exists an $N \in \mathbb{N}$ for which there is a sequence of symplectic isomorphisms

$$\left\{ I_i \colon (E_g^i, \omega|_{E_g^i \times E_g^i}) \xrightarrow{\cong} (V_g, \omega|_{V_g \times V_g}) \right\}_{i=N}^{\infty}$$

uniquely determined by L such that each $\beta_i \colon E_g^i \to E_g^i$ is conjugate via I_i to some $\phi_i \in \mathrm{Sp}(V_g)$ for all $i \geq N$ with $\phi_i \to \phi$. We also show that the β_i preserve the data used in computing ρ, namely the eigenvalues and the conjugacy classes of the A_i restricted to elliptic eigenspaces.

Remark I.4.9. *Refer to [35] for a detailed exposition showing how ρ may be defined on the semisimple elements $A \in \mathrm{Sp}(V)$ and then extended to all of $\mathrm{Sp}(V)$. It's purpose here is to guarantee the E_g^i and E_s^i do not become singular in the limit.*

With these three ingredients and the fact that ρ is multiplicative with respect to direct sums we conclude the proof of theorem V.3.1 in section V.3 setting $\hat{\rho}(L) := \rho^2(\phi_L)$, followed by a brief revisit to the final arguments already covered above (namely applying definition I.3.14 to $\hat{\rho}$) and culminating in the proof of theorem I.4.2.

Part II

Linear Canonical Relations and Isotropic Pairs

II.1 Strata of the Lagrangian Grassmannian

As before, we work over a real symplectic vector space (V^{2n}, ω) and denote our set of admissible Lagrangian subspaces as $\mathcal{L}_{2n} := \Lambda_{2n} \setminus H$ (see definition I.3.4 above for details).

Definition II.1.1. *Define the set $I_k(V)$ for all $0 \leq k \leq n$ to be the Grassmannian of dimension k isotropic subspaces of V,*

$$I_k(V) = \{B \leq V \,|\, B \leq B^{\omega},\ \dim(B) = k\}.$$

We set $I_0(V) := \{0\}$ and call any $(B_1, B_2) \in I_k(V) \times I_k(V)$ an isotropic pair.

Definition II.1.2. *Define the isotropic pair projection map,*

$$\begin{aligned} \mathrm{Pr}_I : \Lambda_{2n} &\to \sqcup_{k=0}^{n} I_k(V) \times I_k(V) \\ L &\mapsto (\ker(L), \mathrm{halo}(L)). \end{aligned}$$

Remark II.1.3. *Despite the significant discontinuities in the above map that are a consequence of the disjoint union, we define it as above for notational convenience since each usage of Pr_I below (unless specified otherwise) is over a fixed stratum.*

Lemma II.1.4. *[39] Each Λ_{2n}^k is a fiber bundle over $I_k(V) \times I_k(V)$ with fiber diffeomorphic to $\mathrm{Sp}(2n-2k)$ and definition II.1.2 as the base projection, see the figure below.*

$$\begin{tikzcd}
\mathrm{Sp}(2n-2k) \cong \Lambda^0_{2n-2k} \arrow[r, hook] & \Lambda^k_{2n} \arrow[d, two heads, "\mathrm{Pr}_I"] \\
& \mathcal{I}_k
\end{tikzcd}$$

Figure II.1: For each $0 \leq k \leq n$, the stratum Λ^k_{2n} forms a fiber bundle.

Remark II.1.5. *We have implicitly used the notational conventions* $\mathrm{Sp}(0) \cong \Lambda^0_n := \{0\}$ *in the above fiber bundle so that in both cases the strata remain well defined fiber bundles (albeit trivial in one manner or another).*

When $k = n$ the projection map Pr_I is a diffeomorphism,

$$\{0\} \hookrightarrow \Lambda^n_{2n} \overset{\mathrm{Pr}_I}{\twoheadrightarrow} \Lambda_n \times \Lambda_n,$$

as each fiber is trivial. One may verify that $\Lambda^n_{2n} \cong \Lambda_n \times \Lambda_n$ is the space of Lagrangian pairs in V and is the only closed smooth stratum in Λ_{2n}.

On the other hand when $k = 0$ we have a bundle with trivial base where definition I.3.3 maps $\mathrm{Sp}(V)$ *onto the lone fiber,*

$$\mathrm{Sp}(V) \overset{\mathrm{Gr}}{\hookrightarrow} \Lambda^0_{2n} \overset{\mathrm{Pr}_I}{\twoheadrightarrow} I_0(V) = \{0\}\,.$$

This yields the Lie group isomorphism $\mathrm{Gr} : (\mathrm{Sp}(V), \cdot) \overset{\cong}{\to} (\Lambda^0_{2n}, \circ)$ *where Lagrangian composition is the Lagrangian composition operation* $(* \circ *)$ *given in definition I.3.1 above.*

We define each admissible strata $\mathcal{L}^k_{2n} := \Lambda^k_{2n} \setminus H$ and let

$$\mathcal{I}_k := \{(B_1, B_2) \in I_k(V) \times I_k(V) \,|\, B_1 \pitchfork B_2^\omega\} = \mathrm{Pr}_I(\mathcal{L}^k_{2n})$$

denote the space of admissible isotropic pairs of dimension k. We see that $\mathcal{L}^k_{2n} \to \mathcal{I}_k$ remains a $\mathrm{Sp}(2n-2k)$ fiber bundle for each $0 \leq k \leq n$ (e.g. one might realize it as the pullback bundle under the inclusion map $\mathcal{I}_k \hookrightarrow I_k(V) \times I_k(V)$). In the extreme cases we see that $\mathcal{L}^0_{2n} = \Lambda^0_{2n}$ and $\mathcal{L}^n_{2n} = \Lambda_n \times \Lambda_n \setminus \hat{\Sigma}_n$ where,

$$\hat{\Sigma}_n := H \cap \Lambda^n_{2n} = \{(B_1, B_2) \in \Lambda_n \times \Lambda_n \mid B_1 \not\pitchfork B_2^\omega \Leftrightarrow B_1 \not\pitchfork B_2\}$$

is the space of Lagrangian pairs which fail to be transverse.

One may observe that lemma II.1.4 implies that the set $\left(\mathcal{L}^{\geq k}_{2n}\right)_{k=0}^n$ is also a stratification of $\mathcal{L}_{2n}$ as the latter is an open and dense subset of Λ_{2n} and therefore a stratified manifold (though not a stratified subset of Λ_{2n} in the sense of definition II.4.1).

Proposition II.1.6. *The codimension of each $\mathcal{L}^k_{2n}$ in $\mathcal{L}_{2n}$ is k^2.*

Proof. It is true that the Grassmannian of isotropic k-planes in $(\mathbb{R}^{2n}, \omega_{std})$ has

$$\dim(I_k(V)) = \frac{k}{2}\left(4n - 3k + 1\right),$$

so that the above fibration and the following routine computation confirm the claim,

$$\begin{aligned}
\mathrm{codim}(\mathcal{L}^k_{2n}) &= dim(\mathcal{L}_{2n}) - dim(\mathcal{I}_k) - dim(\mathrm{Sp}(2n-2k)) \\
&= 2n^2 + n - k(4n - 3k + 1) - (n-k)(2n - 2k + 1) \\
&= k(4n - 2k + 1) - k\left(4n - 3k + 1\right) \\
&= k^2.
\end{aligned}$$

□

II.2 Iterating Linear Canonical Relations

As mentioned in the introduction, the composition map is not continuous everywhere, which motivates the following classic example demonstrating the need for transversality.

Example II.2.1. *Let* $K_i = \mathrm{Gr}(A_i)$ *and* $K_i' = \mathrm{Gr}(A_i^{-1})$ *for* $\{A_i\}_{i=1}^{\infty} \subset \mathrm{Sp}(2n)$ *such that* $K_i \to K = L_{\mathrm{dom}} \times \{0\} \oplus \{0\} \times L_{ran} \in \Lambda_{2n}$ *where both* $L_{\mathrm{dom}}, L_{ran} \in \Lambda_n$ *for* $i = 1, 2$. *Then the set-theoretic composition* $K_i' \circ K_i = K_i \circ K_i' = \triangle_V$ *for all* $i \in \mathbb{N}$ *so that* $lim_{i\to\infty}(K_i' \circ K_i) = \triangle_V$.

Yet, since $K_i' \to K' = L_{ran} \times \{0\} \oplus \{0\} \times L_{\mathrm{dom}}$ *as* $i \to \infty$ *we may form the compositions of the limits,*

$$K' \circ K = L_{\mathrm{dom}} \times \{0\} \oplus \{0\} \times L_{\mathrm{dom}}$$

$$K \circ K' = L_{ran} \times \{0\} \oplus \{0\} \times L_{ran}$$

so that we indeed have a failure of continuity,

$$lim_{i\to\infty} K_i' \circ lim_{i\to\infty} K_i = L_{\mathrm{dom}} \times \{0\} \oplus \{0\} \times L_{\mathrm{dom}} \neq \triangle_V = lim_{i\to\infty}(K_i' \circ K_i).$$

Note here that $\mathrm{ran}(K) = \mathrm{dom}(K')$ *so the pair are in some sense maximally non-transversal.*

Regarding $L \in \mathcal{L}_{2n}$, as mentioned above our condition that $\mathrm{dom}(L) \oplus \mathrm{halo}(L) = V$ is stronger than that needed to prevent discontinuities like the above; $\mathrm{dom}(L) \oplus \mathrm{ran}(L) = V$.

Lemma II.2.2. *If $L \in \mathcal{L}_{2n}$ then $L^l \in \mathcal{L}_{2n}$ for all $l \geq 0$. Additionally each $L \in \mathcal{L}_{2n}$ satisfies $\ker(L^l) = \ker(L)$ and $\mathrm{halo}(L^l) = \mathrm{halo}(L)$ for all $l \geq 1$ making each iteration map $(*)^l : \mathcal{L}_{2n}^k \to \mathcal{L}_{2n}^k$ a bundle map for all $0 \leq k \leq n$. In particular this implies that $L^l = L$ for any $L \in \mathcal{L}_{2n}^n$ and $l \geq 1$.*

Proof. To show this we first observe that $\ker(L) \leq \mathrm{dom}(L^i)$ for any $i \geq 1$ since $(0,0)$ is contained in every canonical relation. Then if $\dim(\ker(L)) = k$ we may write a basis $(d_1, \ldots, d_{2n-2k})$ such that $\langle d_j \rangle_{j=1}^{2n-2k} \oplus \ker(L) = \mathrm{dom}(L)$ and thus each d_j is associated (non-uniquely) via L to some $r_j \in \mathrm{ran}(L)$. The r_j are also linearly independent since if $r_1 = \sum_{j=2}^{2n-2k} c_j r_j$ then $d_1 - \sum_{j=2}^{2n-2k} c_j d_j \in \ker(L)$ which violates $\langle d_j \rangle_{j=1}^{2n-2k} \pitchfork \ker(L)$.

Now again, since $L \notin H$, then $\langle r_j \rangle_{j=1}^{2n-2k} \oplus \mathrm{halo}(L) = \mathrm{ran}(L)$ and we see that $\mathrm{dom}(L) \cap \mathrm{ran}(L) = \mathrm{dom}(L) \cap \langle r_j \rangle_{j=1}^{2n-2k}$. Since $\dim(\mathrm{dom}(L) \cap \mathrm{ran}(L)) \geq 2n - 2k$ we see for dimensional reasons that $\langle r_j \rangle_{j=1}^{2n-2k} \leq \mathrm{dom}(L)$ and therefore

$$\mathrm{dom}(L^2) = (\langle r_j \rangle_{j=1}^{2n-2k}) \oplus \ker(L) = \mathrm{dom}(L).$$

An identical argument shows that $\mathrm{ran}(L^2) = \mathrm{ran}(L)$ as well. To conclude for $i \in \mathbb{N}$ we use this as the base case of a simple inductive argument regarding the domain and range of $L^i \circ L$ and $L \circ L^i$ which suffices to prove the claim for $i \geq 2$. $\square$

We prove in lemma IV.2.1 that this iteration map is continuous for each $i \in \mathbb{N}$, and that it sends differentiable paths to piece-wise differentiable paths in theorem I.4.5. The potential of extending homogeneity of the extended mean index over paths to negative numbers becomes a bit more interesting as it reverses the isotropic pair of L

and thus no longer descends to the identity on the base of each stratum. We address some of the routes which may be taken in the following remark.

Remark II.2.3. *The notion of an inverse in the category of linear relations is a fuzzy one (which is why we assume the powers to be non-negative in corollary I.4.6 and the statements which build up to including definition I.3.1 and lemma IV.2.1) but by far the most natural choice would be the 'reverse' of a linear relation* L*,*

$$\overline{L} := \{(v, w) \in V \times V \,|\, (w, v) \in L\}\,.$$

Then by letting $L^{-l} := \overline{L}^l$ *for* $l \geq 1$ *one may define iteration for all integers, albeit with the undesirable property that* $\ker(L^l) = \mathrm{halo}(L)$ *and* $\mathrm{halo}(L^l) = \ker(L)$ *for all* $l \leq -1$*. In particular this implies negative iteration (using the reverse of* L*) is not a bundle map.*

It seems far more useful in our case to consider a fiber-adapted composition which fixes any $(v, 0), (0, v) \in L$ *for some* $v \in V$ *while continuing to invert the rest of the Lagrangian relation. This composition, when restricted to each fiber, is Lie group isomorphic to a symplectic group,* $(\mathcal{L}^k_{2n})_{(B_1,B_2)} \cong \mathrm{Sp}(B_1^\omega \cap B_2^\omega)$ *(as detailed in proof of the second claim of theorem III.2.5), and these subgroups* $\mathrm{Sp}(B_1^\omega \cap B_2^\omega) \subset \mathrm{Sp}(V)$ *vary smoothly with respect to the isotropic pair* $(B_1, B_2) \in \mathcal{I}_k$*. In particular any* $(B_1, B_2) \in \mathcal{I}_k$ *and* $\phi, \tau \in \mathrm{Sp}(B_1^\omega \cap B_2^\omega)$ *would satisfy* $\mathrm{Gr}(\phi\tau) = \mathrm{Gr}(\phi) \circ \mathrm{Gr}(\tau)$ *and* $\mathrm{Gr}(\phi^i) \in (\mathcal{L}^k_{2n})_{(B_1,B_2)}$ *for every* $i \in \mathbb{Z} \setminus \{0\}$*.*

II.3 Conjugacy Classes of Isotropic Pairs

As above we denote the Grassmannian of isotropic subspaces of dimension k in (V, ω) as $I_k(V)$ and will call any ordered pair $(B_1, B_2) \in I_k(V) \times I_k(V)$ an isotropic pair. Consider the following notion of equivalence.

Definition II.3.1. *$(B_1, B_2) \sim (B_1', B_2')$ if and only if there exists $A \in \mathrm{Sp}(V)$ for which $(A(B_1), A(B_2)) = (B_1', B_2')$. The equivalence classes coincide with the orbits of the group action $\mathrm{Sp}(V) \circlearrowright \mathcal{I}_k$ where $A \cdot (B_1, B_2) = (A(B_1), A(B_2))$.*

For our purposes we have assumed the isotropic pairs have the same dimension; our goal in introducing them is to examine $L \in \Lambda_{2n}$ via the associated isotropic pair $(B_1, B_2) = (\ker(L), \mathrm{halo}(L))$ and since $\dim(\ker(L)) = \dim(\mathrm{halo}(L))$ this assumption is reasonable.

Theorem II.3.2. *[60] The four integers (r, κ, k, n) form a complete set of invariants for isotropic pairs subject to the relations $0 \le r \le \kappa \le k \le n$ and $0 \le \kappa - r \le n - k$ where*

$$(\kappa, r, k, n) = \left(\dim(B_1^\omega \cap B_2), \dim(B_1 \cap B_2), \dim(B_1), \frac{1}{2}\dim(V)\right). \tag{II.3.1}$$

Now denoting $\Lambda_{2n}^k := \{L \in \Lambda_{2n} \mid \dim(\ker(L)) = k\}$ for any $0 \le k \le n$ and noting $H = \{L \in \Lambda_{2n} \mid \kappa(L) \ge 1\}$ we see that the above equivalence relation on $I_k(V) \times I_k(V)$ induces an equivalence relation on Λ_{2n}^k for each $k \le n$ (and therefore on all of Λ_{2n}) where $L \sim L'$ if and only if $(\ker(L), \mathrm{halo}(L)) \sim (\ker(L'), \mathrm{halo}(L'))$. A detail to note is that on $\Lambda^0_{2n} \cong \mathrm{Sp}(V)$ all maps belong to a single equivalence class under this equivalence relation (hence, this notion of equivalence is missing the usual classification of symplectic transformations). We may compare this equivalence relation induced by the $\mathrm{Sp}(2n)$ action on isotropic pairs to a finer relation on Λ_{2n} induced by an essentially identical $\mathrm{Sp}(2n)$ action now acting on Λ_{2n}.

Definition II.3.3. *$L \sim_{\mathrm{Gr}} L'$ if and only if $(v, w) \in L \Leftrightarrow (Av, Aw) \in L'$. The equivalence classes coincide with the orbits of the group action $\mathrm{Sp}(V) \circlearrowright \Lambda_{2n}$ where $A \cdot ((x, y) \in L) \mapsto (Ax, Ay) \in A \cdot L$.*

This equivalence relation in particular splits the single $\sim$ equivalence class of Λ^0_{2n} into the usual conjugacy classes of $\mathrm{Sp}(V)$ while conversely, $\sim$ and $\sim_{Gr}$ are identical on Λ^n_{2n}. The classification and production of normal forms for $L \in \Lambda_{2n}$ with respect to this finer equivalence relation is, to the author's knowledge, incomplete with only partial results (namely the equivalence relation $\sim$) in [59].

II.4 Stratum-Regular Paths

Definition II.4.1. *Given a smooth manifold N, we call any closed subset $C \subset N$ a stratified space if there exists a finite sequence*[19] *of disjoint and locally closed smooth submanifolds $(C_i)_{i=1}^n$ for which $C = \bigcup_{i=1}^n C_i$. We also require that the following three*

[19] This definition is a particularly limited one; more general examples might specify the indexing set to be a partial order while others allow infinite indexing sets provided the family of strata is locally finite. The variety of stratifications that exist in the literature (see [85] for a thorough account from the perspective of algebraic K-theory and cobordisms) is enormous.

statements be equivalent (often called 'frontier conditions');

- $C_i \cap \overline{C_j} \neq \emptyset$
- $C_i \subseteq \overline{C_j}$
- $i \leq j$,

for all $1 \leq i, j \leq n$*. We will signify* C *as a stratified space by listing the strata* $(C_i)_{i=1}^n$.

In our case we will be considering a fairly well-behaved stratification $(\Lambda_{2n}^k)_{k=0}^n$ where for each $0 \leq k \leq n$, the stratum Λ_{2n}^k is locally closed relative to the open set $\Lambda_{2n}^{\leq k}$, and the closure of each stratum is the union of every stratum below it (in dimension) or above it (in index); $\overline{\Lambda_{2n}^k} = \Lambda_{2n}^{\geq k}$.

Remark II.4.2. *Recalling the footnote in the above definition, we remark that our notion of a stratified space corresponds to that of a topological stratification (by manifolds). In particular we do not require any regularity conditions be satisfied for the stratum tangent bundles near their closure, although our stratification* $(\Lambda_{2n}^k)_{k=1}^n$ *likely satisfies a richer definition. One might also consider using all of the invariants* (r, κ, k) *for isotropic pairs in* (II.3.1) *which completely characterize each orbit of* $I_k(V) \times I_k(V)$ *under the obvious symplectic group action [59]. Then, provided our indexing set* $\{(r, \kappa, k)\}_{0 \leq r \leq \kappa \leq k \leq n}$ *may be equipped with some partial order for which the components* $\Lambda_{2n}^{r,\kappa,k}$ *satisfy the appropriate conditions, one would obtain a finer stratification of* Λ_{2n}.

Definition II.4.3. *[3] For any smooth manifolds[20] M, N and smooth submanifold $C \subseteq N$, we call a C^1 map $f : M \to N$ transversal to C (notated $f \pitchfork C$) if for all $p \in f^{-1}(C)$,*

$$\operatorname{Im}(Df_p) + T_{f(p)}C = T_{f(p)}N.$$

More generally, given a stratified space $(C_i)_{i=1}^n \subset N$, we say that $f : M \to N$ is transverse to the stratified space $(C_i)_{i=1}^n$ when it is transverse to each C_i for $1 \leq i \leq n$.

Remark II.4.4. *Given $f \in C^1(M, N)$ and a smooth submanifold $C \subseteq N$, observe that $f \pitchfork C$ implies that $\operatorname{Im}(f) \pitchfork C$ as submanifolds, whereas the converse does not hold. In particular when $\operatorname{codim}(C) > \dim(M)$ the set of all $f \in C^1(M, N)$ which are transverse to C is identical to the set of all $f \in C^1(M, N)$ with $\operatorname{Im}(f) \cap C = \emptyset$, i.e. only the case of vacuous transversality is possible. When $\operatorname{codim}(C) \leq \dim(M)$ the latter is generally a proper subset of the former.*

Proposition II.4.5. *[3] In the following pair of statements both M and N are assumed to be smooth manifolds.*

1. *Given a pair of transversely intersecting smooth submanifolds $A, B \subseteq M$, then $A \cap B$ is a submanifold of M with $\operatorname{codim}(A \cap B) = \operatorname{codim}(A) + \operatorname{codim}(B)$.*

2. *If M is compact, $C \subseteq N$ some smooth submanifold and $f \in C^1(M, N)$ with $f \pitchfork C$, then $\operatorname{codim}(f^{-1}(C)) = \operatorname{codim}(C)$ as submanifolds of M and N respectively.*

[20] We maintain the convention throughout that a smooth manifold is both Hausdorff and second countable. As both are true for all of the relevant manifolds considered herein, this convention will not be necessary outside of the following general propositions.

Proposition II.4.6. *Given a smooth manifold M and $C, D \subset M$ a pair of transverse submanifolds with* $\dim(C) + \dim(D) = \dim(M)$*, then the intersection $C \cap D$ is both countable and discrete.*

Proof. We begin by applying part 1 of proposition II.4.5 to see that $C \cap D$ is a submanifold of dimension zero. Indeed, since the dimensions of C and D are complementary in M we may write the equivalent statement: $\mathrm{codim}(C) + \mathrm{codim}(D) = \dim(M)$ to see that $\mathrm{codim}(C \cap D) = \dim(M)$ or that $\dim(C \cap D) = 0$. Since smooth manifolds are second countable (that is, the topology induced by the smooth structure admits a countable basis) we may fix some countable base $\mathcal{U} = \{U_i\}_{i=1}^{\infty}$ for the subspace topology on the submanifold $C \cap D$. Then, as there exists some open neighborhood $U_i \in \mathcal{U}$ about each $p \in C \cap D$ and $C \cap D$ is Hausdorff (in particular, T_0) every point is topologically distinguishable from which it follows that $|C \cap D| \leq |\mathcal{U}| \leq \aleph_0$.

To show discreteness, we again observe that $C \cap D$ is a countable zero-dimensional submanifold of M. It follows from M being Hausdorff that there exists a countable collection of disjoint open neighborhoods $\{U_p\}_{p \in C \cap D} \subset M$ covering $C \cap D$ and separating each $p \in C \cap D$ so that $C \cap D$ is indeed a discrete subset. □

One might observe that under the conditions of the above proposition, $C \cap D$ must also be discrete as a subset of C and D as well.

Proposition II.4.7. *[3] Let M be a compact smooth manifold, N a smooth manifold* *and C a closed submanifold of N, then the set*

$$\{f \in C^1(M,N) \mid f \pitchfork C\}$$

is an open and dense subset of $C^1(M,N)$.

Proposition II.4.8. *The set $\mathcal{L}_{2n} \subset \Lambda_{2n}$ is open and dense.*

As proposition II.4.8 implies that H is closed, it is also a stratified space in Λ_{2n} with strata $(H_k)_{k=1}^n := (H \cap \Lambda_{2n}^k)_{k=1}^n$;

- Each stratum H_k is closed in $\Lambda_{2n}^{\leq k}$ for $1 \leq k \leq n$ (and therefore locally closed in Λ_{2n}).
- Each is a submanifold exhibited by the smooth fibration $\Lambda_{2n-2k}^0 \hookrightarrow H_k \to \hat{H}_k$ where $\hat{H}_k = Pr_I(H_k) \subset I_k \times I_k$. More concretely, each stratum H_k may be identified with the pullback bundle $H_k = i^*(\Lambda_{2n}^k)$ under the inclusion map $i : \hat{H}_k \hookrightarrow I_k \times I_k$.
- Each satisfies the frontier conditions (in a well behaved manner inherited from $(\Lambda_{2n}^k)_{k=0}^n$) exhibited by the identity $\overline{H_k} = H_{\geq k}$.

