

UC Santa Cruz

UC Santa Cruz Electronic Theses and Dissertations

Title

On An Extension of the Mean Index to the Lagrangian Grassmannian

Permalink

<https://escholarship.org/uc/item/0qz6v0h6>

Author

Grace, Matthew Isaak

Publication Date

2020

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial License, available at <https://creativecommons.org/licenses/by-nc/4.0/>

Peer reviewed|Thesis/dissertation

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

September 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

Copyright © by

Matthew I. Grace

2020

Table of Contents

List of Figures	v
Abstract	vi
Dedication	vii
Acknowledgments	viii
I Preliminaries	1
I.1 Introduction	2
I.2 Historical Context	8
I.3 Definitions and Conventions	21
I.4 Dissertation Outline	28
I.4.1 Outline of Results	28
I.4.2 Outline of Proofs	31
II Linear Canonical Relations and Isotropic Pairs	34
II.1 Strata of the Lagrangian Grassmannian	35
II.2 Iterating Linear Canonical Relations	37
II.3 Conjugacy Classes of Isotropic Pairs	40
II.4 Stratum-Regular Paths	41
III The Set H of Exceptional Lagrangians	50
III.1 The Codimension of H	51
III.2 Admissible Linear Canonical Relations	54
IV The Conley-Zehnder Index and the Circle Map ρ	62
IV.1 Construction of the Conley-Zehnder Index	63
IV.2 Properties of the Circle Map Extension $\hat{\rho}$	65

V	Unbounded Sequences in the Symplectic Group	69
V.1	A Sufficient Condition for Asymptotic Hyperbolicity	70
V.2	A Decomposition for Certain Unbounded Sequences of Symplectic Maps	76
V.2.1	Preparation for the Proof of Theorem V.2.1	78
V.2.2	Proof of Theorem V.2.1	80
V.3	Extending ρ^2 Continuously to \mathcal{L}_{2n}	87
VI	Properties of $\hat{\Delta}$ over Stratum-Regular Paths	90
VI.1	Continuity of Compatible Path Compositions	91
VI.2	The Quasimorphism Bound	98
VI.3	Proof of Proposition VI.1.3 and Lemma VI.1.5	102
VII	Concluding Remarks	109
VII.1	The Extended Mean Index on \mathcal{L}_2	110
VII.2	Smooth Canonical Relations	114
VII.3	Refining the Notion of an Exceptional Lagrangian	116
	Bibliography	121

List of Figures

I.1 An early example of a co-oriented hypersurface in the symplectic group	15
I.2 Commutative diagram giving an alternate definition of the mean index	26
II.1 Showing each stratum Λ_{2n}^k as a fibration	35
V.1 Commutative diagram showing $\hat{\rho}$ as a continuous extension of ρ^2	87
VI.1 Commutative diagram for proof VI.3	107
VII.1 Figure of Λ_2 as a solid torus with boundary identifications	111
VII.2 A depiction of the fundamental group of \mathcal{L}_2	112
VII.3 A depiction of the fundamental group of a covering space over \mathcal{L}_2	112
VII.4 A depiction of a covering space of \mathcal{L}_2	113

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 the 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 $\mathrm{Gr} : \mathrm{Sp}(V, \omega) \rightarrow \Lambda_{2n}$ sends each linear symplectomorphism to its graph. Our central result is in extending the mean index, using this embedding and a formal construction of the mean index in terms of a map $\rho : \mathrm{Sp}(V, \omega) \rightarrow S^1$, from 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 $\mathrm{codim}(\Lambda_{2n} \setminus \mathcal{L}_{2n}) = 2$. Namely, we continuously extend ρ^2 to what we call $\hat{\rho} : \mathcal{L}_{2n} \rightarrow S^1$ so that by applying the aforementioned construction to $\hat{\rho}$, we reduce the existence and continuity of our extended index to the simpler problem of producing the continuous extension $\hat{\rho}$. Our secondary results concern the algebraic properties of the extended index with respect to a set-theoretic composition operation on Λ_{2n} which extends the usual group structure of $\mathrm{Sp}(V, \omega)$ to that of a monoid on \mathcal{L}_{2n} . To derive these we define an open and dense subset $\mathcal{P}_{reg}(\mathcal{L}_{2n}) \subset C^1([0, 1], \mathcal{L}_{2n})$ equipped with an equivalence relation \sim_{comp} and show that the point-wise composition of any two \sim_{comp} equivalent stratum-regular paths is piece-wise differentiable. We then show that, when restricted to $\mathcal{P}_{reg}(\mathcal{L}_{2n})$, the extended index 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

Acknowledgments

The author would like to thank his advisor Viktor Ginzburg for both his unwavering patience and understanding, as well as the consistent and helpful feedback he has provided throughout the last four years.

Part I

Preliminaries

I.1 Introduction

Throughout this dissertation (V, ω) will be a real symplectic vector space of dimension $2n$ and $\Lambda_{2n} := \text{LagGr}(V \times V, \text{Pr}_1^* \omega - \text{Pr}_2^* \omega)$ will denote¹ 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 $\gamma: [0, 1] \rightarrow \text{Sp}(V)$, 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 $\text{Gr}: \text{Sp}(V) \rightarrow \Lambda_{2n}$ (see definition I.3.3) which identifies $\text{Sp}(V)$ in Λ_{2n} with the open and dense embedded submanifold

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

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 $\text{Im}(\text{Gr}) \cong \text{Sp}(V)$, as both $\text{Im}(\text{Gr})$ 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 $\Lambda_{2n} \setminus \text{Im}(\text{Gr})$ is a hypersurface in Λ_{2n} (see proposition II.1.6) whereas we show in theorem I.4.1 that $\text{codim}(H) = 2$.

The mean index is commonly seen in one of two contexts: as a real valued map over all continuous paths² in $\text{Sp}(V)$ or, as it is found in the bulk of the literature (e.g. [5, 82, 78, 60]), restricted to those paths γ originating at the identity (i.e. $\gamma(0) = \text{Id}$, we will often refer to such paths as ‘identity-based’). Besides the fact that most of the algebraic properties of Δ only

¹Since 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, $\text{Sp}(V)$ and Λ_{2n}).

²Continuity of paths will be an assumption implicitly maintained throughout the dissertation save for paths not known to be continuous, in which case we will specify them as such. Differentiability will also be explicitly invoked throughout, save for those instances in which doing so would be unnecessarily redundant.

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 structure³ on $\widetilde{\mathrm{Sp}}(V)$, over which the mean index Δ may be characterized (as in [5]) 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 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 equivalently⁴, 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 [69]) to construct Δ involves one of two similar lifting procedures in terms of a certain continuous map $\rho : \mathrm{Sp}(V) \rightarrow S^1$, 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

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

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

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 [I.4.3](#) through an application of the aforementioned lifting procedure. Specifically, theorem [I.4.3](#) produces a continuous extension $\hat{\rho} : \mathcal{L}_{2n} \rightarrow S^1$ of ρ^2 by identifying $\mathrm{Sp}(V) \cong \mathrm{Im}(\mathrm{Gr})$ so that by simply repeating Milnor's procedure on \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 larger⁵ 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 [\[5, 28, 30, 11\]](#)) we desire to recover some of these properties, but before discussing this we address certain similarities between this dissertation and the thesis [\[33\]](#).

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 [\[33\]](#), authored by Yusuf Gören. In particular, our theorem [V.1.1](#) matches theorem [2.2.6](#) in [\[33\]](#), though are

⁵This 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.

proved using distinct methods. Our theorem III.2.6 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 [33].

Given the parallels outlined above, this dissertation necessarily exhibits significant 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 part 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 [33] this is circumvented rather cleverly by choosing a representative $\gamma : I \rightarrow \mathcal{L}_{2n}$ 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 $\tau(1) = L \in \mathcal{L}_{2n} \setminus \text{Im}(Gr)$, 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 $\text{Im}(Gr)$ at the very last moment. This means that our continuous extension of ρ^2 in theorem I.4.3 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 [5]) 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 theorem II.4.15), 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

⁶In this case a stratum-regular path is one satisfying a transversality condition (see definition I.4.5) with respect to a certain stratification of Λ_{2n} as given in [48].

quasimorphism-type bound for the aforementioned compatible pairs. Despite this, the similar goals that [33] 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 $\gamma : I \rightarrow \mathcal{L}_{2n}$ 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.4.6 mentioned below) which allows us to show in corollary I.4.7 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 (see proposition VI.1.3) and see in theorem I.4.6 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.4.8 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 [41, 48]. 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 $\text{Sp}(2n)$ 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 $\text{codim}(H) = 2$, as well as prove theorem III.2.6 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 [I.4.3](#), 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.