On the other hand, since $\mathcal{L}_{2n}$ is open in Λ_{2n} we cannot consider it as a stratified space in Λ_{2n}. Regardless, as $\mathcal{L}_{2n}$ is an open subset of Λ_{2n} we may consider $\mathcal{L}_{2n}$ simply as a smooth manifold after which the stratification of Λ_{2n} induces a stratification on $\mathcal{L}_{2n}$ with strata $\mathcal{L}_{2n}^k = \Lambda_{2n}^k \setminus H$. Indeed, the subspace topology ensures that each $\mathcal{L}_{2n}^k$

remains disjoint from the other strata, locally closed and also preserves each stratum's closure; i.e. $\overline{\mathcal{L}^{k}_{2n}} = \mathcal{L}^{\geq k}_{2n}$.

Proposition II.4.9. *Let M and N be smooth manifolds where M is closed and $U \subseteq N$ is some open and dense subset with complement $K := N \setminus U$. Then provided* $\operatorname{codim}(K) > \dim(M)$*, the set $C^1(M, U)$ is open and dense in $C^1(M, N)$.*

Proof. We begin by noting that the set $\{f \in C^1(M, N) \,|\, f \pitchfork K\}$ is open and dense in $C^1(M, N)$ by proposition II.4.7, as K is closed. Without the codimension bound on K we have the (generally proper) inclusion;

$$C^1(M, U) = \left\{f \in C^1(M, N) \mid \operatorname{Im}(f) \cap K = \emptyset\right\} \subseteq \left\{f \in C^1(M, N) \,|\, f \pitchfork K\right\},$$

but as mentioned in remark II.4.4, we see when $\operatorname{codim}(K) > \dim(M)$ that each $f \in C^1(M, N)$ and $p \in f^{-1}(K)$ yields

$$\dim\left(df_p(T_pM) + T_{f(p)}K\right) < \dim\left(T_{f(p)}N\right),$$

so that $f \not\pitchfork K$. Consequently, if $\operatorname{codim}(K) > \dim(M)$ we see that $f \pitchfork K$ if and only if $f^{-1}(K) = \emptyset$ so that vacuous transversality is the only possible kind of transversal intersection and thus the above inclusion is no longer proper;

$$\left\{f \in C^1(M, N) \,|\, \operatorname{Im}(f) \cap K = \emptyset\right\} = \left\{f \in C^1(M, N) \,|\, f \pitchfork K\right\}.$$

It follows that $C^1(M, U) = \{f \in C^1(M, N) \,|\, \operatorname{Im}(f) \cap K = \emptyset\}$ is open and dense in $C^1(M, N)$ when $\operatorname{codim}(K) > \dim(M)$. $\square$

Corollary II.4.10. *Consider the following two consequences of proposition II.4.9:*

- $C^1([0,1], \mathcal{L}_{2n})$ *is open and dense in* $C^1([0,1], \Lambda_{2n})$*.*

- $C^1([0,1], \mathcal{L}_{2n}^{\leq k})$ *is open and dense in* $C^1([0,1], \mathcal{L}_{2n})$ *for every* $1 \leq k \leq n$*.*

Proof.

- As $\mathcal{L}_{2n}$ is open and dense by proposition II.4.8, we see that its complement H is a closed stratified space with strata $H_k = \Lambda_{2n}^k \setminus \mathcal{L}_{2n}^k$ for $1 \leq k \leq n$. Then since $\operatorname{codim}(H) = \operatorname{codim}(H_1) = 2$ (as shown in theorem I.4.1), an application of proposition II.4.9 yields the results.

- Since $\operatorname{codim}(\mathcal{L}_{2n}^k) = k^2$ (see proposition II.1.6) we see that $\mathcal{L}_{2n}^0$ is an open submanifold with a codimension one complement $\mathcal{L}_{2n}^{\geq 1}$ so proposition II.4.9 does not apply and $C^1([0,1], \mathcal{L}_{2n}^0)$ is evidently not open (see remark II.4.12 for an example of an interior point in the complement). Though as $\mathcal{L}_{2n}^0 \subset \mathcal{L}_{2n}^{\leq k}$ for all $1 \leq k \leq n$ we see that each $\mathcal{L}_{2n}^{\leq k}$ must be dense as well so for each $1 \leq k \leq n$ we have,

$$\operatorname{codim}(\mathcal{L}_{2n}^{\leq k}) = \dim(\mathcal{L}_{2n}^{\geq k+1}) = (k+1)^2 > 1.$$

  and it follows from proposition II.4.9 that each $C^1([0,1], \mathcal{L}_{2n}^{\leq k})$ is open and dense in $C^1([0,1], \mathcal{L}_{2n})$ for $1 \leq k \leq n$.

□

Remark II.4.11. *The inclusion chains implied by the two statements of corollary II.4.10 above show that* $C^1([0,1], \mathcal{L}_{2n}^{\leq k})$ *are dense and open subsets of* $C^1([0,1], \Lambda_{2n})$ *for all* $1 \leq k \leq n$*.*

Remark II.4.12. *Very simple counter-examples exist to $C^1([0,1],\mathcal{L}^0_{2n})$ being open in $C^1([0,1],\mathcal{L}^{\leq 1}_{2n})$; consider the paths*

$$\gamma^{\pm}(t) = Diag_{2\times 2}\left(\pm(\log(1-t))^{-1}, \log(1-t)\right) \oplus Id_{2n-2} : [0,1) \to Sp(\mathbb{R}^{2n}) \cong \mathcal{L}^0_{2n}$$

with respect to a Darboux basis $(a_i, b_i)_{i=1}^n$. Next define the following paths of canonical relations: $\hat{\gamma}^{\pm} : [0,1] \to \mathcal{L}_{2n}$ for which $\hat{\gamma}^{\pm}(t) = \mathrm{Gr}(\gamma^{\pm}(t))$ and $\hat{\gamma}^{\pm}(1) = L = \langle (a_1,0),(0,b_1)\rangle \oplus \mathrm{Gr}(Id_{2n-2}(V_g))$ where $V_g = \langle (a_i,0),(0,b_i)\rangle_{i=2}^n$. Then $\beta = \gamma^- * \gamma^+ : I \to \mathcal{L}^{\leq 1}_{2n}$ with $\gamma^+(0) \in \mathrm{Sp}^+(V)$ and $\gamma^-(0) \in \mathrm{Sp}^-(V)$. This exhibits a path in the interior of $C^1([0,1],\mathcal{L}^{\leq 1}_{2n}) \setminus C^1([0,1],\mathcal{L}^0_{2n})$. Note that this behavior would potentially not occur in the event of compactifying $Sp(V)$ within the oriented Lagrangian Grassmannian as defined in [1].

Proposition II.4.13. *The subset $\mathcal{P}_{reg}(\mathcal{L}_{2n}) \subset C^1([0,1],\mathcal{L}_{2n})$ is open and dense.*

Proof. To show that $\mathcal{P}_{reg}(\mathcal{L}_{2n})$ is open and dense in $C^1([0,1],\mathcal{L}_{2n})$ consider the sets $D_k := \left\{ f \in C^1([0,1],\mathcal{L}^{\leq k}_{2n}) \mid f \pitchfork \mathcal{L}^k_{2n} \right\}$ of C^1 paths in $\mathcal{L}^{\leq k}_{2n}$ which are transverse to the k^{th} stratum $\mathcal{L}^k_{2n}$. Since $\mathcal{L}^k_{2n}$ is closed in $\mathcal{L}^{\leq k}_{2n}$ we see that $D_k \subset C^1([0,1],\mathcal{L}^{\leq k}_{2n}) \subset C^1([0,1],\mathcal{L}_{2n})$ is an inclusions of open and dense sets for all $1 \leq k \leq n$ by proposition II.4.7 (alternatively proposition II.4.9 would work as well).

Letting $D = \bigcap_{k=2}^n D_k$ be the set of paths transverse to each $\mathcal{L}^k_{2n}$ for $2 \leq k \leq n$, we note it is open and dense as well. Indeed, one may see that it is a finite intersection of open and dense sets or alternatively make the identification $D = C^1([0,1],\mathcal{L}^{\leq 1}_{2n})$ to reach the same conclusion via corollary II.4.10. While we could have considered the total intersection over all $1 \leq k \leq n$ and arrived at the desired result by now, we note that the $k = 1$ case is the only one in which

$$D_k = \left\{\gamma \in C^1(I, \mathcal{L}^{\leq k}) \,|\, \gamma \pitchfork \mathcal{L}_{2n}^k\right\} \neq \left\{\gamma \in C^1(I, \mathcal{L}_{2n}) \,|\, \mathrm{Im}(\gamma) \cap \mathcal{L}_{2n}^k = \emptyset\right\},$$

so that to ensure sufficient exposition we have decided on treating it separately from the cases $2 \leq k \leq n$. We first observe that $\mathcal{P}_{reg}(\mathcal{L}_{2n}) = D \cap D_1$ so that in particular, if $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$ then $\gamma \in D$ and therefore must have image wholly contained within $\mathcal{L}_{2n}^{\leq 1}$. Observing from the above identification we see $D_1 \subset D$ so that $D \cap D_1 = D_1$ we have,

$$\mathcal{P}_{reg}(\mathcal{L}_{2n}) = D_1 = \left\{\gamma \in C^1([0,1], \mathcal{L}_{2n}^{\leq 1}) \,|\, \gamma \pitchfork \mathcal{L}_{2n}^1\right\}.$$

Then since $\mathcal{L}_{2n}^1$ is closed in $\mathcal{L}_{2n}^{\leq 1}$, proposition II.4.7 implies (as in the above cases) that $D_1 = \mathcal{P}_{reg}(\mathcal{L}_{2n})$ is open and dense in $D = C^1([0,1], \mathcal{L}_{2n}^{\leq 1})$ and thus in $C^1([0,1], \mathcal{L}_{2n})$ with an application of corollary II.4.10. As mentioned above in remark II.4.11, this follows from the heredity of both properties with respect to the subspace topology. □

Now all that remains is a brief lemma intended to expedite the proof of our final claim that the intersection set of any stratum-regular path with a higher stratum is finite.

Lemma II.4.14. *Given M a smooth manifold, $A, B \subset M$ submanifolds with $\dim(A) + \dim(B) = \dim(M)$ and A compact, if $A \cap (\overline{B} \setminus B) = \emptyset$ then it must be that $|A \cap B| < \infty$.*

Proof. We suppose that $|A \cap B| = \infty$, so there exists a sequence of distinct points $\{p_i\}_{i=1}^{\infty} \subset A \cap B$ and since A is compact this sequence admits a convergent subsequence

$\{p_{i_j}\}_{j=1}^{\infty}$ for which $p_{i_j} \to p$ as $j \to \infty$ for some $p \in \overline{A \cap B} \subseteq A \cap \overline{B}$. We consider two cases; $p \in A \cap B$ and $p \in A \cap (\overline{B} \setminus B)$, the latter of the two reaching an immediate contradiction due to the hypothesis as $p \in \overline{B} \setminus B = \emptyset$ necessarily doesn't exist.

Now suppose $p \in A \cap B$ and consider any open neighborhood U about p in B. Then since $p_{i_j} \to p \in B$, we see for every open neighborhood U of p that there is some $N_U \in \mathbb{N}$ for which $\{p_{i_j}\}_{j=N_U}^{\infty} \subset U$, yet this clearly violates the fact that the intersection set is discrete as shown in proposition II.4.6. $\square$

Remark II.4.15. *One might describe $\overline{B} \setminus B$ as the points lying in the boundary of B yet not in B itself, in general being a proper subset of the boundary $\partial B = \overline{B} \setminus B^{\circ}$, failing to include those boundary points already contained in B. Even if $p \in \partial B$ in the above proof, since $p_i \in B$ for all $i \in \mathbb{N}$ we would still have a descending chain of open (in B) neighborhoods of p needed to complete the proof.*

Observe that a subset $B \subset M$ is closed if and only if $\overline{B} \setminus B = \emptyset$.

Proposition II.4.16. *For any stratum-regular path γ , the intersection* $\mathrm{Im}(\gamma) \cap (\Lambda_{2n} \setminus \mathrm{Im}(\mathrm{Gr}))$ *is a finite set in $\mathcal{L}_{2n}^{1}$.*

Proof. As we have established, any $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$ has $\gamma \in D_1 = C^1([0,1], \mathcal{L}_{2n}^{\leq 1})$, so that the second statement is true and we need only consider $\gamma : I \to \mathcal{L}_{2n}^{\leq 1}$. In this case, observe that the pair $A := \mathrm{Im}(\gamma)$ is compact and $B := \Lambda_{2n}^{1}$ is closed in $M := \Lambda_{2n}^{\leq 1}$ with complementary dimensions. Thus lemma II.4.14 implies that $|\,\mathrm{Im}(\gamma) \cap \Lambda_{2n}^{1}| < \infty$ for any $\gamma \in D_1$, and thus any $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$.

□

Part III

The Set H of Exceptional Lagrangians

III.1 The Codimension of H

Recall lemma II.1.4 which gives each stratum Λ_{2n}^k as a fiber bundle over $\mathcal{I}_k$.

Theorem I.4.1. The codimension of H in Λ_{2n} is two.

Proof. Recalling proposition II.1.6, that $\operatorname{codim}(\Lambda_{2n}^k) = k^2$, we see from $H \cap \Lambda_{2n}^0 = \emptyset$ that the trivial bound of $1 \leq \operatorname{codim}(H)$ holds. To sharpen this we consider the equivalence classes $[L] \subset \Lambda_{2n}^1$ induced by the associated isotropic pair class $[(\ker(L), \operatorname{halo}(L))] \subset I_1(V) \times I_1(V)$. There are three equivalence classes contained in $I_1(V) \times I_1(V)$ (and therefore three in Λ_{2n}^1) but the class $(\kappa, r, k) = (0, 0, 1)$ does not intersect H so that only the two classes satisfying $\kappa = k = 1$ need be checked, namely $\dim(\ker(L) \cap \operatorname{halo}(L)) = r = 0$ or $r = 1$.

Case I: r=1

Let $L \in [L]_1 = \{L \in \Lambda_{2n}^1 \,|\, \kappa(L) = r(L) = 1\}$. With the above lemma we first consider the associated class $[(\ker(L), \operatorname{halo}(L))]$ and write $\operatorname{dom}(L) = \operatorname{halo}(L) = \langle v \rangle$ for any $v \in V$. Then since v is arbitrary and all one dimensional subspaces are isotropic we see that $[(\langle v \rangle, \langle v \rangle)] = \triangle_{I_1(V)} \subset I_1(V) \times I_1(V) \cong \mathbb{RP}^{2n-1} \times \mathbb{RP}^{2n-1}$ implying that $\dim([(\langle v \rangle, \langle v \rangle)]) = 2n - 1$. Thus it follows that

$$\begin{aligned} \dim([L]_1) &= \dim(Sp(2n-2)) + \dim([(\langle v \rangle, \langle v \rangle)]) \\ &= 2n^2 - 3n + 1 + (2n - 1) \\ &= 2n^2 - n = \dim(\Lambda_{2n}) - 2n \end{aligned}$$

so that $\operatorname{codim}([L]_1) = 2n$ in Λ_{2n}.

Case II: r=0

When $L \in [L]_0 = \{L \in \Lambda^1_{2n} \,|\, r(L) = 0, \kappa(L) = 1\}$ we see that

$$[(\ker(L), \mathrm{halo}(L))] = \{(B_1, B_2) \,|\, B_1 \neq B_2, \text{ and } B_2 \leq B_1^\omega\}$$

where the first condition is due to $r(L) = 0$ and the second from $\kappa(L) = 1$. As before there are $2n - 1$ dimensions in freely choosing $B_1 = \langle v \rangle$ while the two conditions imply that $B_2 = \langle w \rangle \neq B_1$ is restricted to the punctured (not in the usual sense, maybe 'line-punctured' would be more descriptive) subspace $B_1^\omega \setminus B_1$. This subset descends under the quotient map $\mathbb{R}^{2n} \to \mathbb{RP}^{2n-1}$ to a punctured (in the usual sense now) projective hyperplane $[B_1^\omega] \setminus \{[B_1]\} \subset \mathbb{RP}^{2n-1}$ so that there are $2n - 2$ dimensions available when choosing B_2. This yields $\dim([(\ker(L), \mathrm{halo}(L))]) = 4n - 3$ so that

$$\dim([L]_0) = \dim(\mathrm{Sp}(2n-2)) + 4n - 3 = 2n^2 + n - 2 = \dim(\Lambda_{2n}) - 2.$$

Since these are the two equivalence classes in the stratum of minimal codimension which intersect H we see that $\mathrm{codim}(H) = \min(2, 2n) = 2$ for all $n \geq 1$. $\square$

III.2 Decomposing Linear Canonical Relations in the Complement of H

Even though the following proposition has already been taken as implicit, as it is a readily available consequence (e.g. as shown in [39]), we prove it here for completeness.

Proposition III.2.1. *The fiber over any $(B_1, B_2) \in \mathcal{I}_k$ is diffeomorphic to the symplectic group,*

$$(\Lambda^k_{2n})_{(B_1,B_2)} \cong \Lambda^0_{2n-2k}(B_1^\omega/B_1 \times B_2^\omega/B_2) \cong \mathrm{Sp}(2n-2k).$$

Proof. The first diffeomorphism was shown in [6] while the second follows after applying definition I.3.3 to the $4(n-k)$ dimension symplectic quotient,

$$(B_1^\omega \times B_2^\omega, \tilde{\omega}) \xrightarrow{q} \left(B_1^\omega/B_1 \times B_2^\omega/B_2,\, \omega_{red} := \pi_1^*\, \omega_1|_{B_1^\omega} - \pi_2^*\, \omega_2|_{B_2^\omega}\right).$$

See theorem III.2.5 for more details. □

Remark III.2.2. *In general, the diffeomorphism,*

$$(\Lambda^k_{2n})_{(B_1,B_2)} \cong \mathrm{Sp}(B_1^\omega/B_1, B_2^\omega/B_2),$$

is far from unique, depending on the identification $B_1^\omega/B_1 \cong B_2^\omega/B_2$. On the contrary we see in theorem III.2.5 that for any $L \notin H$ the two quotient sets coincide, with each possessing a canonical identification with the subspace $B_1^\omega \cap B_2^\omega$. In particular this allows us to give a unique diffeomorphism $(\Lambda^k_{2n})_{(\ker(L),\mathrm{halo}(L))} \cong \mathrm{Sp}(B_1^\omega \cap B_2^\omega)$ for any $L \notin H$.

Remark III.2.3. *When $L \in \mathcal{L}_{2n}$ we will call both the symplectic map ϕ_L and the Lagrangian subspace $\mathrm{Gr}(\phi_L) \leq V_g \times V_g$ the 'graph' portion of L when the context is unambiguous. We allow this abuse of notation as the above symplectic isomorphism is uniquely determined by $(\ker(L), \mathrm{halo}(L))$ which allows us to bypass the quotient construction as seen in proposition III.2.1 when representing the injective and coinjective portion $q(L)$ of L.*

Now we proceed in proving that the $L \notin H$ induce a unique ω-orthogonal decomposition of V which informally splits L into the direct sum of its 'singular' and 'graph' components.

Remark III.2.4. *We will proceed denoting $L_1 := \mathrm{dom}(L)$ and $L_2 := \mathrm{ran}(L)$ for a given $L \in \Lambda_{2n}$.*

Theorem III.2.5. *Given $L \in \mathcal{L}_{2n}$ (that is, $L_1 \cap L_2^\omega = \{0\}$) then we may define a unique symplectic decomposition $V = V_s \oplus V_g$ where $V_s = \ker(L) \oplus \mathrm{halo}(L)$ and $V_g = \mathrm{dom}(L) \cap \mathrm{ran}(L)$. Consequently each such L then yields the unique symplectic map $\phi\colon V_g \circlearrowleft$ for which*

$$L = (L_1^\omega \times \{0\} \oplus \{0\} \times L_2^\omega) \oplus Gr(\phi) \leq (V_s \times \overline{V_s}) \oplus (V_g \times \overline{V_g})$$

where the L_i^ω are transverse Lagrangian subspaces of V_s.

Proof. Given every $(v, w) \in L$ we may construct a map:

$$\begin{aligned} \tilde{\phi}\colon L_1 &\rightarrow L_2/L_2^\omega \\ v &\mapsto [w]. \end{aligned}$$

(i) This map is well defined.

Proof. Given $(v,w),(v,w') \in L$ we have that $[w]=[w'] \Leftrightarrow w-w' \in L_2^\omega$. $\square$

(ii) The kernel of $\tilde{\phi}$ is L_1^ω.

Proof. Since $\tilde{\phi}(v)=[w]=0 \Leftrightarrow (v,w) \in L,\ w \in L_2^\omega \Leftrightarrow (0,w) \in L$. Then by linearity we see $(v,0) \in L$ which by definition means $v \in L_1^\omega$. Conversely if $v \in L_1^\omega$ then $(v,0) \in L \Rightarrow \tilde{\phi}(v)=[0]$ so indeed $\ker(\phi)=L_1^\omega$.

Since the L_i are co-isotropic the map $\tilde{\phi}\colon L_1/L_1^\omega \to L_2/L_2^\omega$ is an isomorphism between symplectic vector spaces where for both $i=1,2$ we have that

$$\dim(L_i/L_i^\omega)=\dim(L_i)-\dim(L_i^\omega)=(2n-k)-k=2n-2k.$$

In fact since the L_i/L_i^ω are reduced co-isotropic subspaces they each possess a canonical symplectic form: $\omega_{red}^i([v],[v']) := \omega(v,v')$ for all $v,v' \in L_i$ which is independent of the choice of representatives $v,v' \in L_i$ for $i=1,2$. $\square$

(iii) The map $\phi\colon (L_1/L_1^\omega, \omega_{red}^1) \to (L_2/L_2^\omega, \omega_{red}^2)$ is symplectic.

Proof. Given any pair $(v,w),(v',w') \in L$ they must satisfy

$$\tilde{\omega}((v,w),(v',w'))=0 \Leftrightarrow \omega(v,v')=\omega(w,w').$$

It follows for any $[v],[v'] \in L_1/L_1^\omega$ and $[w],[w'] \in L_2/L_2^\omega$ such that $(v,w),(v',w') \in L$ that $\omega_{red}([v],[v'])=\omega_{red}([w],[w'])=\omega_{red}(\phi[v],\phi[v'])$ so ϕ is indeed a symplectic map between the two reduced spaces. $\square$

There always exists a pair of symplectic subspaces $V_i \leq L_i$ which are mapped bijectively under the projection maps $\pi_i \colon L_i \to L_i^\omega$ so that $V_i \cong L_i/L_i^\omega$ and $\phi \colon V_1 \to V_2$ satisfies $\phi^*\omega = \omega$. This map, although symplectic depends not only on L but which *pair* of V_i are chosen as well.

Until this point the hypothesis that $\kappa(L) = 0$ has not been needed but is now required to produce a unique $V_g := V_1 = V_2$ on which L induces $\phi \in \mathrm{Sp}(V_g) \cong \mathrm{Sp}(2n-2k)$, motivating the following lemma.

Lemma III.2.5.1. *If $\kappa(L) = 0$ then $V_g := (L_1 \cap L_2, \omega)$ is a symplectic subspace of V and there is a unique isomorphism $(V_g, \omega) \cong (L_i/L_i^\omega, \omega_{red}^i)$ for $i = 1, 2$.*

Proof. First note the following three conditions are sufficient to show the above: **(1)** $\dim(L_1 \cap L_2) = 2n - 2k$, **(2)** V_g is a symplectic subspace of V and **(3)** $\ker(\pi_i) = L_i^\omega$ has trivial intersection with V_g for $i = 1, 2$. if **(1)**-**(3)** are satisfied then $\pi_i|_{V_g}$ is a symplectomorphism for both $i = 1, 2$.

(1) $\dim(L_1 \cap L_2) = 2n - 2k$

Proof. Recalling that $\dim(L_i) = 2n - k$ for some $k \leq n$ meaning $\dim(L_i^\omega) = k$ and we see from $L_1^\omega \leq L_1$ that

$$\kappa(L) = \dim(L_1 \cap L_2^\omega) = 0 \Rightarrow \dim(L_1^\omega \cap L_2^\omega) = 0 \Leftrightarrow \dim(L_1^\omega \oplus L_2^\omega) = 2k$$

and $\dim(L_1 \cap L_2) = 2n - \dim(L_1^\omega \oplus L_2^\omega) = 2n - 2k$. □

(2) V_g is a symplectic subspace of V

Proof. We note $\kappa(L) = \dim(L_1 \cap L_2^\omega) = \dim(L_1^\omega \cap L_2) = 0$ implies via inclusion that

- $L_1 \cap L_2^\omega = \{0\} \Rightarrow (L_1 \cap L_2) \cap L_2^\omega = \{0\}$
- $L_2 \cap L_1^\omega = \{0\} \Rightarrow (L_1 \cap L_2) \cap L_1^\omega = \{0\}$

so that since $2n - 2k + 2k = 2n$ we have the following decomposition,

$$V = V_s \oplus V_g := (L_1^\omega \oplus L_2^\omega) \oplus (L_1 \cap L_2).$$

Indeed since $V_g^\omega = V_s$ and the above spans V we see that the two form a pair of complementary symplectic subspaces depending uniquely on L (or more precisely the isotropic pair associated to L).

□

(3) $V_g \cap L_i^\omega = \{0\}$ for $i = 1, 2$. This is an immediate consequence of the above two inclusions. □

Thus letting $V_g := L_1 \cap L_2$ we see since $(\pi_i^* \, \omega_{red})|_{V_g \times V_g} = \omega|_{V_g \times V_g}$ for $i = 1, 2$ there is precisely one $\phi \in \mathrm{Sp}(V_g)$ such that $\mathrm{Gr}(\phi) \leq V_g \times \overline{V_g}$ is Lagrangian. Conversely, for any isotropic pair (B_1, B_2) and symplectic map $\phi \in \mathrm{Sp}(B_1^\omega \cap B_2^\omega)$ there is a unique $L \in (\Lambda_{2n}^k)_{(B_1, B_2)}$ (i.e. $\mathrm{Pr}_I(L) = (B_1, B_2)$) for which the above construction yields the graph part (see remark III.2.3) $\mathrm{Gr}(\phi) \leq V_g \times V_g$ (where $V_g = B_1^\omega \cap B_2^\omega$ as before).

As to the final claim, since the L_i^ω are isotropic in V of dimension k it follows that the L_i^ω are transverse maximal isotropic subspaces of V_s and therefore Lagrangian subspaces of V_s. Transversality follows from $L \notin H$. □

Remark III.2.6. *Note that the bundle structure of each stratum Λ^k_{2n} are all derived from the underlying fixed symplectic vector space, in particular the fibers are smoothly dependent on the base point; given any $(B_1, B_2) \in \mathcal{I}_k$, the following symplectic vector space is determined uniquely by the base point $(B_1, B_2) \in I_k(V) \times I_k(V)$,*

$$(B_1^\omega/B_1 \times B_2^\omega/B_2, \tilde{\omega}_{red} = pr_1^*\omega_{1,red} - pr_2^*\omega_{2,red})\,.$$

It is this symplectic vector space over which the fiber is defined: $\Lambda^0_{2n-2k}(B_1^\omega/B_1 \times B_2^\omega/B_2) \cong \mathrm{Sp}(2n-2k)$ and since $\tilde{\omega}_{red}$ is the co-isotropic reduction of $\tilde{\omega}|_{B_1^\omega \times B_2^\omega}$, we see for fixed k that the quotient spaces and their induced symplectic form over which the fibers are defined vary smoothly over the base space as subspaces of the ambient vector space.

Consequently when $B_1 \pitchfork B_2^\omega$ we have the above canonical identifications so that certainly $B_1^\omega \cap B_2^\omega$ depends smoothly on the isotropic pair (B_1, B_2) as an element in the $2n-2k$ dimensional symplectic Grassmannian of V.

Part IV

The Conley-Zehnder Index and the Circle Map ρ

IV.1 Construction of the Index Using the Circle Map ρ

In addition to the mean index, the circle function ρ (as defined in definition I.3.7) is more often seen (as in [78]) during constructions of the Conley-Zehnder index.

Definition IV.1.1. *[78] We define $\rho\colon \mathrm{Sp}(2n) \to S^1$ as follows. Given $A \in \mathrm{Sp}(2n)$ let $E = Spec(A) \cap (S^1 \cup \mathbb{R})$ be the collection of real and elliptic eigenvalues of A. For elliptic eigenvalues $\lambda \in E \cap (S^1 \setminus \{\pm 1\})$ define $m^+(\lambda)$ to be the number of positive eigenvalues of the symmetric, non-degenerate two form Q defined on each complex eigenspace E_λ where*

$$\begin{aligned} Q\colon E_\lambda \times E_\lambda &\to \mathbb{R} \\ (z, z') &\mapsto \operatorname{Im} \omega(z, \overline{z'})). \end{aligned}$$

Then letting m^- denote the sum of the algebraic multiplicities for the real negative eigenvalues we have

$$\rho(A) := (-1)^{\frac{1}{2}m^-} \prod_{\lambda \in S^1 \setminus \{\pm 1\}} \lambda^{\frac{1}{2}m^+(\lambda)}.$$

In our case since eigenvalues are unique the term $m^+(\lambda) = 1$. The factor of $\frac{1}{2}$ amounts to a consistent way of choosing a single eigenvalue from each elliptic eigenvalue pair (while still expressing the product over all elliptic eigenvalues).

Proposition IV.1.2. *[78] The map $\rho\colon \mathrm{Sp}(2n) \to S^1$ has the following properties:*

1. *(determinant)* For $A \in U(n) \subset \mathrm{Sp}(2n)$ we have $\rho(A) = Det_{\mathbb{C}}(A)$.

2. *(invariance)* ρ is invariant under conjugation,

$$\rho(B^{-1}AB) = \rho(A),\ \forall B \in \mathrm{Sp}(2n).$$

3. *(normalization)* $\rho(A) = \pm 1$ if A has no elliptic eigenvalues.

4. *(multiplicativity)* If

$$A = B \oplus C \in \mathrm{Sp}(2n) \times \mathrm{Sp}(2m) \subset \mathrm{Sp}(2(n+m)),$$

   then $\rho(A \oplus B) = \rho(A)\rho(B)$.

5. *(homogeneity)* If $A \in \mathrm{Sp}(2n)$ we have that $\rho(A^l) = l \cdot \rho(A) \in \mathbb{R}/\mathbb{Z}$ for any $l \in \mathbb{Z}$.

All of the above properties are inherited by $\hat{\rho}$ when $L \in \Lambda^0_{2n}$ but for non-graph Lagrangian subspaces some properties no longer have an analog. In example VII.1.2 we see that $\hat{\rho}$ is indeed not a circle map on $\Lambda_2 \setminus H$ since there exists a non-contractible loop $\gamma : I \to \Lambda_2 \setminus H$ for which $\hat{\Delta}(\gamma) = 0$.

IV.2 Properties of the Extension $\hat{\rho}$

Despite the above observation, $\hat{\rho}$ does inherit some of the above properties with the caveat that they are all only defined for $L \in \Lambda^0_{2n}$ and are essentially trivial. One non-trivial property is found in the following lemma.

Lemma IV.2.1.

1. *The operation $(*)^l$ (definition I.3.1) is a continuous map for any $l \geq 0$.*

2. *(Homogeneity) Given any $L \in \mathcal{L}_{2n}$ with graph part $\mathrm{Gr}(\phi)$ and assuming theorem V.3.1 is true, that is $\hat{\rho}$ is continuous and $\hat{\rho}(L) = \rho^2(\phi)$, then $\hat{\rho}(L^l) = l \cdot \hat{\rho}(L)$ for all $l \geq 0$.*

Proof 1. When $l > 0$, continuity follows on each $\mathcal{L}_{2n}^k$ after referring to lemma II.1.4, which gives the fibration $\mathrm{Sp}(2n-2k) \hookrightarrow \mathcal{L}_{2n}^k \twoheadrightarrow \mathcal{I}_k$ for each $1 \leq k \leq n-1$. Since L, L^l share the same domain and range for all $l \in \mathbb{N}$ then $(*)^l$ preserves the fibers when $l \geq 1$. Since the fibers vary smoothly over the base we see that $(*)^l$ is continuous since it is continuous on each fiber, inheriting the group operation on $\mathrm{Sp}(V_g)$. When $k = 0$ this operation corresponds to the group operation in $\mathrm{Sp}(2n)$ and when $k = n$ it is the identity map so that $(*)^l$ is continuous on $\mathcal{L}_{2n}$ for $l > 0$.

When $l = 0$ and $0 \leq k \leq n$ the map $(*)^0$ is no longer a bundle map as it carries every $L \in \Lambda_{2n}^k$ to $L^0 = \mathrm{Gr}(Id_V)$. Regardless, since the image of each stratum of $(*)^0$ is the single point $\mathrm{Gr}(Id_V) = \triangle_V$ we see that $(*)^0$ is trivially continuous. $\square$

Proof 2. When $l > 0$, as shown in lemma II.2.2, the iterated composition operation restricted to $\mathcal{L}_{2n}$ is a well defined map and we may compute in coordinates L^2 where

$$L = (\ker(L) \times \{0\} \oplus \{0\} \times \mathrm{halo}(L)) \oplus \mathrm{Gr}(\phi)$$

so that we may verify the claim via a Darboux basis adapted to $V = \ker(L) \oplus \mathrm{halo}(L) \oplus (\mathrm{dom}(L) \cap \mathrm{ran}(L))$.

Let $\langle v_i \rangle_{i=1}^k = \ker(L)$ and $\langle w_i \rangle_{i=1}^k = \mathrm{halo}(L)$. Then since $(v_i, 0), (0, w_j) \in L$ for all $i, j \leq k$ it follows that $(v_i, w_j) \in L \circ L$ for all $i, j \leq k$. Namely, $\ker(L) \times \{0\} \oplus \{0\} \times$

$\mathrm{halo}(L) \leq L \circ L$. Next consider some $(v, w) \in \mathrm{Gr}(\phi)$ so that since $\phi \in \mathrm{Sp}(V_g)$ for any $w \in V_g$ there exists a unique $z \in V_g$ for which $(w, z) \in \mathrm{Gr}(\phi)$ implying $(v, z) \in L \circ L$. Intuitively this states that $\mathrm{Gr}(\phi) \circ \mathrm{Gr}(\phi) = \mathrm{Gr}(\phi^2)$. As for $(0, w_i) \in \mathrm{halo}(L)$ in the first L and $(v_i, 0) \in \ker(L)$ belonging to the second L the only resulting vector derived from these in the product is $(0, 0)$ (regardless of whether $L \in H$ or not) so we see for any $l \geq 1$ that

$$L^l := (\ker(L) \times \{0\} \oplus \{0\} \times \mathrm{halo}(L)) \oplus \mathrm{Gr}(\phi^l)$$

and $\hat{\rho}(L^l) = l \cdot \hat{\rho}(L)$ since ρ is homogeneous on $\mathrm{Sp}(V_g)$.

When $l = 0$ we have $\hat{\rho}(L^0) = \hat{\rho}(\mathrm{Gr}(\mathrm{Id})) = 0 \in \mathbb{R}/\mathbb{Z}$, so homogeneity holds in this case too. □

Proposition IV.2.2. *The map* $\hat{\rho} \colon \mathcal{L}_{2n} \to S^1$ *inherits the following properties:*

1. *(determinant) For* $A \in U(n) \subset \mathrm{Sp}(2n)$ *we have that* $\hat{\rho}(\mathrm{Gr}(A)) = (\det_{\mathbb{C}}(A))^2$.

2. *(invariance)* $\hat{\rho}$ *is invariant under conjugation on* Λ^0_{2n},

$$\hat{\rho}(\mathrm{Gr}(B^{-1}AB)) = \hat{\rho}(\mathrm{Gr}(A)) \; \forall B \in \mathrm{Sp}(2n).$$

3. *(normalization)* $\hat{\rho}(\mathrm{Gr}(A)) = 1$ *if* A *has no elliptic eigenvalues.*

4. *(multiplicativity) If* $L = K_1 \oplus K_2 \in \mathcal{L}_{2n} \times \mathcal{L}_{2m} \subset \mathcal{L}_{2(n+m)}$ *then* $\hat{\rho}(K_1 \oplus K_2) = \hat{\rho}(K_1)\hat{\rho}(K_2)$.

(4). We begin by observing when $L = K_1 \oplus K_2 \in \mathcal{L}_{2n} \times \mathcal{L}_{2m}$ that there exists a pair of symplectic bases, $(x_i, y_i)_{i=1}^n$ and $(u_i, v_i)_{i=1}^m$ which together form subspaces which

form a decomposition isomorphic to the induced one: $\mathbb{R}^{2n} \oplus \mathbb{R}^{2m} = \mathbb{R}^{2(n+m)}$. With this decomposition we may write

$$L = (\ker(K_1) \times \{0\} \oplus \{0\} \times \mathrm{halo}(K_1) \oplus \mathrm{Gr}(\phi_{K_1}))$$

$$\oplus$$

$$(\ker(K_2) \times \{0\} \oplus \{0\} \times \mathrm{halo}(K_2) \oplus \mathrm{Gr}(\phi_{K_2})).$$

Since $\mathrm{dom}(K_1), \mathrm{ran}(K_1) \leq \langle x_i, y_i \rangle_{i=1}^n$ and $\mathrm{dom}(K_2), \mathrm{ran}(K_2) \leq \langle u_i, v_i \rangle_{i=1}^m$ then theorem III.2.5 implies the two subspaces V_s^j and V_g^j determined by each of the K_1 and K_2 have pairwise trivial intersection thereby refining the decomposition,

$$\mathbb{R}^{2(n+m)} = (V_s^1 \oplus V_g^1) \oplus (V_s^2 \oplus V_g^2) = \mathbb{R}^{2n} \oplus \mathbb{R}^{2m}.$$

This implies $\ker(L) = \ker(K_1) \oplus \ker(K_2)$, $\mathrm{halo}(L) = \mathrm{halo}(K_1) \oplus \mathrm{halo}(K_2)$ and $\mathrm{Gr}(\phi_L) = \mathrm{Gr}(\phi_{K_1}) \oplus \mathrm{Gr}(\phi_{K_2})$, in particular this translates to $\phi_L = \phi_{K_1} \oplus \phi_{K_2} \in \mathrm{Sp}(2n) \times \mathrm{Sp}(2m) \subset \mathrm{Sp}(2(n+m))$. To conclude the proof we assume theorem V.3.1 is true so that $\hat{\rho}$ is multiplicative and $\hat{\rho}(L) = \rho^2(\phi_L) = \rho^2(\phi_{K_1})\rho^2(\phi_{K_2}) = \hat{\rho}(K_1)\hat{\rho}(K_2)$.

$\square$

Part V

Unbounded Sequences in the Symplectic Group

V.1 A Sufficient Condition for Asymptotic Hyperbolicity

The following theorem states that any $A \in \mathrm{Sp}(2n)$ with $\mathrm{Gr}(A)$ sufficiently near $L = L_1 \times \{0\} \oplus \{0\} \times L_2$ with $L_1 \pitchfork L_2$ has only hyperbolic eigenvalues.

Theorem V.1.1. *Suppose $\{A_i\}_{i=1}^{\infty} \subset \mathrm{Sp}(2n)$ is a sequence of symplectomorphisms each with distinct eigenvalues such that $\mathrm{Gr}(A_i) \underset{i\to\infty}{\longrightarrow} L_1 \times \{0\} \oplus \{0\} \times L_2 \in \mathcal{L}_{2n}$, $L_i \in \Lambda_n$ and $L_1 \cap L_2 = \{0\}$. Then there exists $K \in \mathbb{N}$ such that*

$$Spec(A_i) \cap (S^1 \setminus \{\pm 1\}) = \emptyset \tag{V.1.1}$$

for all $i \geq K$.

Proof. Suppose there exists some $\lambda_i \in Spec(A_i)$ such that $\lambda_i \in S^1 \setminus \{\pm 1\}$ for all $i \geq K \in \mathbb{N}$. Then since there are no multiple roots such a λ_i belongs to a unique symplectic eigenvalue pair $\left\{\lambda_i, \overline{\lambda_i}\right\}$ with $|\lambda_i| = 1$ which are stable in the sense that they remain elliptic away from the exceptional points ± 1 since the points ± 1 are the only values at which a (unique) elliptic eigenvalue pair may become hyperbolic or in general meet another pair to form a quadruple upon passing to the limit (this is certainly not true without uniqueness). For more details on how to go about 'ordering' the eigenvalues of a sequence of maps so that individual sequences of eigenvalues may be coherently formed (as done above implicitly) refer to section V.2.

This pair of eigenvalues has eigenvectors $x_i \pm i y_i \in \mathbb{C}^{2n}$ with a convergent subsequence of corresponding real eigenspaces $E_i = \langle x_i, y_i \rangle$ with some subsequence yielding

$E_i \to E$ as $i \to \infty$ (due to the compactness of the symplectic Grassmannian), on which A_i is conjugate to a rotation for all $i \in \mathbb{N}$. We obtain individual limit vectors by letting $x := \lim_{i\to\infty} \frac{x_i}{|x_i|}$ and $y = \lim_{i\to\infty} \frac{y_i}{|y_i|}$ so we may write $\langle x, y\rangle = E$.

The E_i are symplectic so there exists a decomposition $V = E_i \oplus F_i$ where the F_i form a sequence of some symplectic complements to each E_i and $A_i = \psi_i \oplus \phi_i \colon E_i \oplus F_i \to E_i \oplus F_i$ where ψ_i and ϕ_i are symplectic for each $i \in \mathbb{N}$.

Lemma V.1.2. *The Limit Lagrangian's Kernel and Halo.*

Let $\{A_i\}_{i=1}^{\infty} \subset \mathrm{Sp}(2n)$ denote a sequence of symplectic maps and consider the following;

1. *For $\{A_i\}_{i=1}^{\infty}$ such that $\mathrm{Gr}(A_i) \to L \in \mathcal{L}_{2n} \setminus \mathcal{L}_{2n}^0$ then $v \in \ker(L) \Leftrightarrow A_i v \to 0$ and $v \in \mathrm{halo}(L) \Leftrightarrow A_i^{-1} v \to 0$ as $i \to \infty$.*

2. *For $\{A_i\}_{i=1}^{\infty}$ such that such that $\mathrm{Gr}(A_i) \to L = L_1 \times \{0\} \oplus \{0\} \times L_2$ it is true that $L_1 \cap A_i^{-1} L_2 = \{0\}$ for sufficiently large i.*

Note that these both hold regardless of whether $L \in H$ or not.

Proof. We prove the first claim of each part, after which the lemma follows via contradiction:

1. $v \in \ker(L) \Leftrightarrow A_i v \to 0$. We observe $v \in \ker(L)$ if and only if $(v, 0) \in L$ so that since $\mathrm{Gr}(A_i) \to L$ it must be that $(v, A_i v) \to (v, 0)$.

2. $v \in \operatorname{halo}(L) \Leftrightarrow A_i^{-1}v \to 0$. Again $v \in \operatorname{halo}(L)$ if and only if $(0, v) \in L$. Yet $(A_i^{-1}v, v) \in \operatorname{Gr}(A_i)$ so as above we see that $v \in \operatorname{halo}(L) \Leftrightarrow A_i^{-1}v \to 0$.

We now prove the second claim in each part of the lemma: given some sequence $\{A_i\}_{i=1}^{\infty} \subset \operatorname{Sp}(2n)$ for which $\operatorname{Gr}(A_i) \to L \in \mathcal{L}_{2n}^n$ and suppose there exists a sequence $\{v_i\}_{i=1}^{\infty} \subset V$ and some $N \in \mathbb{N}$ for which $v_i \to v \neq 0$ and $v_i \in L_1 \cap A_i^{-1}L_2$ for all $i \geq N$.

Then for any $i \geq N$ we have $v_i \in L_1$ implies $|A_i v_i| \to 0$ and $v_i \in A_i^{-1}L_2$ implies that $w_i := \frac{A_i v_i}{|A_i v_i|} \in L_2$ and therefore $|A_i^{-1}w_i| \to 0$. This means that

$$|A_i^{-1}w_i| = \frac{|v_i|}{|A_i v_i|} \to 0$$

and since $|v_i| \to |v| \neq 0$ we have the contradiction $|A_i v_i| \to \infty$. $\square$

Continuing on from the above we now know for large i that the A_i induce a sequence $L_1 \oplus A_i^{-1}L_2 = V$ of Lagrangian splittings. In that case there exists a unique decomposition for any sequence $u_i = v_i + w_i \in E_i$ with $v_i \in L_1$ and $w_i \in A_i^{-1}L_2$ such that $|A_i v_i| \to 0$ and $|A_i w_i| \to \infty$. Denote $\psi_i := A_i|_{E_i}$ so that $|\psi_i v_i| \to 0$ and $|\psi_i w_i| \to \infty$ and consider the following lemma.

Lemma V.1.3. *Given a sequence of elliptic eigenspaces $E_i \to E$ and symplectic maps $\{\psi_i\}_{i=1}^{\infty}$ as above then $E \cap L_1 \neq \{0\}$.*

Proof. Convergence of the E_i is a consequence of the fact that $\dim(E_i)$ is constant for all i and therefore converges to E along some subsequence [47]. Suppose $E \cap L_1 = \{0\}$, that is every sequence $u_i = v_i + w_i \in E_i$ has $w_i \not\to 0$. We would have for every

$u_i = v_i + w_i \in E_i$ that $|\psi_i u_i| = |\psi_i v_i + \psi_i w_i| = |\psi_i w_i - (-\psi_i v_i)| \geq ||\psi_i w_i| - |\psi_i v_i|| \to \infty$ and since $|\psi_i v_i| \to 0$ we see $|\psi_i u_i| \to \infty$ for any sequence $u_i \in E_i$ with $w_i \not\to 0$.

Now $\dim(E_i) = 2$ so that $\omega_i := \omega|_{E_i \times E_i}$ is an area form on E_i for each $i \in \mathbb{N}$ so we may choose some sequence of balanced neighborhoods $U_i \subset E_i$ about zero on which

$$\int_{U_i} \omega_i = 1.$$

Note: These U_i may grow without bound but it is of no consequence since we will not pass to the limit.

Then we have for any $M > 1$ a K for which any normalized sequence $u_i \in E_i$, $u_i \to u \neq 0$ has $|\psi_i u_i| > M$ when $i \geq K$. This implies in particular that for each $M > 1$ there exists a $K \in \mathbb{N}$ for which $MU_i \subset \psi_i U_i$ for every $i \geq K$. It follows then that

$$1 = \int_{U_i} \omega_i < \int_{MU_i} \omega_i \leq \int_{\psi_i U_i} \omega_i$$

for every $i \geq K$. Each ψ_i is a symplectomorphism so we have a contradiction by choosing any $M > 1$, as some K exists for which

$$1 = \int_{U_K} \omega_K = \int_{\psi_K U_K} \omega_K \geq \int_{MU_K} \omega_K > 1. \tag{V.1.2}$$

□

This lemma leads us to the following lemma we need to prove theorem V.1.1.

Lemma V.1.4. *Given the eigenspaces $E_i \to E$ and $\{\psi_i\}$ as above then $E \cap L_2 \neq \{0\}$.*

Proof. We have already established there exists a sequence $v_i \to v$ for which $A_i v_i \to 0$ so we consider any $u_i = v_i + w_i \in E_i \cap (L_1 \oplus A_i^{-1} L_2)$ with $w_i \not\to 0$. Then $A_i u_i = A_i v_i + A_i w_i$ so that each E_i is A_i invariant and each $A_i w_i \in L_2 \cap E_i$. Then since $|A_i v_i| \to 0$ as $i \to \infty$ we see $\lim_{i\to\infty} \frac{A_i u_i}{|A_i w_i|} = \lim_{i\to\infty} \frac{A_i w_i}{|A_i w_i|} \to w \in \text{halo}(L) \cap E$. □

The above lemma and corollary show there exists a sequence of bases $\langle v_i, w_i \rangle \in E_i$ on each real elliptic eigenspace such that $v_i \to v \in L_1$ and $w_i \to w \in L_2$.