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 18th century [\[50\]](#) and has consistently appeared in various physical and mathematical contexts through to the modern day. In the first half of

the 19th 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 [39], 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* [85], now a standard introductory example for the concept of virtual work⁷, as well as a Hellenic treatise on optics⁸ [50]. 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 [50] 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 20th century the establishment of Morse theory [70] 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 spaces [70], 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

⁷The 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 [8].

⁸Heron’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 [15], 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 [43], 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 [74], 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 [76] 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 functional⁹. The main obstacle Rabinowitz and his

⁹This description is a paraphrasing of Helmut Hofer during his 60th birthday conference. In the video he

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 [14], 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 $n \geq 2$, 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 $\gamma : [0, 1] \rightarrow \text{Sp}(2n)$ 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 [23, 24] which built on his earlier results [21, 22] 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 homology¹⁰. This

is speaking of Rabinowitz’s 1978 paper and it’s influence on his decision to specialize in symplectic/contact geometry [72].

¹⁰The descriptor ‘Floer homology’ is rather vague, even if one ignores analogues developed later. Floer’s work

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 [20] and continued on by other authors in the papers [57, 86, 58], 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 $\gamma : [0, 1] \rightarrow L$ may be written as a signed intersection count of the associated path $\gamma_T : [0, 1] \rightarrow \Lambda_n$ (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 [67]. The manner in which the Maslov index was formed would foreshadow the methods later used in [82] 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), each of which

exhibited homology theories associated to Lagrangian intersections [23] as well as non-degenerate symplectic endomorphisms [21, 24] (and this is entirely omitting his work on three-manifolds).

¹¹The idea of using a signed intersection count and some co-oriented hypersurface in $\mathrm{Sp}(2n)$ (as in [82]) to define an index for linear-symplectic paths was nearly contemporaneous, as can be seen in figure 1.1 which depicts the analogous ‘Maslov cycle’ in $\mathrm{Sp}(2)$, consisting of those symplectomorphism with an eigenvalue equal to 1.

inducing an isomorphism of the relevant fundamental group¹² 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 [82] 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 $\text{Sp}(2n)$).

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 [14] 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 [79] 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 [14] does not mark the first appearance of the mean and Conley-Zehnder indices. The mean index is foreshadowed as far back as 1885 in Poincaré's rotation number [75] later to be associated with Morse theory by Hedlund [41] 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 [9], 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

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

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 ([\[17, 16, 82, 58\]](#) 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 [\[79\]](#) mentioned above, which appears in Yakubovich and Starzhinskii's text [\[91\]](#). 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 [\[26\]](#). 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 [\[14\]](#) 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 [\[13\]](#). 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 [\[18\]](#) that the image $\text{Symp}(M, \omega) \hookrightarrow \text{Diff}(M)$ 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 [\[34\]](#), 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 [\[68, 46, 47, 83\]](#).

Generated using Mathematica 10 © [90]

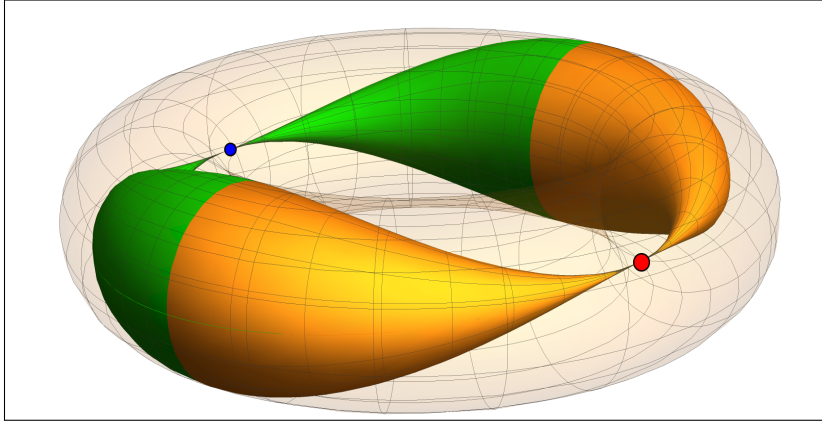


Figure I.1: The above is a depiction of $\mathrm{Sp}(2, \mathbb{R})$ as a solid (open) torus; the elements Id , $-\mathrm{Id}$ are marked by a red/blue point while the surface of symplectomorphisms with eigenvalue $+1$, -1 are drawn in orange/green respectively. The figure is of interest as an analogous figure (not shown here to avoid violating copyright law) appeared in the 1955 paper [26] 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 the procedure given in [82] for defining the Conley-Zehnder index in terms of a signed intersection count (à la the Maslov class by Arnol'd [2]) by several decades).

Two papers of note which utilized Morse theory for Hamiltonian systems (pre-Floer homology) were published by Ekeland [17], and Ekeland and Hofer [16], 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 $\{1, 2, \dots, 2n\}$ -valued nullity index is defined, foreshadowing the role it would play in the paper [58] 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 [19, 20, 21, 22, 23, 24, 25] 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 [20]) and undoubtedly by Gromov's introduction of pseudoholomorphic curves (see the paper [19] 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 [80, 10], 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 [14] 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 [61, 57] 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 [55, 56]). In the first, published with Zehnder, they defined the index for non-degenerate linear Hamiltonian systems in two dimensions [61] filling in the gap left in [14], while the second paper finally extended the index to degenerate fundamental solutions of linear Hamiltonian systems (provided they have symmetric and continuous coefficients) [57]. Remarkably, a nearly identical result in [86] was published that same year by Viterbo. In 1992, Salamon and Zehnder published their hallmark paper [82] 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 ax-

iomatic formulation¹³ of the Conley-Zehnder index under which the index is unique, which coincidentally was also shown the same year by Barge and Ghys in [5].

A paper [78] 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 [91]). Much of the material covered until this point (and much more) would later be compiled and published in 1999 as notes [81] 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 [58] 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 [82, 58] 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 [25] followed two decades later by the Conley conjecture (under identical assumption), which was shown to hold in [29, 31] (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

¹³An analogous list of axioms for the mean index was also published by Barge and Ghys in the same paper [5], although our index unfortunately does not satisfy these axioms (see remark I.3.13).

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 [5, 6] for the former and [82, 5] 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 [82, 5] (among many others), wherein the mean index is defined using the same map $\rho : Sp(2n) \rightarrow S^1$ 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 paragraph¹⁴. 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 $\gamma : [0, 1] \rightarrow Sp(2n)$ which satisfies the following equation for all such γ ,

$$\Delta(\gamma) = \lim_{k \rightarrow \infty} \frac{\mu_{cz}(\gamma^k)}{k}. \quad (\text{I.2.1})$$

Remark I.2.1. Here μ_{cz} denotes the Conley-Zehnder index. We refer the reader to lemmas I.3.8, 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 $Sp(2n)$ 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

¹⁴In fact, the etymology of the term ‘mean’ index is a reference to this ‘weighted average’, expressed in equation (I.2.1).

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 [17, 16, 82] holding for whichever ‘admissible’ paths γ each paper is considering. In 1997, Yiming Long established inequality (I.2.2) for all paths [58], constructing a $\mathbb{C} \setminus \{0\}$ family of Maslov-type indices, $(\mu_\omega, \eta_\omega) \in \mathbb{Z} \times \{0, 1, \dots, 2n\}$, 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 [54].