Lemma V.1.5. *Let E_i be a two dimensional real eigenspace for a complex eigenvalue $\lambda \in S^1 \setminus \{\pm 1\}$ which exists for sufficiently large $i \in \mathbb{N}$. Then given any sequence $\{v_i\}$ with each $v_i \in E_i$ and $\frac{v_i}{|v_i|} \to v \neq 0$ we claim,*

$$v \in \ker(L) \Leftrightarrow v \in \text{halo}(L).$$

Proof. Consider $\sigma_i \in \text{Sp}(E_i)$ such that $\psi_i = \sigma_i^{-1} \circ R(\theta_i) \circ \sigma_i$ where $R(\theta_i)$ denotes the rotation of the $(x_i', y_i') = (\sigma_i x_i, \sigma_i y_i)$ plane E_i by $\lambda_i = e^{i\theta_i}$ so that we may compute for any sequence $\{v_i\}_{i=1}^{\infty}$,

$$\begin{aligned} |\psi_i^{-1} v_i| = |(\sigma_i^{-1} \circ R(\theta_i) \circ \sigma_i)^{-1}(v_i)| &= |(\sigma_i^{-1} \circ R(\theta_i)^{-1} \circ \sigma_i)(v_i)| \\ &= |(\sigma_i^{-1} \circ R(-\theta_i) \circ \sigma_i)(v_i)| \end{aligned}$$

and we see ψ_i^{-1} is simply the opposite rotation of ψ_i conjugated by the same matrix $\sigma_i \in \text{Sp}(E_i)$. recalling the sequence $v_i \in E_i$, there exists $a_i, b_i \in \mathbb{R}$ such that $v_i =$

$a_i x_i + b_i y_i$ so if we let $v_i' = \sigma_i(v_i)$ then $v_i' := \sigma_i v_i = a_i x_i' + b_i y_i'$ and thus

$$\begin{aligned}
\psi_i v_i &= (\sigma_i^{-1} \circ R(\theta_i) \circ \sigma_i)(v_i) \\
&= (\sigma_i^{-1} \circ R(\theta_i))(v_i') \\
&= (\sigma_i^{-1} \circ R(\theta_i))(a_i x_i' + b_i y_i') \\
&= \sigma_i^{-1}\left(a_i \cos(\theta_i) x_i' + a_i \sin(\theta_i) y_i'\right) + \sigma_i^{-1}\left(b_i \cos(\theta_i) y_i' - b_i \sin(\theta_i) x_i'\right) \\
&= \left(a_i \cos(\theta_i) - b_i \sin(\theta_i)\right) x_i + \left(a_i \sin(\theta_i) + b_i \cos(\theta_i)\right) y_i
\end{aligned}$$

and with squared norm

$$\begin{aligned}
|\psi_i v_i|^2 &= \left(a_i^2 \cos^2(\theta_i) - 2a_i b_i \sin(\theta_i)\cos(\theta_i) + b_i^2 \sin^2(\theta_i)\right) \\
&\quad + \left(b_i^2 \cos^2(\theta_i) + 2a_i b_i \sin(\theta_i)\cos(\theta_i) + a_i^2 \sin^2(\theta_i)\right) \\
&= a_i^2 + b_i^2 \to 0 \Leftrightarrow \lim_{i\to\infty} \frac{v_i}{|v_i|} \in \ker(L).
\end{aligned}$$

On the other hand

$$\begin{aligned}
\psi_i^{-1} v_i &= (\sigma_i^{-1} \circ R(\theta_i)^{-1})(v_i') \\
&= (\sigma_i^{-1} \circ R(-\theta_i))(a_i x_i' + b_i y_i') \\
&= \sigma_i^{-1}\left(a_i \cos(\theta_i) x_i' - a_i \sin(\theta_i) y_i'\right) + \sigma_i^{-1}\left(b_i \cos(\theta_i) y_i' + b_i \sin(\theta_i) x_i'\right) \\
&= \left(a_i \cos(\theta_i) + b_i \sin(\theta_i)\right) x_i + \left(b_i \cos(\theta_i) - a_i \sin(\theta_i)\right) y_i
\end{aligned}$$

so that

$$|\psi_i^{-1} v_i|^2 = a_i^2 + b_i^2 \to 0 \Leftrightarrow \lim_{i\to\infty} \frac{v_i}{|v_i|} \in \operatorname{halo}(L)$$

and it's evident that $|\psi_i v_i|^2 = |\psi_i^{-1} v_i|^2$. Since the last conclusions above follow from lemma V.1.2 we see for any sequence $v_i \in E_i$ where $v = \lim_{i\to\infty} \frac{v_i}{|v_i|}$ that

$v \in \ker(L) \Leftrightarrow v \in \mathrm{halo}(L)$ where the E_i are a sequence of two dimensional elliptic eigenspace which persists for arbitrarily large i (by hypothesis) in a sequence of symplectic maps $\{A_i\}_{i=1}^{\infty}$. $\square$

We see now any sequence $\{A_i\}_{i=1}^{\infty} \subset \mathrm{Sp}(2n)$ for which $\mathrm{Gr}(A_i) \to L \in \mathcal{L}_{2n}^{n}$ has only real pairs or the usual symplectic quadruples away from the unit circle for sufficiently large i. $\square$

V.2 Decomposing Certain Unbounded Sequences of Symplectic Maps

Theorem V.2.1. *Consider any sequence $\{A_i\}_{i=1}^{\infty} \subset \mathrm{Sp}(V)$ where each A_i has distinct eigenvalues and for which*

$$\mathrm{Gr}(A_i) \underset{i\to\infty}{\to} L = \ker(L) \times \{0\} \oplus \{0\} \times \mathrm{halo}(L) \oplus \mathrm{Gr}(\phi) \in \mathcal{L}_{2n}^{k},$$

where the graph part (see remark III.2.3) ϕ of L has semisimple eigenvalues.

Then for sufficiently large $i \in \mathbb{N}$, there exists some subsequence of unique, A_i invariant symplectic decompositions $V = E_s^i \oplus E_g^i$ with which we may write $A_i = \alpha_i \oplus \beta_i \in \mathrm{Sp}(E_s^i) \times \mathrm{Sp}(E_g^i)$ and for which $\mathrm{Gr}(\alpha_i) \to \ker(L) \times \{0\} \oplus \{0\} \oplus \mathrm{halo}(L)$.

Additionally there exists an $N \in \mathbb{N}$ for which there is a sequence of symplectic isomorphisms

$$\left\{ I_i \colon (E_g^i, \omega|_{E_g^i \times E_g^i}) \overset{\cong}{\to} (V_g, \omega|_{V_g \times V_g}) \right\}_{i=N}^{\infty},$$

uniquely determined by L such that each $\beta_i\colon E_g^i \to E_g^i$ is conjugate via I_i to some $\phi_i \in \mathrm{Sp}(V_g)$ for all $i \geq N$ with $\phi_i \to \phi$. We also show that the β_i preserve the data used in computing ρ, namely the eigenvalues and the conjugacy classes of the A_i restricted to elliptic eigenspaces.

Outline. We have broken the proof into the five following lemmas which together give the desired result. Recall from theorem III.2.5 that L induces the symplectic splitting $V_s \overset{\omega}{\oplus} V_g = V$ for which $L|_{V_s \times V_s} = \ker(L) \times \{0\} \oplus \{0\} \times \mathrm{halo}(L)$ and $L|_{V_g \times V_g} = \mathrm{Gr}(\phi)$. We will be working over a subsequence of the A_i but for brevity's sake will notate this subsequence as A_i.

Lemma V.2.2. For some $N \in \mathbb{N}$ there exists a subsequence of symplectic splittings $\left\{E_s^i, E_g^i\right\}_{i=N}^{\infty}$ of V by A_i-invariant subspaces E_s^i, E_g^i where $\dim(E_s^i) = \dim(V_s) = 2k$ and $\dim(E_g^i) = \dim(V_g) = 2n - 2k$ for all $i \geq N$. Additionally each subsequence converges;

$$E_s^i \underset{i\to\infty}{\to} E_s \text{ and } E_g^i \underset{i\to\infty}{\to} E_g.$$

such that $E_s = \ker(L) \oplus \mathrm{halo}(L)$.

Using the sequence above, we observe that for $i \geq N$, the maps have the following decomposition $A_i = \alpha_i \oplus \beta_i\colon E_s^i \times E_g^i \to E_s^i \times E_g^i$.

Lemma V.2.3. Given the above decomposition, we have $\mathrm{Gr}(\alpha_i) \to \ker(L) \times \{0\} \oplus \{0\} \times \mathrm{halo}(L)$ and $\mathrm{Gr}(\beta_i) \to \mathrm{Gr}(\beta)$ for some $\beta \in \mathrm{Sp}(E_g)$ as elements in the appropriate dimension isotropic Grassmannian of $V \times \overline{V}$.

Lemma V.2.4. There exists an $N \in \mathbb{N}$ such that for all $i \geq N$ the subsequence of graph portion domains $\{E_g^i\}_{i=N}^{\infty}$ has:

- $\ker(L) \cap E_g^i = \{0\}$
- $Proj_{\mathrm{halo}(L)}(E_g^i) = \{0\}$.

In addition, both properties persist in the limit; $\ker(L) \cap E_g = \{0\}$ and $Proj_{\mathrm{halo}(L)}(E_g) = \{0\}$.

Observe that the second claim of lemma V.2.4 implies $E_g^i, E_g \leq \ker(L) \oplus V_g = \mathrm{dom}(L)$ for all $i \geq N$.

Lemma V.2.5. For all $i \geq N$ there exists a unique sequence of symplectic isomorphisms $I_i \colon E_g^i \to V_g$ such that $I_i = Proj|_{E_g^i}$ and $I_i \underset{i\to\infty}{\to} I \colon E_g \to V_g$ where the function $Proj \colon \mathrm{dom}(L) \to \mathrm{dom}(L)/\ker(L) \cong V_g$ is the coisotropic reduced space of $\mathrm{dom}(L)$ uniquely identified by theorem III.2.5 with V_g.

Lemma V.2.6. Defining $\phi_i := I_i \circ \beta_i \circ I_i^{-1} \in \mathrm{Sp}(V_g)$ then $\phi_i \to \phi \colon V_g \to V_g$ where $\mathrm{Gr}(\phi)$ is the graph part of L. Then for sufficiently large i the pair ϕ_i and β_i share the same eigenvalues and each pair of elliptic eigenvalues quadruples have matching Krein type.

V.2.1 Prerequisites

We first recall that each $A \in \mathrm{Sp}(2n)$ yields a direct sum of V via symplectic generalized eigenspaces, that is

$$V = \bigoplus_{\lambda \in Spec(A) \cap D_2^+} E_{[\lambda]} \text{ where } D_2^+ = \{z \in \mathbb{C} \,|\, 0 < |z| \leq 1, im(z) \geq 0\}$$

where $E_{[\lambda]}$ is the real eigenspace associated to the quadruple $(\lambda, \lambda^{-1}, \bar{\lambda}, \bar{\lambda}^{-1})$ satisfying $E^{\mathbb{C}}_{[\lambda]} = E_\lambda \oplus E_{\lambda^{-1}} \oplus E_{\bar{\lambda}} \oplus E_{\bar{\lambda}^{-1}} \leq V^{\mathbb{C}}$ and E_λ denotes the (generalized) complex eigenspace associated to λ. Note that if λ is an eigenvalue of $A \in \mathrm{Sp}(V)$ then $\lambda \neq 0$, the specification that $\lambda \in D_2^+$ is simply a convenient way of picking a candidate from each quadruple as well as providing a unique limit point for unbounded eigenvalues (since the representative chosen from that quadruple tends to 0).

When A has distinct eigenvalues this further restricts the possibilities for the above eigenspaces; We have already seen that the $E_{[\lambda]}$ for $\lambda \in (S^1 \cup \mathbb{R}) \setminus \{0, \pm 1\}$ are real two dimensional symplectic subspaces on which $A|_{E_{[\lambda]}}$ is either conjugate to a rotation by $\lambda \in S^1$ or to a hyperbolic transformation for $\lambda \in \mathbb{R} \setminus \{0, \pm 1\}$. The eigenvalue quadruples with $|\lambda| \neq 1$ and $im(\lambda) \neq 0$ manifest as a pair of A invariant real eigenspaces associated to the conjugate pairs $(\lambda, \bar{\lambda})$ and $(\frac{1}{\lambda}, \frac{1}{\bar{\lambda}})$. A symplectic normal form for $A \in \mathrm{Sp}(2n)$ restricted to this 4 dimensional real vector space after picking some $\lambda = re^{i\theta}$ from a quadruple is given by the following with (x, x', y, y') a

Darboux basis,

$$\begin{pmatrix} R(re^{-i\theta})^{-1} & 0 \\ 0 & R(re^{-i\theta})^t \end{pmatrix} = \begin{pmatrix} \frac{1}{r}R(e^{i\theta}) & 0 \\ 0 & rR(e^{i\theta}) \end{pmatrix},$$

where

$$R(re^{i\theta}) = r \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix}.$$

so that each A_i is the direct sum of a combination of the above symplectic eigenspaces.

We proceed by considering the eigenvalues of each A_i for all $i \in \mathbb{N}$ as a sequence of tuples $(\lambda^i) := (\lambda^i_1, \ldots, \lambda^i_{2n}) \in \mathbb{C}^{2n}$ treated as an unordered list. The space of unordered $\mathbb{C}$ tuples of length $2n$ may be identified with the orbit space $\mathbb{C}^{2n}/S_{2n}$ where S_{2n} is the permutation group on $2n$ elements and the group action on $\mathbb{C}^{2n}$ is given by $\sigma((\lambda_1, \ldots, \lambda_{2n})) = (\lambda_{\sigma(1)}, \ldots, \lambda_{\sigma(2n)})$ for any permutation $\sigma \in S_{2n}$. Following [47] the topology induced on the space of unordered $\mathbb{C}$ tuples of length $2n$ as constructed above is identical to the one generated by the following metric,

$$d((\lambda), (\tau)) = min_{\sigma \in S_{2n}} max_{i \leq 2n} |\lambda_{\sigma(i)} - \tau_i| \tag{V.2.1}$$

with the helpful property that $\mathbb{C}^{2n}/S_{2n}$ is homeomorphic to $\mathbb{C}^{2n}$.

For convenience if we impose some ordering of the A_1 eigenvalues we may use a recursive process to yield an *essentially* unique representative for every subsequent element (since distinct permutations may both be a minimum in the above metric). Given any order for (λ^1) we choose the order of the i^{th} eigenvalue list $(\lambda^i) = (\lambda^i_1, \ldots, \lambda^i_{2n})$ for any $i \geq 2$ by choosing a permutation which minimizes the above metric with respect to the previous element, i.e. $(\lambda^i_{\tau(j)}) = (\lambda^{i-1}_j)$ where τ is the minimizing permutation chosen by the above minimax metric. This allows us to treat the tuple as ordered given the order of the first.

With this notion we let $\left\{\lambda^i_j\right\}_{j=1}^l$ denote the l representatives from each eigenvalue quadruple of A_i which lies in the closed upper half disc so that $V = \bigoplus_{j=1}^l E_{[\lambda^i_j]}$ for all $i \in \mathbb{N}$. As mentioned above each sequence $E_{[\lambda^i_j]}$ eventually has constant dimension for large i and by compactness each possesses a limit E_j for all $j \leq l$, potentially with lower dimension if distinct eigenvectors converge to each other in the limit. In our case the eigenvalues may converge but the dimension of the eigenspaces will be preserved due to the requirement that ϕ have semi-simple eigenvalues, precluding this possibility [47].

Now that we have established the prerequisite notions, we are ready to begin the proofs.

V.2.2 Proof

Lemma V.2.2. *For some $N \in \mathbb{N}$ there exists a subsequence of symplectic splittings $\left\{E^i_s, E^i_g\right\}_{i=N}^\infty$ of V by A_i-invariant subspaces E^i_s, E^i_g where $\dim(E^i_s) = \dim(V_s) = 2k$ and $\dim(E^i_g) = \dim(V_g) = 2n - 2k$ for all $i \geq N$. Additionally each subsequence converges;*

$$E^i_s \underset{i\to\infty}{\to} E_s = \ker(L) \oplus \operatorname{halo}(L) = V_s \text{ and } E^i_g \underset{i\to\infty}{\to} E_g.$$

Proof. We proceed by distinguishing two possibilities for the behavior of the sequence of A_i when restricted to each $E_{[\lambda_j^i]}$, either $||A||_{E_{[\lambda_j^i]}} = sup_{v \in E_{[\lambda_j^i]}} \frac{|A_i v|}{|v|} \to \infty$ or $||A||_{E_{[\lambda_j^i]}} = sup_{v \in E_{[\lambda_j^i]}} \frac{|A_i v|}{|v|} \to c_j \in \mathbb{R}$. In the first case this implies the existence of a sequence $\{v_i\}_{i=1}^{\infty}$ where $v_i \in E_{[\lambda_j^i]}$ for each i such that $v_i \to v \neq 0$ yet $\frac{|A_i v_i|}{|v_i|} \to \infty$ so that $|A_i v_i| \to \infty$. Then by setting $v_i' = \frac{v_i}{|A_i v_i|}$ then $v_i' \to 0$ and we see

$$\frac{|A_i v_i'|}{|v_i'|} = \frac{|A_i v_i|}{|v_i|} \to \infty. \tag{V.2.2}$$

Thus we see that $(v_i', A_i v_i') \in \mathrm{Gr}(A_i|_{E_{[\lambda_j^i]}})$ for each i and upon passing to the limit $(v_i', A_i v_i') = (v_i', \frac{A_i v_i}{|A_i v_i|}) \to (0, w) \in L$ where $\lim_{i \to \infty} \frac{A_i v_i}{|A_i v_i|} = w \in \mathrm{halo}(L)$ since the $\frac{A_i v_i}{|A_i v_i|}$ belong to a compact set. Recall the above metric and ordering scheme which allows us for some fixed j to identify a unique element λ_j^i for each $i \geq 2$ so that we may form a single sequence for each of the l eigenvalue quadruple representatives $\left\{\lambda_j^i\right\}_{i=1}^{\infty} \underset{i \to \infty}{\to} \lambda_j$ and the associated quadruple eigenspaces $\left\{E_{[\lambda_j^i]}\right\}_{i=1}^{\infty}$ (note that in the limit the eigenvalue is allowed to vanish). We define

$$S = \left\{ j \leq l \mid \ ||A_i|||_{E_{[\lambda_j^i]}} \underset{i \to \infty}{\to} \infty \right\} \subset \{1, 2, \ldots, l\}$$

which certainly satisfies $\{1, 2, \ldots, l\} = S \cup S^c$ so that we may define the symplectic A_i invariant subspaces based on this condition,

$$E_s^i := \bigoplus_{j \in S} E_{[\lambda_j^i]}, \qquad E_g^i := \bigoplus_{j \notin S} E_{[\lambda_j^i]}.$$

We know that $V = E_s^i \overset{\omega}{\oplus} E_g^i$ for each i and so $A_i = \alpha_i \oplus \beta_i \in \mathrm{Sp}(E_s^i) \times \mathrm{Sp}(E_g^i)$. Both subspaces must have constant dimension since the condition defining the two sets is binary and defined using asymptotically stable behavior of the A_i (namely, divergence of norm on eigenspaces and the dimensions of eigenspaces). Since they reside in a compact space and there are no eigenvectors converging to some single vector in the limit, there exists a subsequence for which both converge to E_s and E_g, symplectic subspaces of V. For $v \in E_s$ it's true that $\lim_{i\to\infty} |A_i v| \to 0, \infty$ and for $v \in E_g$ that $lim_{i\to\infty}|A_i v| \to d \in (0,\infty)$ so that certainly $E_s \cap E_g = \{0\}$ and since these are a complete set of eigenspaces we have $V = E_s \oplus E_g$.

Now we state an auxiliary lemma for lemma V.2.2.

Lemma V.2.2.1. *Equations* (V.1.2) *and* (V.2.2) *imply* $\ker(L) \oplus \mathrm{halo}(L) \leq E_s$.

If the above is true, we simply note that $\dim(E_g) \geq \dim(V_g) = 2n - 2k$ so that $\dim(E_s) = 2k$ and $\dim(E_g) = 2n - 2k$ and the first lemma of the proof is finished by the proceeding one.

$\square$

Proof. Lemma V.2.2.1. We may construct a sequence of isotropic subspaces $\langle w_j^i \rangle_{j=1}^k \leq E_s^i$ assuming each of the w_j^i converge such that

$$lim_{i\to\infty} \langle w_j^i \rangle_{j=1}^k = \langle w_j \rangle_{j=1}^k = \mathrm{halo}(L),$$

since each $w \in \mathrm{halo}(L)$ necessarily is in the limit of the E_s^i, otherwise $A|_{E_g^i}$ would not be bounded. Then since each E_s^i is symplectic there exists a sequence $(\tilde{w}_j^i) \leq E_s^i$ such that $\omega(w_j^i, \tilde{w}_j^i) = 1$ for all $j \leq k$ and all $i \in \mathbb{N}$. Then we consider for each $j \leq k$ the sequence of two dimensional symplectomorphisms $A_i|_{\langle w_j^i, \tilde{w}_j^i \rangle}$ on which $|A_i^{-1} w_j^i| \to 0$. Using the argument found in lemma V.1.3 as well as (V.1.2) and (V.2.2) the sequence $v_j^i := \frac{A_i^{-1}w_j^i}{|A_i^{-1}w_j^i|}$ satisfies $|A_i v_j^i| \to \infty$ and since $\omega|_{\langle w_j^i, \tilde{w}_j^i\rangle \times \langle w_j^i, \tilde{w}_j^i\rangle}$ is an area form preserved by $A_i|_{\langle w_j^i, \tilde{w}_j^i\rangle}$ for each i we must have for each $j \leq k$ some sequence $\{z_j^i\}_{i=1}^{\infty}$ where each $z_j^i \in E_s^i$ and $z_j^i \to z_j \neq 0$ such that $|A_i z_j^i| \to 0$ for all $j \leq k$. Thus $\langle z_j \rangle_{j=1}^k = \ker(L)$ implying $\ker(L) \leq E_s$ so that since $\ker(L) \cap \mathrm{halo}(L) = \emptyset$ we see that $\ker(L) \oplus \mathrm{halo}(L) \leq E_s$ and since $A_i|_{E_g^i}$ converges we see for dimensional reasons (as mentioned before the proof began) that $E_s = \ker(L) \oplus \mathrm{halo}(L)$. $\square$

Lemma V.2.3. *Given the above decomposition, we have* $\mathrm{Gr}(\alpha_i) \to \ker(L) \times \{0\} \oplus \{0\} \times \mathrm{halo}(L)$ *and* $\mathrm{Gr}(\beta_i) \to \mathrm{Gr}(\beta)$ *for some* $\beta \in \mathrm{Sp}(E_g)$ *as elements in the appropriate dimension isotropic Grassmannian of* $V \times \overline{V}$.

Proof. As we mentioned above there is a subsequence for which $E_g^i \to E_g$ so that since $\mathrm{Gr}(\alpha_i) \oplus \mathrm{Gr}(\beta_i) = \mathrm{Gr}(A_i) \to L$ we see that $\mathrm{Gr}(\beta_i) \to K_g \leq E_g \times E_g$. If $K_g \neq \mathrm{Gr}(\beta)$ for any $\beta \in \mathrm{Sp}(E_g)$ then $\ker(K_g) \neq \{0\}$ which violates the known dimension of $\ker(L)$. Indeed since $v \in \ker(K_g)$ implies $v \in \ker(L)$ when $K_g \leq L$, this shows that $K_g = \mathrm{Gr}(\beta)$ for some $\beta \in \mathrm{Sp}(E_g)$.

As for $\mathrm{Gr}(\alpha_i) = \{(v, A_i v) \,|\, v \in V_s^i\}$ we may write a convergent sequence of $2k$ tuples $(v_j^i, w_j^i)_{j=1}^k \subset E_s^i$ for all i such that $\langle v_j^i \rangle_{j=1}^k \to \ker(L)$ and $\langle w_j^i \rangle_{j=1}^k \to \mathrm{halo}(L)$. Then we see that $(v_j^i, A_i v_j^i) \to (v_j, 0) \in \ker(L) \times \{0\}$ and $(A_i^{-1} w_j^i, w_j^i) \to (0, w_j) \in \{0\} \times \mathrm{halo}(L)$ so that indeed $\mathrm{Gr}(\alpha_i) \to K_s = \ker(L) \times \{0\} \oplus \{0\} \times \mathrm{halo}(L)$. $\square$

Lemma V.2.4. *There exists an* $N \in \mathbb{N}$ *such that for all* $i \geq N$ *the subsequence of graph portion domains* $\{E_g^i\}_{i=N}^{\infty}$ *has:*

- $\ker(L) \cap E_g^i = \{0\}$

- $Proj_{\text{halo}(L)}(E_g^i) = \{0\}$.

In addition, both properties persist in the limit; $\ker(L) \cap E_g = \{0\}$ *and* $Proj_{\text{halo}(L)}(E_g) = \{0\}$.

Proof. We begin by proving the following lemma,

Lemma V.2.4.1. *Given* A_i *as before with* $\text{dom}(L) \pitchfork \text{halo}(L)$ *and* $j \leq l$ *fixed. Then we claim* $|| \, A_i|_{E_{[\lambda_j^i]}} \, ||$ *is bounded as* $i \to \infty$ *if and only if there exists an* $N \in \mathbb{N}$ *for which* $Proj_{\text{halo}(L)}(E_{[\lambda_j^i]}) = \{0\}$ *for all* $i \geq N$.

Proof. First suppose $F_i := Proj_{\text{halo}(L)}(E_{[\lambda_j^i]})$ is a sequence of subspaces such that $F_i \neq \{0\}$ for all but finitely many $i \in \mathbb{N}$. Then there exists a sequence $\{w_i\} \in E_{[\lambda_j^i]}$ with $w_i \to w \neq 0$ so that since $V = \text{halo}(L) \oplus \text{dom}(L)$ we may write $w_i = f_i + g_i \to f + g = w$ where each $f_i \in \text{halo}(L)$ and $g_i \in \text{dom}(L)$. Then since $|A_i^{-1} f_i| \to 0$ we may define

$$v_i := \frac{A_i^{-1}(f_i + g_i)}{|A_i^{-1} f_i|} = A_i^{-1} \left(\frac{f_i + g_i}{|A_i^{-1} f_i|} \right) \in E_{[\lambda_j^i]}$$

and see that $|A_i v_i| = \frac{|f_i + g_i|}{|A_i^{-1} f_i|}$. Since $g_i \to g$ with $|g| < \infty$ and the same for f_i it must be $|A_i v_i| \to \infty$ and $|| \, A_i|_{E_{[\lambda_j^i]}} \, ||$ is unbounded.

Alternatively if for some $N \in \mathbb{N}$ for which $Proj_{\text{halo}(L)}(E_{[\lambda_j^i]}) = \{0\}$ for all $i \geq N$ we know $E_{[\lambda_j^i]} \leq \text{dom}(L)$ and therefore any converging sequence $v_i \in E_{[\lambda_j^i]}$ may be uniquely written as $v_i = k_i + g_i \in \ker(L) \oplus V_g = \text{dom}(L)$ for any $i \geq N$ where

$k_i + g_i \to k + g = v$. Then since $|A_i k_i| \to 0$ and $A_i g_i \to \phi(g) \in V_g$ by definition, we see that the operator norm of A_i over $E_{[\lambda_j^i]}$ is bounded. $\square$

Now since $||A_i|||_{E_g^i}$ is bounded by construction then sufficiently large i and the above lemma shows that $E_g^i \leq \ker(L) \oplus V_g$. We conclude this portion with a corresponding lemma regarding the kernel.

Lemma V.2.4.2. *If $||A_i||_{E_{[\lambda_j^i]}}$ is bounded as $i \to \infty$ there exists an $N \in \mathbb{N}$ for which $E_{[\lambda_j^i]} \cap \ker(L) = \{0\}$ for all $i \geq N$.*

The following proof for this lemma concludes the proof of lemma V.2.4.

$\square$

Proof. Lemma V.2.4.2. Assume $||A_i||_{E_{[\lambda_j^i]}} \to c \in \mathbb{R}$, then $\mathrm{Gr}(A_i) \to L$ implies that the sequence $\mathrm{Gr}(\,A_i|_{E_{[\lambda_j^i]}}) \to K \leq E_{[\lambda_j]} \times E_{[\lambda_j]}$ such that $K \leq L$ and since the norm remains bounded the limit is the graph of a symplectic map; $\mathrm{Gr}(\,A_i|_{E_{[\lambda_j^i]}}) \to \mathrm{Gr}(f) \leq E_{[\lambda_j]} \times E_{[\lambda_j]}$ for some $f \in \mathrm{Sp}(E_{[\lambda_j]})$. Since symplectic maps are non-singular we have $\ker(L) \cap E_{[\lambda_j]} = \{0\}$ so that since $\ker(L) \leq E_s$ and $E_s^i \oplus E_g^i = V$ we see for some $N \in \mathbb{N}$ that $\ker(L) \cap E_{[\lambda_j^i]} = \{0\}$ for all $i \geq N$ as a consequence. $\square$

Lemma V.2.5. *For all $i \geq N$ there exists a unique sequence of symplectic isomorphisms $I_i \colon E_g^i \to V_g$ such that $I_i = Proj|_{E_g^i}$ and $I_i \underset{i\to\infty}{\to} I \colon E_g \to V_g$, where the function $Proj \colon \mathrm{dom}(L) \to \mathrm{dom}(L)/\ker(L) \cong V_g$ is the coisotropic reduced space of $\mathrm{dom}(L)$ uniquely identified by theorem III.2.5 with V_g.*

Proof. Since E_g^i is a $2n-2k$ dimensional subspace of $\mathrm{dom}(L)$ with $\ker(L) \cap E_g^i = \{0\}$ for sufficiently large i then there eventually exists a unique symplectic map $\mathrm{dom}(L)/\ker(L) \cong E_g^i$ (refer to theorem III.2.5) for each i. We denote the above isomorphisms (The co-isotropic reduction of $\mathrm{dom}(L)$ restricted to E_g^i) by $I_i \colon E_g^i \to \mathrm{dom}(L)/\ker(L)$ and see that $V_g \cong \mathrm{dom}(L)/\ker(L) \cong E_g^i$ uniquely via this restriction of the co-isotropic reduction map to E_g^i for large i. The continuity of the coisotropic reduction with respect to a varying subspace of dimension $2n-2k$ (which by lemma V.2.4 eventually has trivial intersection with $\ker(L)$ and $\mathrm{halo}(L)$, the latter implied by $Proj(E_g^i) = \{0\}$) for sufficiently large i, implies that the sequence of isomorphisms converge; $I_i \to I \colon E_g \to V_g$. $\square$

Lemma V.2.6. *Defining* $\phi_i := I_i \circ \beta_i \circ I_i^{-1} \in \mathrm{Sp}(V_g)$ *then* $\phi_i \to \phi \colon V_g \to V_g$ *where* $\mathrm{Gr}(\phi)$ *is the graph part of* L*. Then for sufficiently large* i *the pair* ϕ_i *and* β_i *share the same eigenvalues and each pair of elliptic eigenvalues quadruples have matching Krein type.*

Proof. By using the above identification from theorem III.2.5 between V_g and the reduced domain $\mathrm{dom}(L)/\ker(L)$ we may define $\phi_i := I_i \circ \beta_i \circ I_i^{-1} \in \mathrm{Sp}(V_g)$ so that the ϕ_i and β_i are conjugate. Then since $\mathrm{Gr}(A_i) = \mathrm{Gr}(\alpha_i) \oplus \mathrm{Gr}(\beta_i) \to L$ and $\beta_i \to \beta$ as $i \to \infty$ we see from the continuity of the projection and inclusion maps that

$$\lim_{i\to\infty} I_i \circ \beta_i \circ I_i^{-1} = I \circ \beta \circ I^{-1} = \phi \in \mathrm{Sp}(V)$$

and indeed ϕ and β are conjugate by I as well as ϕ_i and β_i via I_i for sufficiently large

i.

A subtle yet critical note here is that this notion of conjugacy occurs between distinct domains so we verify manually that $\rho^2(\phi_i) = \rho^2(\beta_i)$ for $i \geq N$. First if $\lambda_i \in Spec(\beta_i)$ with eigenvector v^i_λ we let $w^i_\lambda = I_i(v^i_\lambda) \in V_g$ so that $(I_i \circ \beta_i \circ I_i^{-1})(w^i_\lambda) = (I_i \circ \beta_i)(v^i_\lambda) = I_i(\lambda_i v^i_\lambda) = \lambda_i w^i_\lambda$ and the two indeed share the same eigenvalues with I_i mapping eigenvectors of β_i to eigenvectors of ϕ_i.

The remaining concern is regarding the preservation of the conjugacy class of the elliptic eigenvalues since they are precisely the eigenvalues which have any effect on ρ^2. We must verify that the symmetric bilinear form written below maintains the same number of positive eigenvalues under each I_i, that is if $E_\lambda \leq V^{\mathbb{C}}$ is an elliptic eigenspace for β_i we let

$$\begin{aligned} Q_i \colon E_{\lambda_i} \times E_{\lambda_i} &\to \mathbb{R} \\ (z, z') &\mapsto im\left(\omega(z, \overline{z'})\right), \end{aligned}$$

so that the corresponding bilinear form for ϕ_i will be given by $Q_i \circ (I_i^{-1} \times I_i^{-1})$ defined over the eigenspace $F_{\lambda_i} = I_i(E_{\lambda_i})$.

$$\begin{aligned} Q_i \circ (I_i^{-1} \times I_i^{-1}) \colon F_{\lambda_i} \times F_{\lambda_i} &\to \mathbb{R} \\ (z, z') &\mapsto im\left(\omega(I_i^{-1}(z), \overline{I_i^{-1}(z')})\right). \end{aligned}$$

We have implicitly extended I_i^{-1} to a complex symplectic map in the natural way $(I_i^{-1}(iv) := iI_i^{-1}(v))$ so that $I_i^{-1}(\bar{z}) = \overline{I_i^{-1}(z)}$ meaning

$$im\left(\omega(I_i^{-1}(z), \overline{I_i^{-1}(z')})\right) = im\left(\omega(I_i^{-1}(z), I_i^{-1}(\overline{z'}))\right).$$

Then since each I_i is simply the co-isotropic reduction of $\mathrm{dom}(L)$ restricted to E_i we see that $\omega(I_i^{-1}z, I_i^{-1}z') = \omega(z, z')$ for any $z, z' \in F_\lambda^i$ and the two bilinear forms coincide via I_i and therefore share the same number of positive eigenvalues for a given elliptic eigenvalue pair $[\lambda_i]$ given sufficiently large i. It follows that $\rho^2(\beta_i) = \rho^2(\phi_i)$ for any $i \geq N$ and thus $\lim_{i\to\infty} \rho^2(\beta_i) = \rho^2(\phi)$. □

V.3 A Proof that ρ^2 May be Extended Continuously

Theorem V.3.1. *Define the map $\hat{\rho}\colon \mathcal{L}_{2n} \to S^1$ by $\hat{\rho}(L) := \rho^2(\phi)$ for any $L \in \mathcal{L}_{2n}$ possessing the unique decomposition,*

$$L = \ker(L) \times \{0\} \oplus \{0\} \times \mathrm{halo}(L) \oplus \mathrm{Gr}(\phi) \leq (V_s \times V_s) \oplus (V_g \times V_g).$$

Then the map $\hat{\rho}$ is continuous and the diagram below commutes.

$$\begin{tikzcd} \mathcal{L}_{2n} \arrow[dr, "\hat{\rho}"] & \\ \mathrm{Sp}(2n) \arrow[u, hook, "Gr"] \arrow[r, "\rho^2"'] & S^1 \end{tikzcd}$$

Proof. We first refer above to our implicit use of theorem III.2.5 decomposing L since $\mathrm{dom}(L) \cap \mathrm{halo}(L) = \{0\}$. For every $L = \mathrm{Gr}(A) \in \mathcal{L}_{2n}^0 \cong \mathrm{Sp}(2n)$ we know $\phi = A\colon V \circlearrowleft$ so that $\hat{\rho}(L) = \rho^2(A)$ and the above diagram indeed commutes and $\hat{\rho}$ extends ρ^2 via the graph map, it remains to show continuity.

Note that for any $L \in \mathcal{L}_{2n}^n$ that $\rho(L) = \pm 1$ for sufficiently large i (V.1.1) so that $\hat{\rho}(L) = \lim_{i\to\infty} \rho^2(A_i) = 1$ for any sequence A_i such that $\mathrm{Gr}(A_i) \to L \in \mathcal{L}_{2n}^n$ and ρ^2 may be (rather trivially) continuously extended to $\mathcal{L}_{2n}^n$.

For $L \in \mathcal{L}_{2n}^k$ with $1 \le k \le n-1$ from theorem V.2.1 we see there exists some $N \in \mathbb{N}$ for which $A_i = \alpha_i \oplus \beta_i \in \mathrm{Sp}(E_s^i) \times \mathrm{Sp}(E_g^i) \cong \mathrm{Sp}(V_s) \times \mathrm{Sp}(V_g)$ (the first coordinate of this isomorphism is arbitrary but since α_i diverges it is of no concern, the second coordinate isomorphism is unique for large i via theorem V.2.1) such that

$$\mathrm{Gr}(A_i) = \mathrm{Gr}(\alpha_i) \oplus \mathrm{Gr}(\beta_i) \to \ker(L) \times \{0\} \oplus \{0\} \times \mathrm{halo}(L) \oplus \mathrm{Gr}(\beta).$$

For any $(v, w) \in \mathrm{Gr}(\beta)$ we may decompose $v = v_k + v_g \in \ker(L) \oplus V_g$ so that since $(v_k, 0) \in \ker(L) \times \{0\}$ we have that $(v_g, w) \in L$. Then since we know $Proj_{\mathrm{halo}(L)}(E_g^i) = \{0\}$ for $i \ge N$ then $w = w_h + w_g \in \mathrm{halo}(L) \oplus V_g$ with $w_h = 0$ so that $(v_g, w_g) \in \mathrm{Gr}(\beta) \le L$ since $v_g \in V_g$ means we have $w_g = \phi(v_g)$, i.e. the two graphs are seen to coincide after removing $\ker(L)$ components from the domain in $\mathrm{Gr}(\beta)$. This is simply an excessive confirmation that the normal form given in theorem III.2.5 is identical to the limit of $\mathrm{Gr}(\alpha_i) \oplus \mathrm{Gr}(\beta_i)$ after what amounts to some column operations on $\mathrm{Gr}(\beta)$.

Now since $\mathrm{Gr}(\alpha_i) \to \ker(L) \times \{0\} \oplus \{0\} \times \mathrm{halo}(L)$ with $\ker(L) \cap \mathrm{halo}(L) = \{0\}$ then theorem V.1.1 shows that $\rho^2(\alpha_i) = 1$ for $i \ge N$ and thus $\rho^2(A_i) = \rho^2(\alpha_i)\rho^2(\beta_i) = \rho^2(\beta_i)$ for all $i \ge N$. Then since ρ is continuous on $\mathrm{Sp}(V_g)$ we see that the right-most equality below holds;

$$\lim_{i\to\infty} \rho^2(A_i) = \lim_{i\to\infty} \rho^2(\beta_i) = \lim_{i\to\infty} \rho^2(\phi_i) = \rho^2(\phi)$$

and $\hat{\rho}(L) := \rho^2(\phi)$ is indeed continuous. $\square$

Theorem I.4.2. There exists a unique, real valued continuous function $\hat{\Delta}$ constant on fixed endpoint homotopy classes of paths in $\mathcal{L}_{2n}$ such that for any path $\gamma\colon I \to \mathrm{Sp}(2n)$ we have that $\hat{\Delta}(\mathrm{Gr}(\gamma)) = 2\Delta(\gamma)$.

Proof. We begin by restating the construction of the extended mean index for arbitrary paths $\gamma\colon I \to \mathcal{L}_{2n}$ using the unique continuous map $\hat{\theta}\colon I \to \mathbb{R}$ satisfying $\hat{\theta}(0) \in [-\pi, \pi)$ and $(\hat{\rho} \circ \gamma)(t) = e^{i\hat{\theta}(t)}$ for all $t \in I$. Then if we let $\hat{\Delta}(\gamma) = \frac{\hat{\theta}(1)-\hat{\theta}(0)}{2\pi}$ it is both continuous and well defined modulo fixed endpoint homotopy classes by construction and since $\hat{\rho} = \rho^2$ on $\Lambda^0_{2n} \cong \mathrm{Sp}(2n)$ its clear that any $\gamma \subset \mathrm{Sp}(2n)$ has $\hat{\theta}(t) = 2\theta(t)$ and thus,

$$\hat{\Delta}(\mathrm{Gr}(\gamma)) = \frac{2\theta(1) - 2\theta(0)}{2\pi} = 2\Delta(\gamma).$$

$\square$

Part VI

Properties of $\hat{\Delta}$ over Stratum-Regular Paths

VI.1 Continuity of Compatible Path Compositions

We proceed in continuing the work we began in section II.4, first by defining compatibility among the stratum-regular paths (shown in proposition II.4.8 to be an open and dense subset of $C^1([0,1], \mathcal{L}_{2n})$). Recalling definition II.1.2 we note that $\mathrm{Pr}_I : \mathcal{L}_{2n} \to \bigsqcup_{k=0}^{n} \mathcal{I}_k$ is the map sending each $L \in \mathcal{L}_{2n}^k$ to the associated isotropic pair $(\ker(L), \mathrm{halo}(L)) \in \mathcal{I}_k$ when $1 \leq k = \dim(\ker(L)) \leq n$. When $k = 0$ we recall that Pr_I is identically zero on the open and dense subset $\mathrm{Im}(\mathrm{Gr}) = \mathcal{L}_{2n}^0 \subset \mathcal{L}_{2n}$, where zero in this case refers to the trivial isotropic pair $\mathcal{I}_0 := \{0\}$ as in remark II.1.5. In other words, since $L \in \Lambda_{2n}^0$ if and only if there exists $A \in \mathrm{Sp}(V)$ for which $L = \mathrm{Gr}(A)$, then for all $A \in \mathrm{Sp}(V)$, we have $\mathrm{Pr}_I(\mathrm{Gr}(A)) = \{0\} =: \mathcal{I}_0$.

Definition VI.1.1. *Given any $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$ we define*

$$[\gamma] := \{\tau \in \mathcal{P}_{reg}(\mathcal{L}_{2n}) \mid \mathrm{Pr}_I(\tau(t)) = \mathrm{Pr}_I(\gamma(t)), \forall t \in [0,1]\}$$

and call this the set of γ-compatible paths.

Remark VI.1.2. *One may observe that this definition is trivially satisfied for all but finitely many $t \in [0,1]$ when restricted to the stratum-regular γ, as such paths reside in Λ_{2n}^0 for all but finitely many $t \in [0,1]$. In this manner we might 'trim down' our definition to a finitary one by giving an equivalent definition in terms of the finite partitions induced by each $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$ as given in definition VI.1.4 below. The above definition is a vestigial feature better suited to the larger class of paths as described in remark I.4.8.*

We now show that γ-compatibility is an equivalence relation on the set of stratum-regular paths.

Proposition VI.1.3. *Given any stratum-regular $\gamma, \tau \in [\gamma]$ we have that $\gamma \in [\gamma]$ and $[\gamma] = [\tau]$. Additionally, since every $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$ is contained within $[\gamma]$, we see that the collection of subsets,*

$$\{[\gamma] \,|\, \gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})\}$$

is a disjoint covering of $\mathcal{P}_{reg}(\mathcal{L}_{2n})$.

It follows that γ-regularity is an equivalence relation on $\mathcal{P}_{reg}(\mathcal{L}_{2n})$, see proof VI.3 for more details.

We note that proposition VI.1.3 implies that any statement regarding a pair of elements $\tau, \xi \in [\gamma]$ (e.g. the quasimorphism bound) is independent of whatever γ is chosen to represent the compatibility class. This implies we may safely set $\xi = \gamma$ in such cases without loss of generality.

Definition VI.1.4. *Each compatibility class $[\tau]$ induces the following two objects.*

1. *A finite partition $(t_i)_{i=1}^{M}$ of the unit interval:*

   $$0 \leq t_1 < t_2 < \cdots < t_M \leq 1,$$

   where each time t_i is such that for every representative $\gamma \in [\tau]$, at least one of the index images $\gamma_{\mathcal{K}}([t_i, t_i + \epsilon))$, $\gamma_{\mathcal{K}}((t_i - \epsilon, t_i])$ is non-constant for all $\epsilon > 0$.

2. The collection $\{U_i\}_{i=0}^M$ *of open subsets of* $[0,1]$ *given by,*

$$U_i = \begin{cases} [0,t_1) & i=0 \\ (t_i,t_{i+1}) & 1 \le i \le M-1 \\ (t_M,1] & i=M \text{ if } t_M < 1 \\ (t_{M-1},1) & i=M \text{ if } t_M = 1 \end{cases}$$

are defined such that any $\gamma \in [\tau]$ *is constant on each.*

The pair $\{U_i\}_{i=0}^M, \{t_i\}_{i=1}^M$ *of subset collections together form a disjoint cover of* $[0,1]$*,*

$$\bigcup_{i=0}^M U_i \cup \bigcup_{i=1}^M \{t_i\} = [0,1].$$

over which $\mathrm{Im}(\gamma)$ *is partitioned into* M *open paths in* $\Lambda_{2n}^0 \cong \mathrm{Sp}(V)$ *and* M *singletons* $\{t_i\}$ *when* $\gamma(t_i) \in \mathcal{L}_{2n}^1$*.*

In other words this partition marks each time t_i at which the index of γ has a discontinuity (each of which are jump discontinuities from the 0^{th} to the 1^{st} stratum), which happens finitely often due to our regularity condition (see proposition II.4.16). More precisely, letting $\gamma_{\mathcal{K}} = \dim(\ker(\gamma(t)))$ we have,

$$\gamma_{\mathcal{K}}^{-1}(1) = \{t_i\}_{i=1}^M \text{ and } \gamma_{\mathcal{K}}(t) = 0 \text{ for } t \in [0,1] \setminus \{t_i\}_{i=1}^M = \bigcup_{i=0}^M U_i.$$

The following lemma VI.1.5 is a technical step necessary in both stating and proving corollary I.4.6 and theorems I.4.5, I.4.7. This lemma's proof has been relegated to section VI.3 in addition to the proof for proposition VI.1.3 (see proof VI.3) as they each constitute necessary, albeit technical steps.

Lemma VI.1.5. *Consider any compatibility class* $[\gamma] \subset \mathcal{P}_{reg}(\mathcal{L}_{2n})$ *and representative* $\gamma \in [\gamma]$ *and recall the associated partition* $\{t_i\}_{i=1}^M$ *of the unit interval. If we let* $\phi_i := \gamma|_{U_i} : U_i \to \mathrm{Sp}(V)$ *then there exists a sequence of bounded paths* $\tilde{\phi}_i : [t_i, t_{i+1}] \to \mathrm{Sp}(V)$ *(so that* $\Delta(\tilde{\phi}_i)$ *is well defined) for which* $\Delta(\tilde{\phi}_i) = \hat{\Delta}(\phi_i)$[21] *for each* $0 \le i \le M$*. Consequently* $\hat{\Delta}$ *may be decomposed via concatenation with respect to the partition,*

$$\hat{\Delta}(\gamma) = \sum_{i=1}^{M} \hat{\Delta}(\phi_i) = \sum_{i=1}^{M} \Delta(\tilde{\phi}_i).$$

Theorem I.4.5. Given any $\gamma, \tau \in [\gamma] \subset \mathcal{P}_{reg}(\mathcal{L}_{2n})$ both $\gamma \circ \tau$ and $\tau \circ \gamma$ are well defined piece-wise differentiable paths in $\mathcal{L}_{2n}$. In particular both $\gamma \circ \tau$ and $\tau \circ \gamma$ are stratum-regular and γ-compatible for all $\tau \in [\gamma]$ as is $\gamma^i \in [\gamma]$ for any $i \ge 1$.

Proof. We proceed for any stratum-regular γ and any γ-compatible τ invoking definition VI.1.4 on $[\gamma]$ to obtain the partition $\{t_i\}_{i=1}^M$ and decomposition $\{U_i\}_{i=0}^M$. Then from the definition of γ-regularity we first note that $\mathrm{Pr}_I(\gamma(t)) = \mathrm{Pr}_I(\tau(t))$ for all $t \in [0,1]$ which is trivially satisfied on the U_i where $\mathrm{Pr}_I(\gamma(t)) = \mathrm{Pr}_I(\tau(t)) = \{0\}$. More importantly we have $\mathrm{Pr}_I(\gamma(t_i)) = \mathrm{Pr}_I(\tau(t_i)) = (\gamma_{\mathrm{ker}}(t_i), \gamma_{\mathrm{halo}}(t_i)) \in \mathcal{I}_1 \cong \mathbb{RP}^{2n-1} \times \mathbb{RP}^{2n-1} \setminus \hat{H}_1$. We will proceed referring to these isotropic pairs in terms of the representative γ.