In addition to satisfying inequality (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 $\Delta : G \rightarrow \mathbb{R}$ 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 [5] which established the crucial fact that Δ may be characterized axiomatically as the unique continuous quasimorphism $\Delta :$

$\widetilde{Sp}(2n) \rightarrow \mathbb{R}$ which is both homogeneous and continuous; see the 2008 paper [6] 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 $H_{bc}^2(\widetilde{Sp}(2n), \mathbb{R})$ (as defined in [73]). 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 [11] 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 [28] 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 [30] establishes various conditions (some regarding the mean indices of the periodic points) that a Hamiltonian ‘pseudo-rotation’¹⁵

¹⁵We adopt for the above paragraph the nomenclature that any Hamiltonian with finitely many periodic

of complex projective space must satisfy. The results reached in [12], this time regarding any symplectic manifold admitting a pseudo-rotation, again utilize the 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, 71, 51]. The underlying index theory common to all of these references may be found in [32, 60] where the former focuses on Lusternik-Schirelmann 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,

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

We will distinguish set theoretic from linear subspace inclusion by using ‘ \leq ’ for the latter. A Lagrangian subspace $L \leq V \times \bar{V}$, also referred to as a linear canonical relation, will be said to have source and target V and \bar{V} respectively¹⁶ and we introduce the following notation, used in [84] for linear relations with the exception being the colorful notation $\text{halo}(L)$ (introduced in [63] yet lately supplanted by the more conventional $\text{indet}(L)$ as in [53]).

Definition I.3.1. Recall for a symplectic vector space (V^{2n}, ω) that the Lagrangian Grass-points is called a pseudo-rotation

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

mannian $\text{LagGr}(V, \omega)$ is defined as follows (where L^ω denotes the symplectic orthogonal),

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

where $\text{Gr}_n(V)$ is the standard Grassmannian manifold of n -planes in the vector space V .

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

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

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

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

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

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 $L^0 := \{(v, v) \in V \times V \mid v \in V\}$ denote the diagonal Δ_V which is the identity in the monoid of linear Lagrangian relations [88]; for any $L \in \Lambda_{2n}$, $\Delta_V \circ L = L = L \circ \Delta_V$.

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

$$\begin{aligned} \text{Sp}(V) &\xrightarrow[\text{Gr}]{} \Lambda_{2n} \\ A &\mapsto \{(v, Av) \in V \times \bar{V} \mid v \in V\}. \end{aligned}$$

As shown in [41], 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 $\text{Sp}(2n)$. We partition Λ_{2n} into $n + 1$ pair-wise disjoint sets defined for each $0 \leq k \leq n$ (note that $\Lambda_{2n}^0 = \text{Im}(\text{Gr}) \cong \text{Sp}(V)$),

$$\Lambda_{2n}^k := \{L \in \Lambda_{2n} \mid \dim(\ker(L)) = \dim(\text{halo}(L)) = k\}.$$

As shown in [48], these $n + 1$ sets are in fact smooth submanifolds that form a stratification¹⁷ of Λ_{2n} ,

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

where $\Lambda_{2n}^{\geq r} := \bigcup_{k=r}^n \Lambda_{2n}^k$. Note that since $\Lambda_{2n}^{\geq 1} = \Lambda_{2n} \setminus \text{Im}(\text{Gr})$ 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,

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

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 [62] to

¹⁷The 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.

completely characterize the conjugacy classes of (co)isotropic pairs (see definition II.1.1 for more details on isotropic pairs).

As first determined in [45] following [66], 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 spaces¹⁸ X, Y and Z and a pair of linear canonical 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 $\text{dom}(L') \oplus \text{ran}(L) = Y$ or equivalently, $(L' \times L) \cap (\{0_X\} \times \Delta_Y \times \{0_Z\}) = \{0\}$.

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 $\ker(L') \cap \text{halo}(L) = \{0\}$ whereas we require (under a more restrictive assumption) that $\ker(L') \cap \text{ran}(L) = \{0\}$, 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; $* \circ * : \text{LagGr}(X \times \bar{Y}) \times \text{LagGr}(Y \times \bar{Z}) \rightarrow \text{LagGr}(X \times \bar{Z})$. 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 case 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 [35] 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

¹⁸The 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 [7]

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 [87] and later refined in [53]). 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 [28, 11, 30, 12] 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. [77] For a connected Lie group G with $\pi_1(G) \cong \mathbb{Z}$ we call any smooth $\rho : G \rightarrow S^1$ 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. $\rho(\text{Id}) = 1$.

As shown in [77] our map ρ is a normalized circle map on $\text{Sp}(2n)$.

Lemma I.3.8. [40] Given a pointed topological space (X, x_0) for which a universal cover \tilde{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 \tilde{X} .

As is becoming clear, the vast majority of applications for the mean index rely on the fact that $\pi_1(\text{Sp}(2n)) \cong \mathbb{Z} \cong \pi_1(S^1)$ and that ρ is a circle map. We proceed with a characterization from Rawlins (based on Milnor’s [69]) of the universal cover for such Lie groups.

Lemma I.3.9. [77] 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

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

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

$$\begin{array}{ccc} \tilde{G} & \overset{\Delta}{\dashrightarrow} & \mathbb{R} \\ & \searrow \tilde{\rho} & \downarrow e^{it} \\ & & S^1 \end{array}$$

Figure I.2: The above diagram commutes.

We leave the explicit reconciliation of definition I.3.10 and definition I.3.14 below to [77], but it's not a stretch to see from lemma I.3.9 that any path $\gamma : I \rightarrow G$ with $\gamma(0) = \text{Id}$, when lifted: $\tilde{\gamma} : I \rightarrow \tilde{G}$, 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. [77] 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 $\eta : G \times G \rightarrow \mathbb{R}$ for

which $\eta(\text{Id}, \text{Id}) = 0$ 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 $\tilde{\rho} : \widetilde{\mathcal{L}_{2n}} \rightarrow S^1$ 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 $\rho_* : \pi_1(\mathcal{L}_{2n}) \rightarrow \mathbb{Z}$ 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 $\gamma: [0, 1] \rightarrow \mathrm{Sp}(V)$ there exists a unique, continuous lift $\theta: [0, 1] \rightarrow \mathbb{R}$ 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 $\gamma: [0, 1] \rightarrow \mathrm{Im}(\mathrm{Gr}) \cong \mathrm{Sp}(V)$.

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). 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 below.

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 [I.4.3](#) 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 $\gamma: I \rightarrow \mathcal{L}_{2n}$, when composed with $\hat{\rho}$, lifts to a unique continuous $\hat{\theta}: I \rightarrow \mathbb{R}$ 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 $j: \tilde{\mathcal{L}}_{2n} \rightarrow \mathcal{L}_{2n}$ and lift this to obtain $\hat{\Delta}: \tilde{\mathcal{L}}_{2n} \rightarrow \mathbb{R}$ 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 I.4.3. There exists a unique continuous map $\hat{\rho}: \mathcal{L}_{2n} \rightarrow S^1$ such that for all $\phi \in \text{Sp}(V)$ we have $\hat{\rho}(\text{Gr}(\phi)) = \rho^2(\phi)$.

Already in dimension two the purpose of squaring ρ is clear; there exist sequences A_i^\pm for which $\text{Tr}(A_i^\pm) > 2$ for all $i \in \mathbb{N}$ (equivalent to hyperbolicity) and $\text{Gr}(A_i^\pm) \rightarrow L \notin H$ 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 $\Lambda_{2n} \setminus \text{Im}(\text{Gr})$, 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.4.4. When approaching H via the elliptic transformations significant discontinuities arise even in low dimensions, in [\[33\]](#) an example is given for some fixed $L \in H$ in which an $S^1 \setminus \{\pm 1\}$ family of sequences of symplectic maps $\{A_i^\theta\}_{i=1}^\infty \in \text{Sp}(2)$ are constructed such that each has $\lim_{i \rightarrow \infty} \text{Gr}(A_i^\theta) = L$ yet $\rho(A_i^\theta) = \theta$ for all $i \in \mathbb{N}$.

The following lemma is a partial result towards theorem [I.4.6](#).

Lemma IV.2.1. The extended circle map $\hat{\rho}$ is homogeneous on \mathcal{L}_{2n} ; given any $L \in \mathcal{L}_{2n}$ then $\hat{\rho}(L^l) = 2l \cdot \hat{\rho}(L) \in \mathbb{R}/\mathbb{Z}$ 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.4.5 (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.4.5) to the stratified space $(\mathcal{L}_{2n}^k)_{k=1}^n$.

we show in theorem II.4.15 that $\mathcal{P}_{reg}(\mathcal{L}_{2n})$ 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.4.6. 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, VI.1.5.