Then recalling theorem III.2.5 we see that any $L, L' \in \mathcal{L}_{2n}^1$ for which $\mathrm{Pr}_I(L) =$

[21] This equality omits the scaling factor of 2 that is an artifact of squaring ρ. We have decided to omit it as it's entirely inconsequential for the proof being uniform over all ϕ_i and over all intervals.

$\mathrm{Pr}_I(L')$ yields a pair of unique decompositions,

$$\begin{aligned} L &= B_1 \times \{0\} \oplus \{0\} \times B_2 \oplus \mathrm{Gr}(\phi) \\ L' &= B_1 \times \{0\} \oplus \{0\} \times B_2 \oplus \mathrm{Gr}(\phi') \end{aligned}$$

where $\phi, \phi' \in \mathrm{Sp}(B_1^\omega \cap B_2^\omega) \cong \mathrm{Sp}(2n-2)$. Then, with respect to this decomposition we have

$$\begin{aligned} L \circ L' &= B_1 \times \{0\} \oplus \{0\} \times B_2 \oplus \mathrm{Gr}(\phi\phi') \\ L' \circ L &= B_1 \times \{0\} \oplus \{0\} \times B_2 \oplus \mathrm{Gr}(\phi'\phi) \end{aligned}$$

where the products $\phi\phi', \phi'\phi$ are the usual group operation on

$$Sp\left(B_1^\omega \cap B_2^\omega, \tilde{\omega}|_{B_1^\omega \cap B_2^\omega}\right).$$

Now in terms of the 0^{th} stratum intersections, lemma VI.1.5 produces a collection of symplectic paths $\phi_i, \eta_i : U_i \to \mathrm{Sp}(V)$ for each $0 \leq i \leq M$ such that each $\phi_i := \gamma|_{U_i}$ and $\eta_i := \tau|_{U_i}$. It follows from the group isomorphism $\mathcal{L}_{2n}^0 \cong \mathrm{Sp}(V, \omega)$ that the differentiability of τ and γ implies that their composites $\tau \circ \gamma$ and $\gamma \circ \tau$ are each differentiable on U_i for all $0 \leq i \leq M$ (since Lagrangian composition and group multiplication coincide on the 0^{th} stratum).

Without loss of generality we will continue only proving the result for $\tau \circ \gamma$ as the corresponding proof for $\gamma \circ \tau$ is identical. All that remains to be shown is that the composite path $\tau \circ \gamma$ is continuous at the t_i. Considering any t_{i-1} for $i \leq M$ at which $\gamma_{\mathcal{K}}$ is discontinuous we may consider some sequence $\{a_l\}_{l=1}^\infty$ (without loss of generality) for which $t_{i-1} < a_l$ for all $l \in \mathbb{N}$ and $a_l \to t_i^-$ as $l \to \infty$. The one sided limit is sufficient since the two-sided limit coincides with $(\tau \circ \gamma)|_{U_i}$ as $t \to t_i^+$ which follows after reversing the time of the path in a neighborhood about t_i.

Recall that $\gamma(t_i)$ and $\tau(t_i)$ share the symplectic splitting $V = V_s(t_i) \overset{\omega}{\oplus} V_g$ where $V_g(t_i) := \mathrm{dom}(\gamma(t_i)) \cap \mathrm{ran}(\gamma(t_i))$ is the $2n-2$ dimensional symplectic vector space over which the graph parts $\phi(t_i), \eta(t_i)$ of $\gamma(t_i), \tau(t_i)$ are defined, while $V_s = V_g^\omega$ is the 2 dimensional symplectic subspace on which $\phi_l := \phi_i(a_l), \eta_l := \eta_i(a_l)$ diverge as $l \to \infty$.

Then since

$$\lim_{l\to\infty} \mathrm{Gr}(\eta_l|_{V_g}) = \tau(t_i)|_{V_g \times V_g} \text{ and } \lim_{l\to\infty} \mathrm{Gr}(\phi_l|_{V_g}) = \gamma(t_i)|_{V_g \times V_g},$$

and Lagrangian composition coincides with the symplectic group operation via the graph map on $\mathcal{L}_{2n}^0$ we see that,

$$\lim_{l\to\infty} \mathrm{Gr}(\eta_l \phi_l)|_{V_g \times V_g} = \lim_{l\to\infty} (\tau(a_l) \circ \gamma(a_l))|_{V_g \times V_g} = \tau(t_i) \circ \gamma(t_i)|_{V_g \times V_g}.$$

So that since $(\gamma \circ \tau)(a_l) = \phi_i(a_l)\eta_i(a_l)$ for all $l \in \mathbb{N}$ the product does indeed converge to the correct Lagrangian subspace of $V_g \times V_g$ (the graph portion) and all that remains to be shown is that the vanishing (diverging) subspaces of the product $\eta(a_l)\phi(a_l)$ converge in the limit $l \to \infty$ to $\gamma_{\mathrm{ker}}(t_i)$ and $\gamma_{\mathrm{halo}}(t_i)$ respectively. In particular we desire,

$$\lim_{l\to\infty} \mathrm{Gr}(\eta_l \phi_l))|_{V_s \times V_s} = \tau(t_i) \circ \gamma(t_i)|_{V_s \times V_s}.$$

That is to say, since $V_s \oplus V_g = V$ and,

$$\tau(t_i) \circ \gamma(t_i) = \tau(t_i) \circ \gamma(t_i)|_{V_s \times V_s} \oplus \tau(t_i) \circ \gamma(t_i)|_{V_g \times V_g}, \tag{VI.1.1}$$

we see that $\mathrm{Gr}(\eta_i \phi_i)$ converges to $\tau(t_i) \circ \gamma(t_i)$ on U_i provided proposition VI.1.7 holds.

To proceed in proving proposition VI.1.7 we must be certain that the limit of our symplectic product graphs $\mathrm{Gr}(\eta_l \phi_l)$ does not end up in H as $l \to \infty$.

Proposition VI.1.6. *The limit $K := \lim_{l \to \infty} \mathrm{Gr}(\eta_l \phi_l) \notin H$.*

Proof. We suppose $K \in H$, then there exists a sequence of normalized vectors $v_l \to v \neq 0$ in V such that $v \in \ker(K) \cap \mathrm{ran}(K)$. It follows from the kernel membership that $|(\eta_l \phi_l)(v_l)| \to 0$ as $l \to \infty$ and from the range that $|\phi_l^{-1} \eta_l^{-1}(v_l)| < \infty$. If we let $w_l = \frac{\phi_l(v_l)}{|\phi_l(v_l)|}$ we see from membership in the kernel that one of two possibilities exist:(a) $|\eta_l(w_l)| \to 0$ or (b) $|\phi_l(v_l)| \to 0$ and $|\eta_l(w_l)| < \infty$. In both cases $|\eta_l \phi_l(v_l)| < \infty$ which ensures that $\lim_{l \to \infty} v_l = v \notin \mathrm{halo}(K)$ (Recall that the v_l are normalized).

It follows that $[v]_{\mathrm{ran}} \neq 0$ where $[v]_{\mathrm{ran}} \in \mathrm{ran}(K)/\mathrm{halo}(K)$ yet $[v]_{\mathrm{dom}} = 0$ in $\mathrm{dom}(K)/\ker(K)$ which yields our contradiction. Namely the map

$$K_{\mathrm{Gr}} : \mathrm{dom}(K)/\ker(K) \to \mathrm{ran}(K)/\mathrm{halo}(K),$$

is guaranteed to be symplectic and in particular, a well defined function so that $K_{\mathrm{Gr}}([0]_{\mathrm{dom}}) \neq [0]_{\mathrm{ran}}$ provides a (rather significant) contradiction (see proposition III.2.1 for more details on the aforementioned map K_{Gr}). □

To finish the proof we prove the aforementioned proposition.

Proposition VI.1.7.

$$\lim_{l\to\infty} \mathrm{Gr}(\eta_l\phi_l))|_{V_s\times V_s} = \tau(t_i)\circ\gamma(t_i)|_{V_s\times V_s}. \tag{VI.1.2}$$

Proof. Since $K := \lim_{l\to\infty}\mathrm{Gr}(\eta_l\gamma_l) \notin H$ we may invoke theorem V.2.1 to obtain a sequence of symplectic splittings $E^l_S \overset{\omega}{\oplus} E^l_g = V$ for which E^l_s converges to some $\tilde{V}_s$. Additionally equation (VI.1.1) shows that $E^l_g \to V_g$ as $l\to\infty$. Then since $V_g = V_s^\omega$ by theorem III.2.5 we see that it must be that $\tilde{V}_s = V_s$ so that $\tilde{V}_s = \ker(K)\oplus\mathrm{halo}(K) = V_s$ and we need only show that the elements of v_s vanish (diverge) appropriately. That is to say $\ker(K) = \ker((\tau\circ\gamma)(t_i))$ and $\mathrm{halo}(K) = \mathrm{halo}((\tau\circ\gamma)(t_i))$.

Since these are one dimensional subspaces we may do this via computation. Beginning by letting $\langle v\rangle = \ker(\gamma(t_i))$ we see that $|(\eta_l\phi_l)(v)| \not\to 0$ if and only if $v\in\mathrm{halo}(\tau(t_i))$ since $|\phi_l(v)|\to 0$. Then since $\ker(\gamma(t_i))\cap\mathrm{halo}(\gamma(t_i)) = \{0\}$ and $\mathrm{halo}(\gamma(t_i)) = \mathrm{halo}(\tau(t_i))$ this can't be possible and the only alternative is that $|(\eta_l\phi_l)(v)|\to 0$ and $\ker((\tau\circ\gamma)(t_i)) = \ker(K)$.

Similarly for $\langle w\rangle = \mathrm{halo}(\gamma(t_i))$ we see that $|(\eta_l\phi_l)^{-1}(w)| \not\to 0$ if and only if $w\in\ker(\tau(t_i))$ which as above, can't be true as it implies non-transversality of the kernel and halo so we obtain the corresponding result $\mathrm{halo}((\tau\circ\gamma)(t_i)) = \mathrm{halo}(K)$ as a consequence. $\square$

As detailed above the theorem follows from equation (VI.1.1) since proposition VI.1.7 is true. That is, $(\gamma\circ\tau)(t_i) = K = \ker(\gamma)\times\{0\}\oplus\{0\}\times\mathrm{halo}(\gamma)\oplus\mathrm{Gr}(\phi_i\eta_i)$ and the path $\gamma\circ\tau$ converges to the appropriate $K\in\mathcal{L}^1_{2n}$ and so is continuous at each of the

t_i, and therefore is piece-wise differentiable on all of $[0, 1]$. $\square$

VI.2 The Quasimorphism Bound

We begin by stating an immediate corollary of the previous section and lemma VI.1.5, namely that of homogeneity.

Corollary I.4.6. The extended mean index $\hat{\Delta}$ is homogeneous over any stratum-regular γ,

$$\hat{\Delta}(\gamma^l) = l \cdot \hat{\Delta}(\gamma).$$

Recall that theorem I.4.5 implies γ^l is the piece-wise differentiable path given by point-wise composition as written in definition I.3.2.

Proof. This proof will be pleasantly brief compared to the previous one. We begin with lemma VI.1.5 to see for any $l \in \mathbb{N}$ and $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$ that

$$\hat{\Delta}(\gamma^l) = \sum_{i=0}^{M} \hat{\Delta}(\phi_i^l) = \sum_{i=0}^{M} \Delta(\tilde{\phi}_i^l) = \sum_{i=0}^{M} l \cdot \Delta(\tilde{\phi}_i) = l \cdot \hat{\Delta}(\gamma).$$

$\square$

Now we recover the quasimorphism bound on the compatibility classes.

Theorem I.4.7. For any $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$ and $\tau \in [\gamma]$ the Lagrangian mean index $\hat{\Delta}$ satisfies

$$|\hat{\Delta}(\gamma \circ \tau) - \hat{\Delta}(\tau) - \hat{\Delta}(\gamma)| < C$$

where C depends only on $[\gamma]$, i.e. the above bound is uniform on each compatibility class[22].

Proof. Consider two compatible paths $a, b : [0,1] \to \mathcal{L}_{2n}$ which by definition have identical unit interval partitions for all $i \leq M$ and for which we write the associated graph portions $\alpha_i, \beta_i : U_i \to \mathrm{Sp}(V)$ and bounded approximations $\tilde{\alpha}_i, \tilde{\beta}_i : K_i \to \mathrm{Sp}(V)$, defined over the compact subsets $K_i = cl(U_i)$ described in proof VI.3 below. All that is necessary to know here is that the $\alpha_i : U_i \to \mathrm{Sp}(V)$, being unbounded lie outside the realm of the usual applications of the mean index so in we define the 'bounded approximation' $\tilde{\alpha}_i : [t_i, t_{i+1}] \to \mathrm{Sp}(V)$ for which there exists some ϵ_i for which $\tilde{\alpha}_i(t) = \alpha_i(t)$ for $t \in [t_i + \epsilon_i, t_{i+1} - \epsilon_i]$. Yet, near the boundaries where $\alpha_i(t)$ has a divergent component we essentially isolate and pause it in $\tilde{\alpha}_i$ at $t_i + \epsilon_i$ and $t_i - \epsilon_i$, after which the graph parts of $\tilde{\alpha}_i$ and α_i coincide such that $\Delta(\alpha_i) = \Delta(\tilde{\alpha}_i)$ (guaranteed for sufficiently small ϵ_i by a two sided application of theorem V.2.1), see proof VI.3 for more details.

[22] Whether there exists a bound which is uniform over all compatible pairs is unlikely to be shown via this method, as arbitrary numbers of intersections with higher strata are possible so that any attempt and taking the supremum over all compatibility classes without somehow normalizing with respect to this intersection count would likely lead to $C \to \infty$ (as C is a linear function of $M = |\,\mathrm{Im}(\gamma) \cap \mathcal{L}^1_{2n}|$.)

Now recalling lemma VI.1.5,

$$\hat{\Delta}(a) = \sum_{i=1}^{M} \Delta(\tilde{\alpha}_i)$$
$$\hat{\Delta}(b) = \sum_{i=1}^{M} \Delta(\tilde{\beta}_i)$$
$$\hat{\Delta}(a \circ b) = \sum_{i=1}^{M} \Delta(\tilde{\alpha}_i \tilde{\beta}_i)$$

where the product $\tilde{\alpha}_i \tilde{\beta}_i$ is the usual group operation on $\mathrm{Sp}(V)$ identified with $\mathcal{L}^0_{2n}$ via the graph map.

We now reconcile the fact that these paths are highly unlikely to originate at the same point, let alone the origin. We construct a trio of paths $p_{\tilde{\alpha}_i}, p_{\tilde{\beta}_i}, p_{\tilde{\alpha}_i\tilde{\beta}_i} : I \to \mathrm{Sp}(V)$ for each $i \leq M$ which originate at the identity and which terminate at the initial points of $\tilde{\alpha}_i$, $\tilde{\beta}_i$ and $\tilde{\alpha}_i\tilde{\beta}_i$ respectively.

We recall lemma I.3.9 as given in the introduction wherein the universal cover of $\mathrm{Sp}(2n)$ may be written as

$$\widetilde{Sp}(2n) = \left\{ (g, c) \in G \times \mathbb{R} \,|\, \rho(g) = e^{ic} \right\}.$$

with group action given by,

$$(g_1, c_1) \cdot (g_2, c_2) = (g_1 g_2, c_1 + c_2).$$

We consider the homotopy classes of paths originating at the identity and terminating at $\tilde{\alpha}_i(0)$ denoted $\widetilde{\mathcal{P}}_{Id,\tilde{\alpha}_i(0)}(\mathrm{Sp}(2n))$, on which we wish to choose $p_{\tilde{\alpha}_i}$ such that $|\Delta(p_{\tilde{\alpha}_i})|$ is minimized. If we suppose $\Delta([p_{\tilde{\alpha}_i}])$ is not minimal in $\widetilde{\mathcal{P}}_{Id,\tilde{\alpha}_i(0)}(\mathrm{Sp}(2n))$ then the group action of $\pi_1(\mathrm{Sp}(2n))$ on the universal cover along with lemma I.3.9 implies there exists a deck transformation (identified with concatenation) induced by some nonzero $[l] \in \pi_1(\mathrm{Sp}(2n))$ and some $[p'_{\tilde{\alpha}_i}] \in \widetilde{\mathcal{P}}_{Id,\tilde{\alpha}_i(0)}(\mathrm{Sp}(2n))$ for which $[p_{\tilde{\alpha}_i}] = [p'_{\tilde{\alpha}_i}] * [l]$. It follows that $\Delta([p_{\tilde{\alpha}_i}]) = \Delta([p'_{\tilde{\alpha}_i}] * [l])$ so that $|\Delta([p_{\tilde{\alpha}_i}])| = |\Delta([p'_{\tilde{\alpha}_i}]) + \Delta([l])| > |\Delta([p'_{\tilde{\alpha}_i}])|$ since the mean index of a non-trivial loop is non-zero. Repeat this procedure on any $[p_{\tilde{\alpha}_i}]$ enough times and we will reach the point where the homotopy class $[p_{\tilde{\alpha}_i}]$ has minimal (absolute value) mean index. We repeat the same procedure for both $[p_{\tilde{\beta}_i}]$ and $[p_{\tilde{\alpha}_i\tilde{\beta}_i}]$ and assume they too have minimal (absolute value) mean index.

Explicitly the definition of ρ implies that the worst case scenario would be a path with all S^1 eigenvalues each with the maximal number of type I Krein type eigenvalues (equivalently, the quadratic form $\operatorname{Im}\omega(z, \overline{z'})$) over the generalized eigenspace of said eigenvalue has $2n$ positive real eigenvalues). In either case for any path $\tilde{\alpha}_i$ that has undergone the above minimization procedure we see $\rho(p_{\tilde{\alpha}_i}) = (\theta t)^n$ so that $|\Delta(p_{\tilde{\alpha}_i}(t))| = n\theta$. It follows that the maximal rotation of $\theta = 2\pi$ yields $|\Delta(p_{\tilde{\alpha}_i})| \leq n$. Indeed, by the previous argument if $|\Delta(p_{\tilde{\alpha}_i})| > n$ this would imply the existence of some identity based loop q and path from Id to $\tilde{a}_i(0)$ for which $p_{\tilde{\alpha}_i} = p'_{\tilde{\alpha}_i} * q$ which violates our minimality assumption. The same holds for $\tilde{\beta}_i$ and $\tilde{\alpha}_i \circ \tilde{\beta}_i$ as well,

$$|\Delta(p_{\tilde{\alpha}_i})|, |\Delta(p_{\tilde{\beta}_i})|, |\Delta(p_{\tilde{\alpha}_i\tilde{\beta}_i})| \leq n.$$

It follows that we have three concatenated paths, $l_{\tilde{\alpha}_i} := \tilde{\alpha}_i * p_{\tilde{\alpha}_i}, l_{\tilde{\beta}_i} := \tilde{\beta}_i * p_{\tilde{\beta}_i}$ and $l_{\tilde{\alpha}_i\tilde{\beta}_i} := (\tilde{\alpha}_i\tilde{\beta}_i) * p_{\tilde{\alpha}_i\tilde{\beta}_i}$ with concatenated domain J_i being the union of the $K_i = Cl(U_i)$ with the domain of the identity-based connecting paths. Then since Δ is additive under concatenation we see

$$\begin{aligned}
|\Delta(\tilde{\alpha}_i) - \Delta(l_{\tilde{\alpha}_i})| &\leq n \\
|\Delta(\tilde{\beta}_i) - \Delta(l_{\tilde{\beta}_i})| &\leq n \\
|\Delta(\tilde{\alpha}_i\tilde{\beta}_i) - \Delta(l_{\tilde{\alpha}_i\tilde{\beta}_i})| &\leq n.
\end{aligned}$$

Then since lemma VI.1.5 shows that $\hat{\Delta}(\alpha_i) = \Delta(\tilde{\alpha}_i), \hat{\Delta}(\beta_i) = \Delta(\tilde{\beta}_i)$ and $\hat{\Delta}(\alpha_i \circ \beta_i) = \Delta(\tilde{\alpha}_i\tilde{\beta}_i)$ it follows that the above estimates allow us to write,

$$\begin{aligned}
&\left|\hat{\Delta}(a \circ b) - \hat{\Delta}(a) - \hat{\Delta}(b)\right| \\
\leq\ & \sum_{i=1}^{M} \left|\hat{\Delta}(\alpha_i \circ \beta_i) - \hat{\Delta}(\alpha_i) - \hat{\Delta}(\beta_i)\right| \\
=\ & \sum_{i=1}^{M} \left|\Delta(\tilde{\alpha}_i\tilde{\beta}_i) - \Delta(\tilde{\alpha}_i) - \Delta(\tilde{\beta}_i)\right| \\
\leq\ & \sum_{i=1}^{M} 3n + |\Delta(l_{\tilde{\alpha}_i \circ \tilde{\beta}_i}) - \Delta(l_{\tilde{\alpha}_i}) - \Delta(l_{\tilde{\beta}_i})| \\
=\ & 3Mn + \sum_{i=1}^{M} |\Delta(l_{\tilde{\alpha}_i \circ \tilde{\beta}_i}) - \Delta(l_{\tilde{\alpha}_i}) - \Delta(l_{\tilde{\beta}_i})| \\
<\ & M(3n + c) < \infty.
\end{aligned}$$

Thus for every compatibility class $[\gamma]$, and pair of compatible paths $\gamma, \tau \in [\gamma]$ the quasimorphism property holds for $C = M(3n + c)$. □

Remark VI.2.1. *The size of the bound is almost certainly a gross overestimate and an artifact of the partitioning procedure into some arbitrary number of symplectic paths. Regardless the uniform bound c we have over each* $\mathrm{Sp}(V)$ *indeed translates over in this special case.*

The sensitivity to parameterization is now becoming so prominent it would be fair to characterize it as 'distressing'; every $\tau \in [\gamma]$ is likely to immediately leave the set of γ-compatible paths upon any perturbation of the time scale. Similarly any perturbation of τ as an element in $\mathcal{P}_{reg}(\mathcal{L}_{2n})$ will have the same result, a loss of compatibility with γ. It follows that arbitrary perturbations of τ in the space of differentiable paths $C^1([0,1], \mathcal{L}_{2n})$ do not preserve compatibility class.

If we restrict ourselves to those perturbations τ^ϵ which preserve each base path, $\mathrm{Pr}_I(\tau(t)) = \mathrm{Pr}_I(\tau^\epsilon(t))$ (or in our case the sequence $\{\mathrm{Pr}_I(\tau((t_i))\}_{i=1}^M$) then τ^ϵ will remain both stratum-regular and γ compatible, which may be the best sort of stability we could hope for.

VI.3 Proof of proposition VI.1.3 and lemma VI.1.5

We begin with the promised proof of proposition VI.1.3.

Proof: Proposition VI.1.3. We consider the two paths as before and observe that any stratum-regular ξ is γ-compatible if and only if it is τ-compatible;

$$\xi \in [\gamma] \Leftrightarrow \xi \in [\tau],$$

so the two sets must be identical.

In regards to the covering of the stratum-regular paths we note that every stratum-regular path is contained within its own compatibility class implying that the above collection indeed covers $\mathcal{P}_{reg}(\mathcal{L}_{2n})$ (so that by proposition II.4.13 we see this is an equivalence relation on an open and dense subset of $C^1([0,1],\mathcal{L}_{2n})$). Now consider any pair γ, τ of stratum-regular paths for which $[\gamma]\neq[\tau]$ and suppose there exists $\xi\in[\gamma]\cap[\tau]\neq\emptyset$. Since γ and τ are not compatible (yet ξ is compatible with both) there exists some t_0 for which $\mathrm{Pr}_I(\xi(t_0))=\mathrm{Pr}_I(\tau(t_0))\neq\mathrm{Pr}_I(\gamma(t_0))=\mathrm{Pr}_I(\xi(t_0))$. No such ξ exists and we have,

$$[\gamma]\neq[\tau]\Rightarrow[\gamma]\cap[\tau]=\emptyset.$$

Now letting γ,τ,ξ be stratum-regular we see that the binary relation of γ-compatibility is indeed an equivalence relation,

– Reflexive: $\gamma\in[\gamma]$.

– Symmetric: $\gamma\in[\tau]\Leftrightarrow\tau\in[\gamma]$.

– Transitive: $\gamma\in[\tau]$, $\tau\in[\xi]\Rightarrow\gamma\in[\xi]$.