Corollary I.4.7. 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.4.8. 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.4.9. 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 $\phi_{\gamma(t)} \in \text{Sp}(\text{dom}(\gamma(t)) \cap \text{ran}(\gamma(t)))$ (see theorem III.2.6 below) 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.4.5 above.

I.4.2 Outline of Proofs

We first note that part VI contains the proofs regarding the stratum-regular paths including corollary I.4.7 as well as the quasimorphism-type bound stated in theorem I.4.8. In addition to these, the part concludes with proofs for both the technical lemma VI.1.5 and theorem I.4.6, together being critical in proving the preceding results. In regards to the proof of our central theorem I.4.3 (that is, continuously extending ρ^2 as described above) it will come as a fairly straightforward consequence of the following three theorems.

Theorem III.2.6. For a given $L \in \mathcal{L}_{2n}$ there exists a unique symplectic decomposition of

$V = V_s \oplus V_g$ and $\phi \in \text{Sp}(V_g)$ such that

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

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

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

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

(i.e. $L_{\text{dom}}, L_{\text{ran}} \in \Lambda_n$ 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 $\{A_i\}_{i=1}^\infty \subset \text{Sp}(V)$ where each A_i has distinct eigenvalues and for which

$$\text{Gr}(A_i) \xrightarrow{i \rightarrow \infty} L \in \mathcal{L}_{2n},$$

where the graph part¹⁹ of L , $\phi_L \in \text{Sp}(V_g)$ 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 $A_i = \alpha_i \oplus \beta_i \in \text{Sp}(E_s^i) \times \text{Sp}(E_g^i)$ such that $\text{Gr}(\alpha_i) \rightarrow \ker(L) \times \{0\} \oplus \{0\} \oplus \text{halo}(L)$.

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

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

uniquely determined by L such that each $\beta_i: E_g^i \rightarrow E_g^i$ is conjugate via I_i to some $\phi_i \in \text{Sp}(V_g)$ for all $i \geq N$ with $\phi_i \rightarrow \phi$. 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.

¹⁹See remark III.2.3.

Remark I.4.10. Refer to [36] 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 I.4.3 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 ,

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

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,

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

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

$$\begin{array}{ccc} \text{Sp}(2n - 2k) \cong \Lambda_{2n-2k}^0 & \hookrightarrow & \Lambda_{2n}^k \\ & & \downarrow \text{Pr}_I \\ & & \mathcal{I}_k \end{array}$$

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

Remark II.1.5. We have implicitly used the notational conventions $\mathrm{Sp}(0) \cong \Lambda_n^0 := \{0\}$ 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,

$$\{0\} \hookrightarrow \Lambda_{2n}^n \xrightarrow{\mathrm{Pr}_I} \Lambda_n \times \Lambda_n,$$

as each fiber is trivial. One may verify that $\Lambda_{2n}^n \cong \Lambda_n \times \Lambda_n$ 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,

$$\mathrm{Sp}(V) \xrightarrow{\mathrm{Gr}} \Lambda_{2n}^0 \xrightarrow{\mathrm{Pr}_I} I_0(V) = \{0\}.$$

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

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

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

denote the space of admissible isotropic pairs of dimension k . We see that $\mathcal{L}_{2n}^k \rightarrow \mathcal{I}_k$ 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 $\mathcal{L}_{2n}^0 = \Lambda_{2n}^0$ and $\mathcal{L}_{2n}^n = \Lambda_n \times \Lambda_n \setminus \hat{\Sigma}_n$ where,

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

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

One may observe that lemma II.1.4 implies that the set $\left(\mathcal{L}_{2n}^{\geq k}\right)_{k=0}^n$ 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.4.2 below).

Proposition II.1.6. The codimension of each \mathcal{L}_{2n}^k 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

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

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

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

□

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

Yet, since $K'_i \rightarrow K' = L_{\text{ran}} \times \{0\} \oplus \{0\} \times L_{\text{dom}}$ as $i \rightarrow \infty$ we may form the compositions of the limits,

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

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

so that we indeed have a failure of continuity,

$$\lim_{i \rightarrow \infty} K'_i \circ \lim_{i \rightarrow \infty} K_i = L_{\text{dom}} \times \{0\} \oplus \{0\} \times L_{\text{dom}} \neq \Delta_V = \lim_{i \rightarrow \infty} (K'_i \circ K_i).$$

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

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

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 $\text{halo}(L^l) = \text{halo}(L)$ for all $l \geq 1$ making each iteration map $(*)^l : \mathcal{L}_{2n}^k \rightarrow \mathcal{L}_{2n}^k$ 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 $\ker(L) \leq \text{dom}(L^i)$ 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, \dots, d_{2n-2k}) such that $\langle d_j \rangle_{j=1}^{2n-2k} \oplus \ker(L) = \text{dom}(L)$ and thus each d_j is associated (non-uniquely) via L to some $r_j \in \text{ran}(L)$. The r_j are also linearly independent as 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 $\langle d_j \rangle_{j=1}^{2n-2k} \pitchfork \ker(L)$.

Again, since $L \not\subseteq H$, we observe that $\langle r_j \rangle_{j=1}^{2n-2k} \oplus \text{halo}(L) = \text{ran}(L)$ and we see that $\text{dom}(L) \cap \text{ran}(L) = \text{dom}(L) \cap \langle r_j \rangle_{j=1}^{2n-2k}$. Since $\dim(\text{dom}(L) \cap \text{ran}(L)) \geq 2n - 2k$ we see for dimensional reasons that $\langle r_j \rangle_{j=1}^{2n-2k} \leq \text{dom}(L)$ and therefore

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

An identical argument shows that $\text{ran}(L^2) = \text{ran}(L)$ 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 and homogeneous for each $i \in \mathbb{N}$, and that it sends differentiable paths to piece-wise differentiable paths in theorem [I.4.6](#). The potential of extending homogeneity of the extended mean index over paths to negative numbers becomes a bit more interesting as the natural inverse in the category of linear relations would reverse 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.4.7](#) and the statements which build up to it including definition [I.3.2](#) and lemma [IV.2.1](#)) but by far the most natural choice would be the ‘reverse’ of a linear relation L ,

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

Then by letting $L^{-l} := \bar{L}^l$ for $l \geq 1$ one may define iteration for all integers, albeit with the undesirable property that $\ker(L^l) = \text{halo}(L)$ and $\text{halo}(L^l) = \ker(L)$ 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, $(\mathcal{L}_{2n}^k)_{(B_1, B_2)} \cong \text{Sp}(B_1^\omega \cap B_2^\omega)$ (as detailed in remark [III.2.2](#)), and these subgroups $\text{Sp}(B_1^\omega \cap B_2^\omega) \subset \text{Sp}(V)$ 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 $\phi, \tau \in \text{Sp}(B_1^\omega \cap B_2^\omega)$ would satisfy $\text{Gr}(\phi\tau) = \text{Gr}(\phi) \circ \text{Gr}(\tau)$ and $\text{Gr}(\phi^i) \in (\mathcal{L}_{2n}^k)_{(B_1, B_2)}$ 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 and 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 $A \in \text{Sp}(V)$ for which $(A(B_1), A(B_2)) = (B'_1, B'_2)$. The equivalence classes coincide with the orbits of the group action $\text{Sp}(V) \circlearrowleft \mathcal{I}_k$ 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 $(B_1, B_2) = (\ker(L), \text{halo}(L))$, recalling that $\dim(\ker(L)) = \dim(\text{halo}(L))$, we see that only considering those pairs of equal dimension is justified.

Theorem II.3.2. [63] The four integers (r, κ, k, n) form a complete set of invariants for isotropic pairs subject to the relations $0 \leq r \leq \kappa \leq k \leq n$ and $0 \leq \kappa - r \leq n - k$ where,

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

Now denoting $\Lambda_{2n}^k := \{L \in \Lambda_{2n} \mid \dim(\ker(L)) = k\}$ for any $0 \leq k \leq n$ and noting $H = \{L \in \Lambda_{2n} \mid \kappa(L) \geq 1\}$ 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 $k \leq n$ (and therefore on all of Λ_{2n}) where $L \sim L'$ if and only if $(\ker(L), \text{halo}(L)) \sim (\ker(L'), \text{halo}(L'))$. A detail to note is that on $\Lambda_{2n}^0 \cong \text{Sp}(V)$ 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. $L \sim_{Gr} L'$ if and only if $(v, w) \in L \Leftrightarrow (Av, Aw) \in L'$. The equivalence classes coincide with the orbits of the group action $\mathrm{Sp}(V) \curvearrowright \Lambda_{2n}$ 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_{2n}^0 into the usual conjugacy classes of $\mathrm{Sp}(V)$ while conversely, \sim and \sim_{Gr} are identical on Λ_{2n}^n . 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 [62].

II.4 Stratum-Regular Paths

Remark II.4.1. We should note that the material below is elementary, and that any proofs explicitly shown are present only for the sake of completeness. One particular reason for this (beyond the obvious fact that just about any textbook addressing transversality and stratifications of manifolds is likely to cover and generalize every claim herein) is that our domain is assumed to be the unit interval, thus simplifying the proceeding arguments in two ways. The first is a consequence of our domain's dimension; it is the minimal dimensional domain one might consider when inspecting non-trivial C^r spaces,²⁰ and as a consequence many of the subtleties encountered when proving analogous transversality arguments fail to manifest. The second is compactness of the domain, which make the usually vital question of

²⁰By trivial we refer to the differentiable function spaces $C^r(K, N)$ where N is a smooth manifold (of dimension greater than zero) and K is some subset of \mathbb{N} . In particular, these 'function spaces' are little more than the set of countable sequences in a manifold N equipped with a certain topology.

which topology the space $C^r(M, N)$ is equipped with unnecessary as the two most commonly used (these being the compact-open/weak topology and Whitney/strong topology) become equivalent when the domain M is compact [3].

We may now proceed in giving the particular definition of a stratified space we will be using throughout.

Definition II.4.2. Given a smooth manifold N , we call any closed subset $C \subset N$ a stratified space if there exists a finite sequence 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 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 $(C_i)_{i=1}^n$.

Remark II.4.3. We note that the above definition is far from canonical, as there exist a veritable zoo²¹ of definitions for stratified spaces in the literature (see [89] for a thorough account from the perspective of algebraic K-theory and cobordisms).

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); $\overline{\Lambda_{2n}^k} = \Lambda_{2n}^{\geq k}$.

²¹More general definitions may specify the indexing set to be an infinite partial order provided the collection of strata is locally finite, or even define the stratification in terms of a continuous map from the given space to a topologized partial order, as in appendix A.5 of Lurie’s Higher Algebra [64] (cited in [4] as the earliest example of such a definition known to the authors).

Remark II.4.4. In light of the above remark, we note here that our notion of a stratified space generally corresponds to that of a topological stratification. 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 [62]. Then, provided our indexing set $\{(r, \kappa, k)\}_{0 \leq r \leq \kappa \leq k \leq n}$ may be equipped with some partial order such that the intended stratum $\Lambda_{2n}^{r, \kappa, k}$ satisfy the appropriate conditions, this would yield a much finer stratification of Λ_{2n} than that given here.

We will consider these invariants again in the context of proving that the codimension of H is two, see remark III.1.2.

Definition II.4.5. [3] For any smooth manifolds²² M, N and smooth submanifold $C \subseteq N$, we call a C^1 map $f : M \rightarrow N$ transversal to C (notated $f \pitchfork C$) if for all $p \in f^{-1}(C)$,

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

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

Remark II.4.6. Given $f \in C^1(M, N)$ and a smooth submanifold $C \subseteq N$, observe that $f \pitchfork C$ implies that $\text{Im}(f) \pitchfork C$ as submanifolds, whereas the converse does not hold. In particular when $\text{codim}(C) > \dim(M)$ 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 $\text{Im}(f) \cap C = \emptyset$, i.e. only the case of vacuous transversality is possible. When $\text{codim}(C) \leq \dim(M)$ the latter is generally a proper subset of the former.

²²We 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.4.7. [3] Given a smooth manifold M and a pair of transversely intersecting smooth submanifolds $A, B \subseteq M$, then $A \cap B$ is a submanifold of M with $\text{codim}(A \cap B) = \text{codim}(A) + \text{codim}(B)$.

Proposition II.4.8. 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 proposition II.4.7 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: $\text{codim}(C) + \text{codim}(D) = \dim(M)$ to see that $\text{codim}(C \cap D) = \dim(M)$ or equivalently, $\dim(C \cap D) = 0$. 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. \square

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.

Lemma II.4.9. [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.4.10. The set $\mathcal{L}_{2n} \subset \Lambda_{2n}$ is open and dense.

As proposition II.4.10 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 $\Lambda_{2n-2k}^0 \hookrightarrow H_k \rightarrow \hat{H}_k$ where $\hat{H}_k = Pr_I(H_k) \subset I_k \times I_k$.
- 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} (using our choice of definition at least). 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. $\overline{\mathcal{L}_{2n}^k} = \mathcal{L}_{2n}^{\geq k}$.

Proposition II.4.11. 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 $\text{codim}(K) > \dim(M)$, the set $C^1(M, U)$ is open and dense in $C^1(M, N)$.

Proof. We observe that the set $\{f \in C^1(M, N) \mid f \pitchfork K\}$ is open and dense in $C^1(M, N)$ by lemma II.4.9, as K is closed. Without the codimension bound on K we have the (generally proper) inclusion;

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

but as mentioned in remark II.4.6, we see when $\text{codim}(K) > \dim(M)$ that each $f \in C^1(M, N)$ and $p \in f^{-1}(K)$ yields

$$\dim(df_p(T_p M) + T_{f(p)}K) < \dim(T_{f(p)}N),$$

so that $f \not\pitchfork K$. Consequently, if $\text{codim}(K) > \dim(M)$ 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;

$$\{f \in C^1(M, N) \mid \text{Im}(f) \cap K = \emptyset\} = \{f \in C^1(M, N) \mid f \pitchfork K\}.$$

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

Corollary II.4.12. Consider the following two consequences of proposition II.4.11:

- $C^1([0, 1], \mathcal{L}_{2n})$ 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.4.10, 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 $\text{codim}(H) = \text{codim}(H_1) = 2$ (as shown in theorem I.4.1), an application of proposition II.4.11 yields the results.

- Since $\text{codim}(\mathcal{L}_{2n}^k) = k^2$ (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.4.11 does not apply and $C^1([0, 1], \mathcal{L}_{2n}^0)$ is evidently not open (see remark II.4.14 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,

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

Then proposition II.4.11 means 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.4.13. The inclusion chains implied by the two statements of corollary II.4.12 above show that $C^1([0, 1], \mathcal{L}_{2n}^{\leq k})$ are dense and open subsets of $C^1([0, 1], \mathcal{L}_{2n})$ for all $1 \leq k \leq n$.

Remark II.4.14. Very simple counter-examples exist to $C^1([0, 1], \mathcal{L}_{2n}^0)$ being open in the set $C^1([0, 1], \mathcal{L}_{2n}^{\leq 1})$; consider the paths

$$\gamma^\pm(t) = \text{Diag}_{2 \times 2}(\pm(\log(1-t))^{-1}, \log(1-t)) \oplus \text{Id}_{2n-2} : [0, 1] \rightarrow \text{Sp}(\mathbb{R}^{2n}) \cong \mathcal{L}_{2n}^0$$

with respect to a Darboux basis $(a_i, b_i)_{i=1}^n$. Next define the following paths of canonical relations: $\hat{\gamma}^\pm : [0, 1] \rightarrow \mathcal{L}_{2n}$ for which $\hat{\gamma}^\pm(t) = \text{Gr}(\gamma^\pm(t))$ and $\hat{\gamma}^\pm(1) = L = \langle (a_1, 0), (0, b_1) \rangle \oplus \text{Gr}(\text{Id}_{2n-2}(V_g))$ where $V_g = \langle (a_i, 0), (0, b_i) \rangle_{i=2}^n$. Then $\beta = \gamma^- * \gamma^+ : I \rightarrow \mathcal{L}_{2n}^{\leq 1}$ with $\gamma^+(0) \in \text{Sp}^+(V)$ and $\gamma^-(0) \in \text{Sp}^-(V)$. This exhibits a path in the interior of $C^1([0, 1], \mathcal{L}_{2n}^{\leq 1}) \setminus C^1([0, 1], \mathcal{L}_{2n}^0)$.