$\square$

Denoting $\gamma\in[\tau]$ as $\gamma\sim_{comp}\tau$, we see that the set of equivalence classes is well defined and isomorphic to a certain collection of finite sequences with strictly increasing real part,

$$\mathcal{P}_{reg}(\mathcal{L}_{2n})/\sim_{comp}\cong\bigcup_{M=1}^{\infty}\left\{(t_i,L_i)_{i=1}^{M}\in[0,1]\times\mathcal{I}_1\,|\,t_i<t_j\Leftrightarrow i<j\right\}.$$

If we had used the definition as discussed in the remark I.4.8 above (that $|\pi_0(\gamma^{-1}(\mathcal{L}_{2n}^k)|<\infty$ for all $0\leq k\leq n$) this set would be far more complicated due to the non-transverse intersection.

Of minor interest is a diffeomorphism which makes the situation a bit more familiar, $\mathcal{I}_1 \cong \mathbb{RP}^{2n-1} \times \mathbb{RP}^{2n-1} \setminus T$ where

$$T = \left\{ [v, w] \in \mathbb{RP}^{2n-1} \times \mathbb{RP}^{2n-1} \,|\, v \in \langle w \rangle^{\omega} \Leftrightarrow \omega(v, w) = 0 \right\},$$

or the space of one dimensional subspace pairs of V whose direct sum is isotropic, in turn identifying the first stratum as the space of one dimensional subspace pairs $(B_1, B_2) \in \mathbb{RP}^{2n-1} \times \mathbb{RP}^{2n-1}$ for which $B_1 \oplus B_2$ is symplectic.

Now we finally come to proving our key technical lemma VI.1.5. It is essentially a continuous analogue of theorem V.2.1, and while the hypothesis may be easily generalized to $2 \leq k \leq n$ we only prove the lowest case $k = 1$ below in an attempt at efficiency.

Lemma VI.1.5. Consider any compatibility class $[\gamma] \subset \mathcal{P}_{reg}(\mathcal{L}_{2n})$ and representative $\gamma \in [\gamma]$ and recall the associated partition $\{t_i\}_{i=1}^{M}$ of the unit interval. If we let $\phi_i := \gamma|_{U_i} : U_i \to \mathrm{Sp}(V)$ then there exists a sequence of bounded paths $\tilde{\phi}_i : [t_i, t_{i+1}] \to \mathrm{Sp}(V)$ (so that $\Delta(\tilde{\phi}_i)$ is well defined) for which $\Delta(\tilde{\phi}_i) = \hat{\Delta}(\phi_i)$ for each $0 \leq i \leq M$.

Consequently $\hat{\Delta}$ may be decomposed via concatenation with respect to the partition,

$$\hat{\Delta}(\gamma) = \sum_{i=1}^{M} \hat{\Delta}(\phi_i) = \sum_{i=1}^{M} \Delta(\tilde{\phi}_i).$$

Proof. We denote the graph portions of the (transverse) intersections of $\mathrm{Im}(\gamma)$ with $\mathcal{L}^1_{2n}$ by $\Phi_i = \mathrm{Pr}_{Sp}(\gamma(t_i)) \in \mathrm{Sp}(V_g(t_i)) \cong \mathrm{Sp}(2n-2)$ for each $0 \leq i \leq M$ recalling that

$$V_g(t_i) := (\ker(\gamma(t_i)) \oplus \mathrm{halo}(\gamma(t_i)))^{\omega} \in SG_{2n-2}(V),$$

where $SG_{2k}(V) = \{W \in \mathrm{Gr}_{2k}(V) \,|\, \omega|_W$ is nondegenerate$\}$ is the Grassmannian of $2k$ dimensional symplectic subspaces of V for $0 \leq k \leq n$. Assuming that both Φ_i and $\phi_i(t)$ have semi-simple eigenvalues[23] for all $0 \leq i \leq M$ and $t \in (t_i, t_i+\epsilon] \cup [t_{i+1}-\epsilon, t_{i+1})$ for some $\epsilon > 0$ then we see from [47] that the eigenvalues of each $\phi_i(t)$ (with repetition) form a continuous path approaching t_i, t_{i+1} in the space of unordered $\mathbb{C}$-tuples (as in the proof of theorem V.2.1), $(\lambda)_i^- : (t_i, t_i + \epsilon] \to \widetilde{\mathbb{C}}^{2n} := \mathbb{C}^{2n}/S_{2n}$ and $(\lambda)_i^+ : [t_{i+1} - \epsilon, t_{i+1}) \to \widetilde{\mathbb{C}}^{2n}$.

Without loss of generality we will be working with the connected component $(t_i, t_i + \epsilon]$ and in doing so will use $(\lambda)_i$ to refer to the path $(\lambda)_i^-$ above as defined above. We let $\hat{\lambda}_i \in \widetilde{\mathbb{C}}^{2n}$ denote the eigenvalues of Φ_i and recall that the topology induced on the space $\widetilde{\mathbb{C}}^{2n}$ is homeomorphic to $\mathbb{C}^{2n}$ and is generated by the metric defined in equation (V.2.1);

$$d((\lambda), (\tau)) = min_{\sigma \in S_{2n}} max_{i \leq 2n} |\lambda_{\sigma(i)} - \tau_i|.$$

Then since $max_{i \leq 2n} |\lambda_i| \leq |z| = |\sigma z|$ for all $z \in \mathbb{C}^{2n}$ we see that $|*| : \widetilde{\mathbb{C}}^{2n} \to \mathbb{R}$ is a well defined continuous function. This implies that any (λ) is bounded if and only if $|(\lambda_j)_{j=1}^{2n}| < \infty$ for any ordering $(\lambda_j)_{j=1}^{2n}$ of (λ).

This bornological structure allows us to distinguish certain subsets of eigenvalues,

[23] Alternatively the more restrictive assumption that the eigenvalues be distinct in a neighborhood of $\{t_i, t_{i+1}\}$ would also suffice as both subsets are open and dense in $\mathrm{Sp}(V)$.

namely for sufficiently small ϵ the following map is well defined for all $t \in (t_i, t_i + \epsilon]$,

$$\begin{aligned} i_{s,g}(t) : \mathbb{C}^n/S_n &\rightarrow \mathbb{C}^2/S_2 \times \mathbb{C}^{2n-2}/S_{2n-2} \\ \lambda(t) &\mapsto (\lambda_s(t), \lambda_g(t)) \end{aligned}$$

where the $\lambda_s(t)$ is defined as the subset $\{\lambda_i\}_{i=1}^2 \subset (\lambda)_i$ for which given any $\delta > 0$ there exists a $t \in (t_i, t_i + \epsilon]$ such that either $|\lambda_i| < \delta$ or $|\lambda_i| > \delta^{-1}$. In other words λ_s consists of those eigenvalues quadruples/pairs which are unbounded[24] (alternatively vanishing) as $t \in t_i^+$ such that $|\lambda_s(t)| \to \infty$ as $t \to t_i^+$. On the other hand $\lambda_g(t)$ consists of those eigenvalues which remain bounded and non-zero as $t \to t_i^+$. Consequently the $\lambda_g(t)$ converges to a unique set of (unordered) eigenvalues $\lambda_g(t_i) = \lim_{t \to t_i^+} \lambda_g(t)$ which must be identical to $\hat{\lambda}_i = spec(\Phi(t_i)) \in \mathbb{C}^{2n-2}/S_{2n-2}$ due to continuity.

Now consider a continuous map $r : [t_i, t_i + \epsilon] \to [1, \infty)$ for which $\lambda_g(t) \subset \overline{D}^2_{r(t)}$ for each $t \in [t_i, t_i + \epsilon]$ (where D_r^2 is the open disk of radius r), which is possible since $|\lambda_g(t)|$ is bounded as $t \to t_i^+$. We note from the symmetry of symplectic eigenvalue quadruples that $r(t)$ is also large enough such that $\lambda_g(t)$ is contained in the open annulus $A_{r(t)}$ for all $t \in [t_i, t_i + \epsilon]$ (implying that λ_s is contained in the complement) where $A_r = \left\{ z \in \widetilde{\mathbb{C}}^{2n} \,|\, r^{-1} < z < r, \right\}$ for any $r > 1$. Then since the eigenvalue path $\lambda_g(t)$ consists of semi-simple eigenvalues which don't vanish and are bounded as $t \to t_i^+$ then as shown in [47] for sufficiently small ϵ there exists a continuous path of symplectic eigenspaces $E_g : [t_i, t_i + \epsilon] \to SG_{2n-2}(V)$ (symplectic as each eigenspace of

[24] This could also suffice as a definition by letting λ_s be the smallest subset of eigenvalue quadruples/pairs of (λ) for which $|\lambda_s|$ diverges.

$\phi_i(t)$ is symplectic) which correspond to the λ_g. Additionally these are the image of a contour integral of the continuous resolvent, $P_g : [t_i, t_i + \epsilon] \times V \to E_g(t)$. for which $E_g(t_i) = V_g(t_i)$.

We note an analogous total eigenspace $E_s(t)$ for $\lambda_s(t)$ does not necessarily exist since that set contains eigenvalues which merge at the exceptional point $\lambda = 0$. We may write the explicit form of the graph resolvent if we consider the (unordered) $2n$-fold product of the circle of radius R,

$$\widetilde{S^1_R} := (S^1_R \overset{2n \text{ times}}{\times \cdots \times} S^1_R)/S_{2n} \subset \widetilde{\mathbb{C}}^{2n},$$

which is well defined since $|\sigma x| = |x|$ for all $\sigma \in S_{2n}$, i.e. the subset is closed under permutations.

Then we may give each $P_g(t) : V \to E_g(t)$ as,

$$P_g(t, v) = \frac{1}{2\pi i} \int_{\partial A_t} (\phi_i(t)(v) - zId)^{-1} dz,$$

where $\partial A_t = S^2_{r(t)} \cup \overline{S^2}_{r(t)^{-1}}$ is the boundary of the annulus with outer radius $r(t)$ and inner radius $r^{-1}(t)$ (where the bar indicates a reversed orientation).

Now from theorem V.2.1 and $E^l_g = E_g(a_l)$ for all sufficiently large l we see that $E_g(t_i) = V_g(t_i)$ so that if we let $j_i : V \to V_g(t_i)$ denote the projection along $V_s(t_i)$ onto $V_g(t_i)$ (where $V_s(t_i) = \ker(\gamma(t_i)) \oplus \mathrm{halo}(\gamma(t_i)) \cong V_g^\omega$) then the subset of $W \in SG_{2n-2}(V)$ for which $(j_i|_W, \omega|_W)$ is a symplectic isomorphism is open and dense in $SG_{2n-2}(V)$; $j_i|_W$ is an isomorphism if and only if $W \pitchfork V_s(t_i)$.

If the above is satisfied then the following diagram commutes where both the top and bottom sequences, $V_s(t_i) \hookrightarrow E_g(t) \overset{P_i}{\twoheadrightarrow} V_g(t_i)$ and $V_s(t_i) \hookrightarrow V \overset{P_i}{\twoheadrightarrow} V_g(t_i)$ are exact.

$$\begin{tikzcd}
 & E_g(t) \arrow[dr, "P_i|_{E_g(t)}"] \arrow[dd, hook] & \\
V_s(t_i) \arrow[ur, hook] \arrow[dr, hook] & & V_g(t_i) \\
 & V \arrow[ur, "P_i"] &
\end{tikzcd}$$

Figure VI.1: All injective maps are inclusions and the symplectic form is preserved throughout.

It follows that for sufficiently small ϵ the path of symplectic isomorphisms $j|_{E_g(t)} =: I_i(t) : E_g(t) \to V_g(t_i)$ induces a unique continuous path of symplectic maps $\beta_i : [t_i, t_i + \epsilon] \to \mathrm{Sp}(V_g(t_i))$ such that $I(t_i) = Id \in \mathrm{Sp}(V_g(t_i))$ (implying that $\beta_i(t_i) = \Phi_i$). Additionally we see that the following holds for all $t \in [t_i, t_i + \epsilon]$,

– $I(t) \circ \phi_i(t)|_{E_g(t)} \circ I_i(t)^{-1} = \beta_i(t)$.

– $(\rho \circ \phi_i|_{E_g(t)})(t) = (\rho \circ \beta_i)(t)$.

We are about ready to construct our bounded approximation $\tilde{\phi}_i$ but first we apply the above to the interval $[t_{i+1} - \epsilon, t_{i+1})$ where we let $F_g : [t_{i+1} - \epsilon, t_{i+1}] \to SG_{2n-2k}(V)$ denote the continuous path of eigenspaces, $\hat{I}_i(t) : F_g(t) \to V_g(t_{i+1})$ the family of isomorphisms and $\hat{\beta}_i : [t_{i+1} - \epsilon, t_{i+1}] \to \mathrm{Sp}(V_g(t_{i+1}))$ the associated symplectic path. It follows that $\beta_i|_{E_g(t)}$ and $\hat{\beta}_i|_{F_g(t)}$ converge to $\Phi_i \in \mathrm{Sp}(V_g(t_i))$ and $\Phi_{i+1} \in \mathrm{Sp}(V_g(t_{i+1}))$ as $t \to t_i^+, t_{i+1}^-$ respectively.

Recalling the symplectic decomposition $V = E_s^l \oplus E_g^l$, $\phi_l = \alpha_l \oplus \beta_l$ as shown in theorem V.2.1 for all but finitely many l we see for all $\epsilon > 0$ and $N \in \mathbb{N}$ large enough such that $a_N \in (t_i, t_i+\epsilon]$ and $b_N \in [t_{i+1}-\epsilon, t_{i+1})$ that we may fix the pair of elements $\alpha_N \in \mathrm{Sp}(E_s^N)$ and $\hat{\alpha}_N \in \mathrm{Sp}(F_s^N)$ from the divergent sequence of symplectic maps so that $E_s^N \pitchfork E_g(t)$ and $F_s^N \pitchfork F_g(t)$ for all $t \in [t_i, t_i+\epsilon]$ and $[t_{i+1}-\epsilon, t_{i+1}]$ respectively. With these we may now 'pause' either sides' divergent term α_l, α_l' at $l = N \in \mathbb{N}$ large enough such that $\rho(\alpha_N) = \rho(\hat{\alpha}_N) = 1$ and define the 'bounded approximations',

$$\tilde{\phi}_i(t) := \begin{cases} \alpha_N \oplus \phi_i|_{E_g(t)} & t_i \leq t \leq a_N \\ \phi_i(t) & a_N < t < b_N \\ \hat{\alpha}_N \oplus \phi_i|_{F_g(t)} & b_N \leq t \leq t_{i+1}\,. \end{cases}$$

Since $\rho(\alpha_N) = \rho(\alpha_N') = 1$ we see that for all $t \in [t_i.t_{i+1}] = cl(U_i)$ that

$$\rho(\tilde{\phi}_i(t)) = \hat{\rho}(\gamma(t))$$

so that in particular $\rho(\tilde{\phi}_i(t)) = \rho(\phi_i(t))$ for each $t \in U_i$.

Then any pair of lifts, $\hat{\theta}$ of $\hat{\rho}\circ\gamma : [t_i, t_{i+1}] \to \mathbb{R}$ (in $\mathcal{L}_{2n}$) and $\tilde{\theta}$ of $\rho\circ\tilde{\phi}_i : [t_i, t_{i+1}] \to \mathbb{R}$ (in $\mathrm{Sp}(V)$) for which $\hat{\theta}(\frac{t_i+t_{i+1}}{2}) = \tilde{\theta}(\frac{t_i+t_{i+1}}{2}) \in (-\pi, \pi]$ must be identical everywhere; $\hat{\theta}(t) = \tilde{\theta}(t)$ for all $t \in [t_i, t_{i+1}]$. We see,

$$\frac{\hat{\theta}(t_{i+1}) - \hat{\theta}(t_i)}{2\pi} = \frac{\tilde{\theta}(t_{i+1}) - \tilde{\theta}(t_i)}{2\pi},$$

so that $\hat{\Delta}(\phi_i) = \Delta(\tilde{\phi}_i)$ for each $i \leq M$ and the lemma follows from $\hat{\Delta}$'s additivity with respect to concatenation. $\square$

Part VII

Concluding Remarks

VII.1 The extended Mean index on $\mathcal{L}_2$

As mentioned above, Δ may be defined on paths in $\mathrm{Sp}(2n)$ originating at the identity so that since Δ is constant on fixed end-point homotopy classes, $\widetilde{Sp}(2n)$ may be used as the domain instead. This is defined via the association of any γ where $\gamma(0) = Id$ with some $g \in \widetilde{Sp}(2n)$ by setting $g := \tilde{\gamma}(1)$ (where $\tilde{\gamma}$ is the lifted path). This map is clearly onto as $\widetilde{Sp}(2n)$ is path connected and is one-to-one on fixed end-point homotopy classes of paths.

Remark VII.1.1. *The subgroup of $\pi_1(\mathcal{L}_2)$ given in the example below lies in the kernel of $\hat{\Delta}$, in particular when $\hat{\Delta}$ is restricted to the corresponding covering space E there no longer exists some non-trivial loop for which $\hat{\Delta}$ is non-zero (as was the case when defined over all of $\widetilde{\mathcal{L}}_2$). In other words, when defined on E the mean index is zero on a loop in $\mathcal{L}_{2n}$ if and only if that loop, when lifted to E, is trivial. It still remains whether these non-trivial loops arising from the removal of H should be tossed out or if they're too important to be ignored.*

Example VII.1.2. *For* $\mathrm{Sp}(2) \hookrightarrow \mathcal{L}_2$ *we have a nice geometric interpretation:* $\mathrm{Sp}(2) \cong D^2 \times S^1$ *and*

$$\Lambda_2 = S^2 \times S^1/(x,t) \sim (-x,-t) \cong \overline{D^2} \times S^1/\sim$$

where $\sim$ identifies boundary points which are antipodal with respect to only the S^1 term (i.e. $(1,\theta,t) \sim (1,\theta,t+\pi)$). Then

$$\mathcal{L}_2 \cong (S^2 \times S^1/\sim) \setminus \{(0,\theta,\pi/2) \sim (0,\theta+\pi,3\pi/2)\}$$

has $\pi_1(\mathcal{L}_2) = \mathbb{Z}[\eta] * \mathbb{Z}[\tau]$ *where* $[\eta]$ *is the push-forward of the generator for* $\pi_1(\mathrm{Sp}(2))$ *and* τ *corresponds to a loop about the missing circle.*

The minimal covering space $E \underset{Pr}{\to} \mathcal{L}_2$ *is the one satisfying the property that* $\mathrm{Pr}_*(\pi_1(E)) = \left\{\eta^l \tau \eta^{-l} \,|\, l \in \mathbb{Z}\right\}$ *which is given by*

$$E = (\overline{D^2} \times \mathbb{R} \underset{\phi:\, \partial\circlearrowleft}{\sqcup} \overline{D^2} \times \mathbb{R}) \setminus K$$

where the set $K = \partial D^2 \times (\frac{\pi}{2} + \pi\mathbb{Z})$ *and* ϕ *identifies the boundaries of the two solid cylinders via* $(x, t) \sim (x, t + \pi)$.

Generated using Mathematica 10 © [86]

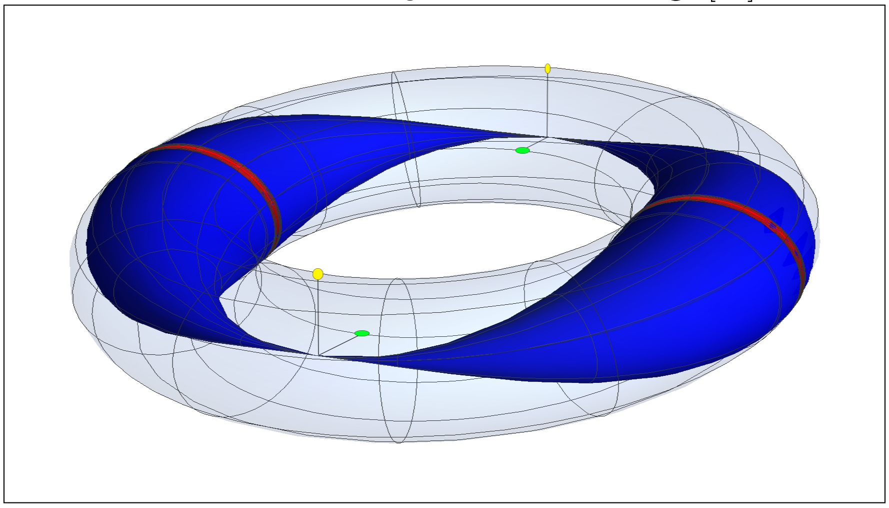

Figure VII.1: The interior of the solid torus corresponds to $\Lambda_2^0 \cong \mathrm{Sp}(2)$ while the rest of Λ_2, namely Λ_2^1, is the quotient of the boundary torus using the aforementioned identification, depicted here by two pair of like colored ellipses (note the lack of rotation in the minor radius). The exceptional set H is depicted in red (drawn at both identified components for clarity), connecting the co-oriented surface of parabolic transformations in blue at a shared ‘circle at infinity’.

An alternative way of visualizing the above is using a quotient of the trivial lens space $S^2 \times S^1$, *roughly described as identifying the equators of antipodal (with respect* *to S^1) spheres where each hemisphere, in terms of the above figure, corresponds with the disc swept out by the minor radius at some fixed angle of the outer radius.*

VII.2 Smooth Canonical Relations

Our extended mean index encounters significant issues even with linear canonical relations, and in particular to apply our Lagrangian mean index to the tangent projections of paths along Lagrangian submanifolds will impose significant constraints on said submanifold, two of which follow below would be necessary to define the mean index of a path on a Lagrangian submanifold.

1. First we require (M, ω) be a symplectic manifold with $\dim(M) = 4l$ for some $l \in \mathbb{N}$ and let L be a Lagrangian submanifold. Then there must be a smooth bundle decomposition $TM = B \oplus C$ where $B \to M$ and $C \to M$ are $2l$ dimensional real vector bundles modeling the source and target equipped with some unique bundle isomorphism $B \cong C$ without which critical concepts such as $\ker(T_x L)$ would be meaningless. It is common to consider smooth relations $M \times M$ for symplectic manifolds M for which this property holds by construction.

2. Next we would need that $T_x L \notin \mathrm{LagGr}(T_x M) \setminus H_x$ for all $x \in L$ (or at least for all $x \in \gamma(I)$, that is the path we are linearizing must stay away from such points) where $H_x = \{L \in \mathrm{LagGr}(T_x M) \,|\, \pi_{B_x}(L) \cap (L \cap (\{0\}_x \times C_x)) \neq \{0\}_x\}$ as well as some further transversality conditions as detailed in [84] (provided we intend

on composing our canonical relations). This leads to the question of whether there exists any global obstructions to a Lagrangian submanifold satisfying these conditions, as well as the possibility that such Lagrangian submanifolds may be rare or non-existent for large classes of manifolds satisfying the first property.

3. Many of the proofs referenced in the introduction rely on Δ being a quasimorphism, and although we have partially recovered a mimicry of the property, the associated group theoretic implications are still lost and perhaps are only recoverable in a more general algebraic setting (e.g. path groupoids).

The question of which smooth canonical relation framework to work in is also an immediate question. Perhaps the most promising is the Wehrheim-Woodward method applied to Lagrangian relations found in [83] and [50]. The *highly selective category* **WW(SLREL)** (far too intricate to describe here) may be described in part by the techniques used in the linear case. Namely, the authors replace the usual Lagrangian Grassmannian with the 'indexed Lagrangian Grassmannian' $\mathrm{LagGr}_\bullet(V)$. As a set $\mathrm{LagGr}_\bullet(V) = \Lambda_{2n} \times \mathbb{N}$ albeit equipped with a topology quite different from the product topology (which they call the Sabot topology). They use a discrete metric $d(L, L') = \mathrm{codim}(L \cap L')$ (codimension relative to L or L') to define a partial order $(L, k) \preceq (L', k') \Leftrightarrow d(L, L') \leq k' - k$ used to produce a basis for the topological space $(\Lambda_{2n} \times \mathbb{N}, \mathcal{T})$ under which the above function becomes continuous, albeit at the very heavy price of not even satisfying the T_1 separation axiom (though fortunately it does remain T_0).

Additionally the properties established regarding *composable tuples* of Lagrangian relations, $(L_1, L_2, \ldots, L_n)$ may in particular yield information in the context of time-dependent flows where the iterated return maps (relations) may be distinct. Another benefit of the potential application of the mean index in **WW(SLREL)** is the ability to coherently form a composition of distinct $L, L' \in \mathcal{L}_{2n}$ (in a manner avoiding the rather ad hoc procedures found in the proofs of section VI) so that bounds of the type $|\hat{\Delta}(L \circ L') - \hat{\Delta}(L) - \hat{\Delta}(L')|$ may be defined in a more natural way.