Theorem II.4.15. 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 $D_k := \left\{ f \in C^1([0, 1], \mathcal{L}_{2n}^{\leq k}) \mid f \pitchfork \mathcal{L}_{2n}^k \right\}$ of C^1 paths in $\mathcal{L}_{2n}^{\leq k}$ which are transverse to the k^{th} stratum \mathcal{L}_{2n}^k . Since \mathcal{L}_{2n}^k is closed in $\mathcal{L}_{2n}^{\leq k}$ we see that $D_k \subset C^1([0, 1], \mathcal{L}_{2n}^{\leq k}) \subset C^1([0, 1], \mathcal{L}_{2n})$ is an inclusion of open and dense sets for all $1 \leq k \leq n$ by lemma II.4.9 (alternatively proposition II.4.11 would work as well).

Letting $D = \bigcap_{k=2}^n D_k$ be the set of paths transverse to each \mathcal{L}_{2n}^k 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 $D = C^1([0, 1], \mathcal{L}_{2n}^{\leq 1})$ to reach the same conclusion via corollary II.4.12. 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

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

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,

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

Then since \mathcal{L}_{2n}^1 is closed in $\mathcal{L}_{2n}^{\leq 1}$, lemma II.4.9 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.4.12. As mentioned above in remark II.4.13, this follows from the heredity of both properties with respect to the subspace topology. \square

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.4.16. Given M a smooth manifold, $A, B \subset M$ submanifolds for which $\dim(A) + \dim(B) = \dim(M)$ 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 $\{p_i\}_{i=1}^{\infty} \subset A \cap B$ so that as A is compact this sequence admits a convergent subsequence $\{p_{i_j}\}_{j=1}^{\infty}$ for which $p_{i_j} \rightarrow p$ as $j \rightarrow \infty$ 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 $p_{i_j} \rightarrow p \in B$, 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.4.8. \square

Remark II.4.17. Observe that a subset $B \subset M$ is closed if and only if $\bar{B} \setminus B = \emptyset$. Also note that $\bar{B} \setminus B$ is in general a proper subset of the boundary $\partial B = \bar{B} \setminus B^\circ$, failing to include those boundary points already contained in B . Regardless, 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.

Proposition II.4.18. For any stratum-regular path γ , the intersection $\text{Im}(\gamma) \cap (\Lambda_{2n} \setminus \text{Im}(\text{Gr}))$ 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 $\gamma : I \rightarrow \mathcal{L}_{2n}^{\leq 1}$. In this case, observe that the pair $A := \text{Im}(\gamma)$ is compact and $B := \Lambda_{2n}^1$ is closed in $M := \Lambda_{2n}^{\leq 1}$ with complementary dimensions. Thus lemma II.4.16 implies that $|\text{Im}(\gamma) \cap \Lambda_{2n}^1| < \infty$ for any $\gamma \in D_1$, and thus any $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$. \square

Part III

The Set H of Exceptional Lagrangians

III.1 The Codimension of H

Before proceeding, we restate some central concepts first covered in section II.1, including lemma II.1.4 in which it was shown that each of the strata Λ_{2n}^k may be identified with a smooth fiber bundle over $I_k(V) \times I_k(V)$ with fiber diffeomorphic to $\text{Sp}(2n - 2k)$ and projection map given in definition II.1.2. Additionally recall proposition II.4.10 wherein the above stratification was shown to induce a stratification of the closed set $H = \bigcup_{k=1}^n H_k$, where each $H_k \subseteq \Lambda_{2n}^k$ comes equipped with the projection map $Pr_I : H_k \rightarrow I_k(V) \times I_k(V)$ inherited from Λ_{2n}^k . As it turns out, the map Pr_I is the key to proving the following theorem while avoiding technical issues that the set H may present²³

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

Proof. Recalling that $\text{codim}(\Lambda_{2n}^k) = k^2$ (shown in proposition II.1.6) we see from $H \cap \Lambda_{2n}^0 = \emptyset$ that the trivial bound of $1 \leq \text{codim}(H)$ holds. To expedite the procedure we consider only the highest dimensional stratum $H_1 \subset \Lambda_{2n}^1$ and furthermore consider the elements up to equivalence under the relation \sim given in definition II.3.1, which we denote $[L] \in \Lambda_{2n}^1 / \sim$. As $[L]$ is induced entirely by the associated isotropic pair class $[Pr_I(L)] = [(\ker(L), \text{halo}(L))] \in (I_1(V) \times I_1(V)) / \sim$, we may recall equation (II.3.1) to list the three equivalence classes contained in $I_1(V) \times I_1(V)$ (and therefore three in Λ_{2n}^1); $(r, \kappa, k) = (0, 0, 1), (0, 1, 1)$ and $(1, 1, 1)$. We see the definition of H precludes any non-empty intersection with the class $(\kappa, r, k) = (0, 0, 1)$ (observe that $L \in H$ if and only if $\kappa_L = 0$, so the class associated to the invariant tuple $(0, 0, k)$ will correspond to \mathcal{L}_{2n}^k for all $1 \leq k \leq n$) so it follows that H intersects the two classes remaining with $\kappa = k = 1$; both $\dim(\ker(L) \cap \text{halo}(L)) = r = 0$ and $r = 1$.

In the above discussion, we have made the implicit assumption that $n \geq 2$, as certain

²³As H is not a manifold but a closed stratified set we proceed, as in [89], by letting $\text{codim}(H) := \min_{1 \leq i \leq n} \text{codim}(H_i)$.

issues arise²⁴ when $n = 1$ as a consequence of the bound $\kappa - r \leq n - k$ given by Lorand in [62]. We verify that for $n \geq 2$, both tuples $(r, \kappa, k) = (0, 1, 1), (1, 1, 1)$ are admissible (i.e. they satisfy the inequalities above). When $n = 2$ we see that the invariant inequality $\kappa - r \leq n - k = 1$ allows both $r = \kappa = 1$ and $r = 0, \kappa = 1$ so that the equivalence classes below are indeed non-empty, and it is obvious this remains true for all $n \geq 2$. For a lengthier discussion on the combinatorics that are involved in these invariants, see remark III.1.2 following the proof.

Case I: $r=0$

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

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

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 $B_1 = \langle v \rangle$, after which the two conditions imply that $B_2 = \langle w \rangle \neq B_1$ (since $r = 0$) is restricted to the ‘line-punctured’ subspace $B_1^\omega \setminus B_1$ (this set is not a quotient, it is the linear subspace B_1^ω missing the one dimensional subspace B_1). This subset descends under the quotient map $\mathbb{R}^{2n} \rightarrow \mathbb{R}\mathbb{P}^{2n-1}$ to a punctured (in the usual sense now) projective hyperplane $[B_1^\omega] \setminus \{[B_1]\} \subset \mathbb{R}\mathbb{P}^{2n-1}$ so that there are $2n - 2$ dimensions available when choosing B_2 . This yields $\dim([\ker(L), \text{halo}(L)]) = 4n - 3$ and thus,

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

We see that the codimension at any L belonging to the above equivalence class is indeed two, although we still need verify that the remaining invariant class $(r, \kappa, k) = (1, 1, 1)$ has codimension at least two (this is admittedly redundant when $n \neq 1$, as the punctured point in the aforementioned projective hyperplane corresponded to this class, and therefore it cannot exceed the dimension found above in the $r = 0$ case).

²⁴Fortunately it is just the proof as written that runs into issues when $n = 1$; the theorem’s claim remains true, see remark III.1.1 below for a description of the proof when $n = 1$.