Alternatively one might use the extended mean index to define the mean index of unbounded paths of symplectomorphisms in graph to $L \in \mathcal{L}_{2n}$, perhaps near unbounded punctures of pseudoholomorphic curves or along Hamiltonian flows on open manifolds.

VII.3 Refining the Notion of an Exceptional Lagrangian $\hat{H} \subset H$

Proposition VII.3.1. *Such a V_g as described in lemma III.2.5.1 with an associated unique $\phi \in \mathrm{Sp}(V_g)$ exist if dim$((L_1^\omega \oplus L_2^\omega) \cap (L_1 \cap L_2)) \leq r =$ dim$(L_1^\omega \cap L_2^\omega)$ where* $\dim(L_1 \cap L_2) = 2n - 2k + r$ *for some $0 \leq r \leq k$. In particular this shows that the hypothesis of theorem III.2.5 is not a necessary one for some $L \in \Lambda_{2n}$ to possess a uniquely determined graph portion. See equation (II.3.1) for the isotropic pair invariants used above.*

Proof. We first claim that such a V_g exists when

$$\dim((L_1 \cap L_2)/(L_1^\omega + L_2^\omega) \cap (L_1 \cap L_2)) \geq 2n - 2k.$$

It is important to note that $L_1 \cap L_2$ is no longer necessarily symplectic. The above certainly implies the existence of a $V_g \leq L_1 \cap L_2$ such that $\dim(V_g) = 2n - 2k$ as well as condition (3) of the proof for theorem III.2.5, that is $V_g \cap L_i^\omega = \{0\}$ since $V_g \subset L_1 \cap L_2$. Whether V_g as constructed is still a viable method for producing this isomorphism remains to be shown. Since $(L_1 \cap L_2)/((L_1^\omega + L_2^\omega) \cap (L_1 \cap L_2))$ carries a unique reduced symplectic form $\omega_{red}([v], [w]) = \omega(v, w)$ we may choose symplectic $\hat{V}_g \leq (L_1 \cap L_2)/((L_1^\omega + L_2^\omega) \cap (L_1 \cap L_2))$ with dimension $2n - 2k$ and choose $V_g \leq L_1 \cap L_2$ such that $V_g \underset{\pi_{1,2}}{\cong} \hat{V}_g$. Then since $\pi_{1,2}^* \omega_{red} = \omega$ on $L_1 \cap L_2 \leq V$ it follows that V_g will be a symplectic subspace of V. Thus the above bound guarantees the existence of V_g such that $\phi \in \mathrm{Sp}(V_g)$.

We observe that

$$\begin{aligned}
\dim((L_1 \cap L_2)/(L_1^\omega + L_2^\omega) \cap (L_1 \cap L_2)) &\geq 2n - 2k \Leftrightarrow \\
(2n - 2k + r) - \dim((L_1^\omega + L_2^\omega) \cap (L_1 \cap L_2)) &\geq 2n - 2k \Leftrightarrow \\
\dim((L_1^\omega + L_2^\omega) \cap (L_1 \cap L_2)) &\leq r = \dim(L_1^\omega \cap L_2^\omega).
\end{aligned}$$

In the context of Lorand's classification equation (II.3.1) we see that $\kappa = 0$ (i.e. $L \in H$) implies $r = 0$ and indeed $\dim((L_1^\omega \oplus L_2^\omega) \cap (L_1 \cap L_2)) = \dim((\ker(L) \oplus \mathrm{halo}(L)) \cap V_g) = 0$. $\square$

Despite the continuity transversality requirement (namely $\ker(L) \cap \mathrm{halo}(L) = \emptyset$)

being equivalent to $r = 0$, we see that here a larger r seems to add room for potential domains to fit in $L_1 \cap L_2$ and survive the quotienting process retaining a dimension of $2n - 2k$. As $r \leq \kappa$ and $\kappa - r \leq n - k$, it appears that a small $\kappa - r$ (perhaps even $\kappa = r$) is preferred in this context (expressing the symplectic quotient map in the fiber as an explicit symplectic map on some subspace of V). This is fairly surprising when compared to the traditional transversality requirement that $r = 0$ with no concern for κ required to ensure continuity of composition.

Example VII.3.2. *An explicit example of the existence of Lagrangian L where $\kappa(L) \neq 0$, yet ϕ is uniquely determined follows below. Due to the nature of the bounds imposed on the invariants found in equation* (II.3.1) *such an L exists only when $n \geq 3$, indeed since any $L \in \Lambda_4^1$ has either $r = 0$ or $r = 1$ then either $L \notin H$ or* $\mathrm{dom}(L) = \mathrm{ran}(L)$ *respectively. Let*

$$L = \langle (e_1, 0), (e_2, 0), (e_3, f_3), (f_3, e_3), (0, e_1), (0, f_2) \rangle \leq \mathbb{R}^6 \times \overline{\mathbb{R}^6}$$

where $(e_i, f_i)_{i=1}^3$ is a Darboux basis and L is Lagrangian.

We observe that $L_1 = \langle e_1, e_2, e_3, f_3 \rangle$, $L_2 = \langle e_1, e_3, f_2, f_3 \rangle$, $L_1^\omega = \langle e_1, e_2 \rangle$ and $L_2^\omega = \langle e_1, f_2 \rangle$. Thus $\dim(L_1 \cap L_2^\omega) = 1$ yet $V_g = \langle e_3, f_3 \rangle \leq L_1 \cap L_2$ satisfies $V_g \cap L_i^\omega = \{0\}$ for $i = 1, 2$. Additionally we see that $\phi = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ under basis (e_3, f_3). Verifying with the above lemma we see that indeed $\dim((L_1^\omega + L_2^\omega) \cap (L_1 \cap L_2)) = 1 \leq r$ since $r = \dim(L_1^\omega \cap L_2^\omega) = 1$.

Conjecture VII.3.3. For each $L \in \Lambda_{2n}$ consider the tuple of isotropic pair invariants

$(k_L, r_L, \kappa_L) = (\dim(\ker(L)), \dim(\ker(L) \cap \mathrm{ran}(L)), \dim(\mathrm{dom}(L) \cap \mathrm{ran}(L)))$ associated to $(\ker(L), \mathrm{halo}(L)) \in \mathcal{I}_{k_L}$ as shown in [59]. Letting $\hat{H} := \{L \in \Lambda_{2n} \,|\, 0 < r_L < \kappa_L\}$ then ρ may be continuously extended to $\hat{\mathcal{L}}_{2n} := \Lambda_{2n} \setminus \hat{H}$.

We see that since $r_L \leq \kappa_L$ by definition and in this notation $H = \{L \in \Lambda_{2n} \,|\, 0 < r_L\}$ that $\hat{H} \subset H$ and $\mathcal{L}_{2n} \subset \hat{\mathcal{L}}_{2n}$. It is unlikely $\mathrm{codim}(\hat{H}) > 2$ although the added elements in $\hat{\mathcal{L}}_{2n}$ may have another benefit. We conclude the thesis with the very essence of anti-climax; the example justifying the extension of the square of ρ as referenced in remark I.3.15. Namely, we exhibit the failure of continuity for ρ when extending to $\mathcal{L}_{2n}$ even for those symplectic maps with all real eigenvalues.

Example VII.3.4. *Given a Darboux basis $(e_i, f_i)_{i=1}^n$ for a symplectic vector space V consider the two following sequences of symplectic matrices,*

$$\begin{aligned} A_k &:= \mathrm{Diag}(\underbrace{1/k, \ldots, 1/k}_{n \text{ times}}, k, \ldots, k) \\ B_k &:= \mathrm{Diag}(-1/k, \ldots, 1/k, -k, \ldots, k). \end{aligned}$$

We observe for each k that $\rho(A_k) = \prod_{i=1}^n 1 = 1$ and $\rho(B_k) = -\prod_{i=2}^n 1 = -1$, i.e. each A_k is positive hyperbolic and each B_k negative hyperbolic. Then it's simple to see that as $k \to \infty$ we have,

$$\begin{aligned} \mathrm{Gr}(A_k) &= \langle (e_1, \frac{e_1}{k}), \ldots, (e_n, \frac{e_n}{k}), (f_1, kf_1), \ldots, (f_n, kf_n) \rangle \\ &\to \langle (e_1, 0), \ldots, (e_n, 0), (0, f_1), \ldots, (0, f_n) \rangle \\ \mathrm{Gr}(B_k) &= \langle (e_1, \frac{-e_1}{k}), \ldots, (e_n, \frac{e_n}{k}), (f_1, -kf_1), \ldots, (f_n, kf_n) \rangle \\ &\to \langle (e_1, 0), \ldots, (e_n, 0), (0, -f_1), \ldots, (0, f_n) \rangle. \end{aligned}$$

Thus both $\mathrm{Gr}(A_k) \to L \leftarrow \mathrm{Gr}(B_k)$ *yet* $\rho(A_k) = 1 \neq -1 = \rho(B_k)$ *while* $\rho^2(A_k) = 1 = \rho^2(B_k)$ *for all* $k \in \mathbb{N}$.

Bibliography

[1] V. I. Arnol'd, Proof of a Theorem of A. N. Kolmogorov on the Preservation of Conditionally Periodic Motions under a Small Perturbation of the Hamiltonian, *Uspehi Mat. Nauk* **18** (1963), no. 5 (113), 13–40 (Russian). *Russian Math. Surveys* **18** (1963), no. 5, 9–36 (English).

[2] V. I. Arnol'd, O Harakteristicheskom Klasse Vhodjashhem v Uslovija Kvantovanija [Characteristic class Entering in Quantization Conditions], *Funktsional. Anal. i Prilozhen.* **1** (1967), no. 1, 1–14, (Russian edition). *Funct. Anal. Appl.*, **1**, (1967), no. 1, 1–13, (English edition).

[3] V. I. Arnol'd, S. M. Gusein-Zade, V. I. Varchenko, Singularities of Differentiable Maps Vol. I, *Birkhäuser, Monographs of Math.* **82** (1985).

[4] J. Barge, E. Ghys, Cocycles d'Euler et de Maslov [Euler and Maslov cocycles], *Mathematische Annalen.* **294** (1992), 235–265.

[5] G. Ben Simon, D. Salamon, Homogeneous Quasimorphisms on the Symplectic Linear Group, *Isr. Jour. of Math.* **175** (2010), 221–224.

[6] S. Benenti, W. Tulczyjew, Symplectic Linear Relations, *Mem. Accad. Sci. Torino Cl. Sci. Fis. Mat. Natur.* **5** (1981), no. 4, 71–140.

[7] E.Binz, G.Schwarz, The Principle of Virtual Work and a Symplectic Reduction of Non-Local Continuum Mechanics, *Rep. on Math. Phys.* **32** (1993), no. 1, 49–69.

[8] R. Bott, On the Iteration of Closed Geodesics and the Sturm Intersection Theory, *Comm. on Pure and Appl. Math.* **9** (1956), 171–206.

[9] K. Cieliebak, A. Floer, H. Hofer, K. Wysocki, Applications of Symplectic Homology II: Stability of the Action Spectrum, *Math. Zeit.* **223** (1996), no. 1, 27–45.

[10] E. Çineli, Conley Conjecture and Local Floer Homology, *Arch. Math.* **111** (2018), no. 6, 647–656.

[11] E. Çineli, V. Ginzburg, B. Gürel, Pseudo-rotations and Holomorphic Curves, Preprint *arXiv:1905.07567v2 [math.SG]* (2019).

[12] C. Conley, E. Zehnder, The Birkhoff-Lewis Fixed Point Theorem and a Conjecture of V. I. Arnold, *Invent. Math.* **73** (1983), no. 1, 33–49.

[13] C. Conley, E. Zehnder, Morse–type Index Theory for Flows and Periodic Solutions for Hamiltonian Equations, *Comm. Pure Appl. Math.* **37** (1984), no. 2, 207–253.

[14] J. Eells, A Setting for Global Analysis, *Bull. Amer. Math. Soc.* **72** (1966), 751–807.

[15] I. Ekeland, H. Hofer, Convex Hamiltonian Energy Surfaces and Their Periodic Trajectories, *Comm. Math. Phys.* **113** (1987), no. 3, 419–469.

[16] Ekeland, I.: Une thèorie de Morse pour les systémes Hamiltoniens convexes [A Morse Theory for Convex Hamiltonian Systems], *Ann. Inst. H. Poincaré, Anal. non linraire* **1** (1984), no. 1, 19–78, (French edition).

[17] Y. M. Eliashberg, The Theorem on the Structure of Wave Fronts and its Applications in Symplectic Geometry, *Functional Analysis and Its Applications* **21** (1987), no. 3, 227–232.

[18] A. Floer, Holomorphic Curves and a Morse Theory for Fixed Points of Exact Symplectomorphisms, *Aspects dynamiques et topologiques des groupes infinis de transformation de la mécanique (Lyon, 1986)*, 49–60, *Travaux en Cours* **25**, *Hermann, Paris*, (1987).

[19] A. Floer, A Refinement of the Conley Index and an Application to the Stability of Hyperbolic Invariant Sets, *Ergodic Theory Dynam. Systems* **7** (1987), no. 1, 93–103.

[20] A. Floer, A Relative Morse Index for the Symplectic Action, *Comm. Pure Appl. Math.* **41** (1988), no. 4, 393–407.

[21] A. Floer, The Unregularized Gradient Flow of the Symplectic Action, *Comm. Pure Appl. Math.* **41** (1988), no. 6, 775–813.

[22] A. Floer, Morse Theory for Lagrangian Intersections, *J. Differential Geom.* **28** (1988), no. 3, 513–547.

[23] A. Floer, Witten's Complex and Infinite-Dimensional Morse Theory, *J. Differential Geom.* **30** (1989), no. 1, 207–221.

[24] A. Floer, Symplectic Fixed Points and Holomorphic Spheres, *Comm. Math. Phys.* **120** (1989), no. 4, 575–611.

[25] I. M. Gelfand, V. B. Lidskii, O Strukture Oblastej Ustojčivosti Linejnyh Kanoničeskih Sistem Differencialnyh Uravnenij s Periodičeskimi Koèfficientami [On the Structure of the Regions of Stability of Linear Canonical Systems of Differential Equations with Periodic Coefficients], *Uspekhi Mat. Nauk* **10** (1955), no. 1(63), 3–40, (Russian edition). *Amer. Math. Soc. Transl.* **2** (1958), no. 8, 143–181, (English edition).

[26] V. Ginzburg, B. Gürel, Action and Index Spectra and Periodic Orbits in Hamiltonian Dynamics, *Geom. Topol.* **13** (2009), no. 5, 2745–2805.

[27] V. Ginzburg, B. Gürel, Local Floer Homology and the Action Gap, *Jour. of Symplectic Geo.* **8** (2010), no. 3, 323–357.

[28] V. Ginzburg, B. Gürel, The Conley Conjecture and Beyond, *Arnold Math. Jour.* **1** (2015), no. 3, 299–337.

[29] V. Ginzburg, B. Gürel, Pseudo-rotations vs. Rotations, Preprint *arXiv:1812.05782v1 [math.SG]* (2018).

[30] V. Ginzburg, B. Gürel, Conley Conjecture Revisited, *Int. Math. Res. Not. IMRN* (2019), no. 3, 761–798.

[31] V. Ginzburg, B. Gürel, Lusternik-Schnirelmann Theory and Closed Reeb Orbits, *Math. Zeit.* **295** (2020), no. 1-2, 515–582.

[32] Y. Gören, Counting Periodic Orbits: Conley Conjecture for Lagrangian Correspondences and Resonance Relations for Closed Reeb Orbits (PhD Thesis), https://ginzburg.math.ucsc.edu/YusufGorenThesis.pdf (2015).

[33] M. Gromov, Pseudo-holomorphic curves in Symplectic Manifolds, *Invent. Math.* **82** (1985), no. 2, 307–347.

[34] V. Guillemin, S. Sternberg, Some problems in Integral Geometry and some Related Problems in Micro-local Analysis, *Amer. Jour. of Math.* **101** (1979), no. 4, 915–955.

[35] J. Gutt, The Conley-Zehnder Index for a Path of Symplectic Matrices, Preprint *arXiv:1201.3728v2 [math.DG]*, (2012).

[36] J. Gutt, Generalized Conley-Zehnder index, *Ann. Fac. Sci. Toulouse Math. (6)* **23** (2014), no. 4, 907–932.

[37] W. Hamilton, On a General Method in Dynamics, *Phil. Trans. R. Soc.* **124** (1834), 247–308.

[38] A. Hatcher, Algebraic Topology, *Cambridge Univ. Press* (2002), 56–78.

[39] H. He, Compactification of Classical Groups, *Comm. Anal. Geom.* **10** (2002), no. 4, 709–740.

[40] G. A. Hedlund, Poincaré's Rotation Number and Morse's Type Number, *Trans. Amer. Math. Soc.* **34** (1932), no. 1, 75–97.

[41] H. Hofer, On Strongly Indefinite Functionals with Applications, *Trans. of the Amer. Math. Soc.* **275** (1983), no. 1, 185–214.

[42] H. Hofer, K. Wysocki, E. Zehnder, Properties of Pseudo-Holomorphic Curves in Symplectisations II: Embedding controls and algebraic invariants. *Geom. Funct. Anal.* **5** (1995), 270–328.

[43] L. Hörmander, Fourier Integral Operators I, *Acta Mathematica* **127** (1971), no. 1-2, 79-183.

[44] M. Hutchings, Quantitative Embedded Contact Homology, *J. Differential Geom.* **88** (2011), no. 2, 231–266.

[45] M. Hutchings, Beyond ECH Capacities, *Geom. Topol.* **20** (2016), no. 2, 1085–1126.

[46] S. Janeczko, Lagrangian Submanifolds in Product Symplectic Spaces, *Journal of Math. Phys.* **41** (2000), no. 8, 5642–5655.

[47] T. Kato, Perturbation Theory for Linear Operators, *Springer Berlin Heidelberg*, *Classics in Mathematics* **132** (1995).

[48] M. Klein, Mathematical Thought from Ancient to Modern Times, Vol. I and II, *New York: Oxford University Press.* (1972), 167–168.

[49] A. N. Kolmogorov, On Conservation of Conditionally Periodic Motions for a Small Change in Hamilton's Function, *Dokl. Akad. Nauk SSSR* **98** (1954), 527–530.

[50] D. Li-Bland, A. Weinstein, Selective Categories and Linear Canonical Relations, *Symmetry Integrability Geom. Methods Appl.* **10** (2014).

[51] C. Liu and Y. Long, An Optimal Increasing Estimate for Maslov-type Indices of Iterations, *Chinese Sci. Bull.* **43** (1997), 2275–2278.

[52] Y. Long, Multiple Solutions of Perturbed Superquadratic Second Order Hamiltonian Systems, *Trans. Amer. Math. Soc.* **311** (1989), no. 2, 749–780.

[53] Y. Long, Forced Oscillations for the Triple Pendulum, *Analysis, et cetera* (1990), 177–208, *Academic Press, Boston, MA.*

[54] Y. Long, Maslov-type Index, Degenerate Critical Points, and Asymptotically Linear Hamiltonian Systems, *Sci. China Ser. A* **33** (1990), no. 12, 1409–1419. (English edition).

[55] Y. Long, A Maslov-type Index Theory for Symplectic Paths, *Top. Meth. Nonl. Anal.* **10** (1997), 47–78.

[56] Y. Long, Index Iteration Theory for Symplectic Paths with Applications to Nonlinear Hamiltonian Systems, *Proc. of the Int. Cong. of Math., Vol. II* (Beijing, 2002), 303–313, *Higher Ed. Press*, (2002).

[57] Y. Long, Index Theory for Symplectic Paths with Applications, *Progress in Mathematics* **207**, Birkhäuser Verlag, Basel, (2002).

[58] Y. Long, E. Zehnder, Morse theory for Forced Oscillations of Asymptotically Linear Hamiltonian Systems, *Stoc. Proc. Phys. and Geom.* (Ascona and Locarno, 1988), World Scientific, Teaneck (1990), 528–563.

[59] J. Lorand, Classifying Linear Canonical Relations (Master's Thesis), *arXiv:1508.04568v1 [math.SG]*, (2014).

[60] J. Lorand, Alan Weinstein, Coisotropic Pairs, Preprint *arXiv:1408.5620v1 [math.SG]*, (2014).

[61] K. Mackenzie, On Extensions of Principal Bundles, *Ann. of Global Anal. and Geo.* **6** (1988), 141–163.

[62] V. P. Maslov, Perturbation Theory and Asymptotic Methods, *Moscow University Press, Moscow*, (1965), (Russian edition).

[63] V. P. Maslov, V. I. Arnol'd, V. S. Buslaev, Théorie des Perturbations Et Méthodes Asymptotiques, *Etudes math.*, Dunod, (1972).

[64] D. McDuff, F. Schlenk, The Embedding Capacity of 4-dimensional Symplectic Ellipsoids, *Ann. of Math. Second Series* **175** (2012), No. 3, 1191–1282.

[65] J. Milnor, On the Existence of a Connection with Curvature Zero, *Comment. Math. Helv.* **32** (1958), 215-223.

[66] M. Morse, The Calculus of Variations in the Large, (Reprint of the 1932 original), *Amer. Math. Soc. Col. Pub.* **18** (1996).

[67] J. Moser, On Invariant Curves of Area-Preserving Mappings of an Annulus, *Nachr. Akad. Wiss. Gttingen Math.-Phys.* **II** (1962), 1–20.

[68] J. Nelson, Reflections on Cylindrical Contact Homology, *Instit. For Adv. Study*, (14 May 2020), https://youtu.be/t5NiA4Xdgk4, (lecture).

[69] B. Nica, Group Quasimorphisms, Vanderbilt University (2006, rev. 2008), http://www.math.mcgill.ca/bnica/quasimorphisms.pdf, (lecture).

[70] R. Palais, Lusternik-Schnirelmann Theory for Banach manifolds, *Topology* **5** (1966), 115–132.

[71] H. Poincaré, Sur les Courbes Définies par les Équations Différentielles, *Journal de Mathématiques* **4** (1885), vol. 1, 167–244.

[72] P. H. Rabinowitz, Periodic Solutions of Hamiltonian Systems, *Comm. Pure Appl. Math.* **31** (1978), no. 2, 157–184.

[73] J. Rawnsley, On the Universal Covering Group of the Real Symplectic Group, *Jour. of Geo. and Phys.* **62** (2012), no. 10, 2044–2058.

[74] J. Robbin, D. Salamon, The Maslov Index for Paths, *Jour. of Topo.* **32** (1993), no. 4, 827–844.

[75] P. Robinson, J. Rawnsley, The Metaplectic representation, Mp^c Structures and Geometric Quantization, *Mem. of the Amer. Math. Soc.* **81** (1989), no. 410.

[76] D. Salamon, Morse theory, the Conley index and Floer homology, *Bulletin of the London Mathematical Society* **22** (1990), no. 2, 113–140.

[77] D. Salamon, Lectures on Floer Homology, *Symp. Geo. and Topo.* (Park City, UT, 1997), 143–229, *IAS/Park City Math. Ser.* **7**, Amer. Math. Soc., Providence, RI, (1999).

[78] D. Salamon, E. Zehnder, Morse theory for Periodic Solutions of Hamiltonian Systems and the Maslov Index, *Comm. Pure Appl. Math.* **45** (1992), no. 10, 1303–1360.

[79] K. Siegel, Higher Symplectic Capacities, Preprint *arXiv:1902.01490v3 [math.SG]*, (2020).

[80] J. Towber, Linear Relations, *Journal of Algebra* **19** (1971), 1–20.

[81] A. Usher, A History of Mechanical Invention (Revised edition), *Cambridge, MA: Harvard Univ. Press*, (1954).

[82] C. Viterbo, A New Obstruction to Embedding Lagrangian Tori, *Invent. Math.* **100** (1990), no. 2, 301–320.

[83] K. Wehrheim, C. Woodward, Functoriality for Lagrangian Correspondences in Floer theory, *Quantum Topol.* **1** (2010), no. 2, 129–170.

[84] A. Weinstein, Symplectic Categories, *Proc. of Geom. Summer School: Lisbon*, *arXiv:0911.4133v1 [math.SG]*, (2009), (lecture).

[85] S. Weinberger, The Topological Classification of Stratified Spaces, *The Univ. of Chicago Press, Ltd., London*, (1994).

[86] Wolfram Research Inc., Mathematica, Version 10, Champaign, IL, (2013).

[87] V. Yakubovich, V. Starzhinskii, Linear Differential Equations with Periodic

Coefficients (D. Louvish, Trans.), Vol. 1 and 2, *John Wiley & Sons New York-Toronto*, Israel Program for Scientific Translations, (1975).