Case II: r=1

Let $L \in [L]_1 = \{L \in \Lambda_{2n}^1 \mid \kappa(L) = r(L) = 1\}$, the last equivalence class in Λ_{2n}^1 by theorem II.3.2. We first consider the associated class $[(\ker(L), \text{halo}(L))]$ and write $\text{dom}(L) = \text{halo}(L) = \langle v \rangle$ for any $v \in V$. Then since v is arbitrary and all one dimensional subspaces are isotropic we see that $[(\langle v \rangle, \langle v \rangle)] = \Delta_{I_1(V)} \subset I_1(V) \times I_1(V) \cong \mathbb{RP}^{2n-1} \times \mathbb{RP}^{2n-1}$ implying that $\dim([\langle v \rangle, \langle v \rangle]) = 2n - 1$. Consequently we see 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 $\text{codim}([L]_1) = 2n$ in Λ_{2n} .

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

Remark III.1.1. As a quick proof of the veracity of the concluding statement above for $n = 1$, we observe that $\text{codim}([L]_0)$ is undefined as $[L]_0 = \emptyset$. As described above, this is due to the fact that the only non-zero stratum Λ_2^1 belongs to the class $(r, \kappa, k) = (1, 1, 1)$ since $\kappa - r \leq (n - k) = 0$ implies that the three classes above reduce to the two cases $(0, 0, 1)$ and $(1, 1, 1)$; the first corresponding to the set \mathcal{L}_2^1 while the second is equal to $H_1 = H$. these pathologies are easily avoid though, as in the $n = 1$ case we need only accept that $[L]_0$ is empty, skip that step and see from the above computations that the remaining class $[L]_1$ has $\text{codim}([L]_1) = 2n = 2$, thereby formally extending the proof to the $n = 1$ case.

□

Remark III.1.2. This seemingly singular behavior at $n = 1$ is actually a shadow of the combinatorial complexity in the collection of all tuples (r, κ, k, n) which satisfy **(a)**: $0 \leq r \leq$

$\kappa \leq k \leq n$ and **(b)**: $0 \leq \kappa - r \leq n - k$. Even a process as straightforward as computing the number of equivalence classes for a given Λ_{2n} turns out to be a rather nuanced procedure. For example, when $k \leq \frac{n}{2}$ it is inequality **(a)** that entirely governs the range of r and κ ; observe that in that case we know that $k \leq n - k$ so that since inequality **(a)** implies that $0 \leq \kappa - r \leq k$, we see inequality **(b)** has no effect on the choice of admissible pairs (r, κ) . Consequently, counting the admissible tuples for $k \leq \frac{n}{2}$ amounts to the familiar counting procedure for ordered tuples with equality;

$$|(r, \kappa, k)|_{0 \leq r \leq \kappa \leq k \leq \lfloor \frac{n}{2} \rfloor} = \sum_{k=0}^{\lfloor \frac{n}{2} \rfloor} \sum_{\kappa=0}^k \sum_{r=0}^{\kappa}.$$

Yet when k passes above $\frac{n}{2}$, the freedom of r to vary independently in $\{0, 1, \dots, \kappa\}$ is no longer growing with k but begins shrinking; $\kappa - r \leq n - k < k$. To bring the $n = 1$ case into the current context, we see for all Λ_{2n} that when $k = n$, inequality **(b)** becomes $\kappa - r = 0$ and so $(i, i, n)_{i=0}^n$ exhausts all equivalence classes belonging to the given stratum Λ_{2n}^n . In the $n = 1$ case, it just happens that this top stratum is the only stratum (other than $\Lambda_2^0 \cong \text{Sp}(2)$).

III.2 Admissible Linear Canonical Relations

Even though the following proposition has already been taken as an implicit fact (being readily available in the literature, e.g. [41],[48]), we prove it here for completeness.

Proposition III.2.1. The fiber over any $(B_1, B_2) \in I_k \times I_k$ is diffeomorphic to the symplectic group,

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

Proof. The first diffeomorphism was shown in [7] while the second follows after taking the $4(n - k)$ dimension symplectic quotient,

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

and using the identification $\Lambda_{2n-2k}^0 \cong \text{Sp}(2n-2k)$ with respect to the associated Lagrangian Grassmannian. See theorem III.2.6 below for more details. \square

Remark III.2.2. In general the right-most diffeomorphism,

$$(\Lambda_{2n}^k)_{(B_1, B_2)} \cong \text{Sp}(B_1^\omega/B_1, B_2^\omega/B_2) \cong \text{Sp}(2n-2k),$$

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.2.6 that for any $L \notin H$ the two quotient sets are particularly well behaved, each possessing a canonical identification with the subspace $B_1^\omega \cap B_2^\omega$ via their co-isotropic reduction maps. In particular this allows us to give a unique diffeomorphism $(\Lambda_{2n}^k)_{(\ker(L), \text{halo}(L))} \cong \text{Sp}(\text{dom}(L) \cap \text{ran}(L))$ for any $(\ker(L), \text{halo}(L)) \notin \text{Pr}_I(H)$, equipping each fiber with a standard diffeomorphism $(\Lambda_{2n}^k)_{(\ker(L), \text{halo}(L))} \cong \text{Sp}(\text{dom}(L) \cap \text{ran}(L))$. One natural consequence of this fact is that we may use these maps to equip each fiber with a group operation compatible with Lagrangian composition (restricted to the given fiber) which would induce the inverse defined in remark II.2.3 above.

Remark III.2.3. When $L \in \mathcal{L}_{2n}$ we will call both the symplectic map ϕ_L and the Lagrangian subspace $\text{Gr}(\phi_L) \leq L$ the ‘graph’ portion of L when the context is unambiguous. We allow this abuse of notation as the aforementioned fiber diffeomorphism $(\Lambda_{2n}^k)_{(\ker(L), \text{halo}(L))} \cong \text{Sp}(\text{dom}(L) \cap \text{ran}(L))$ is uniquely determined by the base point $(\ker(L), \text{halo}(L))$, which allows us to bypass the quotient construction as seen in proposition III.2.1.

Remark III.2.4. There always exists a pair of symplectic subspaces $V_i \leq L_i$ which are mapped bijectively under the projection maps $\pi_i: L_i \rightarrow L_i^\omega$, for which $V_i \cong L_i/L_i^\omega$ and $\phi: V_1 \rightarrow V_2$ is symplectic, yet the ϕ obtained depends not only on L but which pair of V_i are chosen as well. This leaves us with the bare-bones structure²⁵ of $\text{Sp}(V, W)$, where V, W

²⁵In [84], a set of complete invariants for (arbitrary) linear relations $\text{Gr}_k(V \times W)$ for V, W distinct vector

are two symplectic vector spaces of the same dimension. When $\text{Sp}(V, W)$ is compared to the usual group of symplectic automorphisms $\text{Sp}(V)$ (which exhibits a variety of normal forms and decompositions), the sparsity of the structure (up to linear symplectomorphism) on the former group becomes evident. In particular, any discussion of eigenvalues and eigenvectors (ergo any prospect of defining a circle map ρ) are dependent on some arbitrary isomorphism $V \rightarrow W$ fixed as an ‘identity’. Of course this is little better than not having the aforementioned structure theorems and canonical forms at all, as each eigenvalue/eigenvector/normal form will vary with the choice of ‘identity’.

Now we intend to show that the problem described in remark III.2.4 is avoided provided $L \notin H$, as such an L is shown to induce via $\text{Pr}_I(L) = (\ker(L), \text{halo}(L))$, a unique ω -orthogonal decomposition of V which splits L into the direct sum of its ‘singular’ (kernel and halo) and ‘graph’ components. The end result will be that any $L \in \mathcal{L}_{2n}$ may be identified with a symplectic map $\phi_L \in \text{Sp}(\text{dom}(L) \cap \text{ran}(L))$ in a symplectic group uniquely determined by $\text{Pr}_I(L)$.

Remark III.2.5. In both the following proof of theorem III.2.6 in addition to certain portions of part V, we will use the shorthand $L_1 := \text{dom}(L)$ and $L_2 := \text{ran}(L)$ for any $L \in \Lambda_{2n}$ when the notation becomes overly cumbersome. We restate here for clarity that $L \in \mathcal{L}_{2n}$ if and only if $L_1 \cap L_2^\omega = \{0\}$.

Theorem III.2.6. Given $L \in \mathcal{L}_{2n}$ there exists a unique²⁶ symplectic decomposition $V = V_s \oplus V_g$ where $V_s = \ker(L) \oplus \text{halo}(L)$ and $V_g = \text{dom}(L) \cap \text{ran}(L)$. Consequently each L yields

spaces of identical dimension is shown to amount to just three integers, corresponding to the dimension of the kernel, domain and halo. On the other hand, when the two vector spaces are identified the isomorphism classes in $\text{Gr}_k(V \times V)$ amount to direct sums of four types of normal form relations which uniquely represent the class up to a permutation (much like the Jordan normal form).

²⁶Unique with respect to the base of the fiber in which L lies; i.e. the decomposition is uniquely determined by $\text{Pr}_I(L) = (\ker(L), \text{halo}(L))$.

a unique symplectic map $\phi \in \text{Sp}(V_g)$ for which,

$$L = (\ker(L) \times \{0\} \oplus \{0\} \times \text{halo}(L)) \oplus Gr(\phi) \leq (V_s \times \overline{V}_s) \oplus (V_g \times \overline{V}_g),$$

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

Proof. We begin by giving an explicit construction (largely overlapping the results of proposition III.2.1 above) for the unique symplectic map $\phi_L \in \text{Sp}(\text{dom}(L)/\ker(L) \times \text{ran}(L)/\text{halo}(L))$ associated to each $L \in \Lambda_{2n}$.

Proposition III.2.6.1. If $L \in \Lambda_{2n}$ and $\tilde{\phi}$ denotes the induced linear map,

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

defined such that $v \mapsto [w]$ if and only if there is some $w \in L_2$ for which $(v, w) \in L$. Given the above, we claim the following three statements are true.

1. This map is well defined.
2. The kernel of $\tilde{\phi}$ is L_1^ω .
3. The map $\phi: (L_1/L_1^\omega, \omega_{red}^1) \rightarrow (L_2/L_2^\omega, \omega_{red}^2)$ is a symplectic isomorphism.

Proofs (1)-(3):

1. Given $(v, w), (v, w') \in L$ we see that $[w] = [w'] \Leftrightarrow w - w' \in L_2^\omega$.
2. 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 $\ker(\tilde{\phi}) = L_1^\omega$. Since the L_i are co-isotropic the map

$\phi: L_1/L_1^\omega \rightarrow L_2/L_2^\omega$ is an isomorphism between symplectic vector spaces and for both $i = 1, 2$ we have,

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

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

3. 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 that ϕ is indeed a symplectic map between the two reduced spaces.

□

Note that proposition III.2.6.1 holds regardless of whether L is exceptional or not much like proposition III.2.1 above. Regardless, as detailed in remark III.2.2 we will need to assume that $L \in \mathcal{L}_{2n}$ (as we do in the following proposition) to complete the proof of theorem III.2.6.

Proposition III.2.6.2. Given $L \in \mathcal{L}_{2n}$, then the isotropic pair $\text{Pr}_I(L) = (L_1^\omega, L_2^\omega)$ determines the symplectic subspace $V_g := (L_1 \cap L_2, \omega)$ of V . Additionally, there exists a standard isomorphism $(V_g, \omega) \cong (L_i/L_i^\omega, \omega_{red}^i)$ for $i = 1, 2$ whereby we will prove the explicit diffeomorphism given in proposition III.2.1 above and justify remark III.2.3 in speaking of ‘the’ graph part of some $L \in \mathcal{L}_{2n}$.

Proof. Denoting the co-isotropic reduction map of the domain and range as $\pi_i : L_i \rightarrow L_i/L_i^\omega$ for $i = 1, 2$ we make the following claim.

Claim. Given L as above, the three statements below are true.

1. $\dim(L_1 \cap L_2) = 2n - 2k$.
2. V_g is a symplectic subspace of V .
3. $\ker(\pi_i) = L_i^\omega$ has trivial intersection with V_g for $i = 1, 2$.

Proofs (1)-(3):

1. Recalling that $\kappa(L) = \dim(L_1 \cap L_2^\omega) = 0$, we know that $r(L) = \dim(L_1^\omega \cap L_2^\omega) = 0$ as well. Then since $\dim(L_i) = 2n - k$ for some $k \leq n$, we have $\dim(L_i^\omega) = k$ and therefore,

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

It follows that $\dim(L_1 \cap L_2) = 2n - \dim(L_1^\omega \oplus L_2^\omega) = 2n - 2k$.

2. 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, as $V_g^\omega = V_s$ are complementary symplectic subspaces, then together they span V and hence comprise a symplectic-orthogonal splitting; $V = V_g \overset{\omega}{\oplus} V_s$.

3. We see that $V_g \cap L_i^\omega = \{0\}$ for $i = 1, 2$ is an immediate consequence of the above two intersections.

□

With a bit of detail filled in, we'll see the above three results will be sufficient to prove proposition III.2.6.2. We wrap things up by letting $V_g := L_1 \cap L_2$ and see that $(\pi_i^* \omega_{red})|_{V_g \times V_g} = \omega|_{V_g \times V_g}$ for $i = 1, 2$ and the above isomorphisms $\text{dom}(L)/\ker(L) \cong V_g \cong \text{ran}(L)/\text{halo}(L)$ together imply that for each $L \in \mathcal{L}_{2n}$ with graph part $\phi_L : \text{dom}(L)/\ker(L) \rightarrow \text{ran}(L)/\text{halo}(L)$, there is precisely one $\phi \in \text{Sp}(V_g)$ for which $\text{Gr}(\phi) \leq L$ or equivalently, for which $\pi_2 \circ \phi = \phi_L \circ \pi_1|_{V_g} : V_g \rightarrow \text{ran}(L)/\text{halo}(L)$. Conversely, for any non-exceptional (i.e. $B_1 \pitchfork B_2^\omega$) isotropic pair $(B_1, B_2) \in \mathcal{I}_k$ and symplectic map $\phi \in \text{Sp}(B_1^\omega \cap B_2^\omega)$, there is a unique $L_\phi \in (\mathcal{L}_{2n}^k)_{(B_1, B_2)}$ whose graph part²⁷ ϕ_L satisfies $\pi_2 \circ \phi = \phi_L \circ \pi_1$, or equivalently that $\text{Gr}(\phi) \leq L$. This formally establishes the diffeomorphism between each non-exceptional fiber $(\mathcal{L}_{2n}^k)_{(B_1, B_2)}$ and the associated symplectic group $\text{Sp}(B_1^\omega \cap B_2^\omega)$. □

As to the final claim of theorem III.2.6 (that the L_i^ω are transverse Lagrangian subspaces of the symplectic vector space V_s) we observe that $L_1^\omega \cap L_2^\omega = \{0\}$ (as $L \notin H$) and that both L_i^ω are isotropic subspaces of dimension k in V . It follows from $\dim(V_s) = 2k$ that the L_i^ω are maximal non-intersecting isotropic subspaces in V_s , and therefore transversal Lagrangians in V_s . □

Remark III.2.7. Note that the bundle structure of each stratum Λ_{2n}^k are each 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 product of quotient symplectic vector spaces is determined uniquely,

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

It is this symplectic vector space over which the fiber is defined: $\Lambda_{2n-2k}^0(B_1^\omega/B_1 \times B_2^\omega/B_2) \cong$

²⁷See remark III.2.3.

$\mathrm{Sp}(2n - 2k)$. Consequently when $B_1 \pitchfork B_2^\omega$ and we have the above canonical identifications $B_1^\omega/B_1 \cong B_1^\omega \cap B_2^\omega \cong B_2^\omega/B_2$ this too must depend smoothly on the isotropic pair (B_1, B_2) as an element in the $2n - 2k$ dimensional symplectic Grassmannian $\mathrm{SG}_{2n-2k}(V)$.

Part IV

The Conley-Zehnder Index and the Circle Map ρ

IV.1 Construction of the Conley-Zehnder Index

In addition to the mean index, the circle function ρ (as defined in definition I.3.7) is more often seen when constructing the Conley-Zehnder index, e.g. [5, 78, 82, 80].

Definition IV.1.1. [82] We define $\rho: \text{Sp}(2n) \rightarrow S^1$ as follows. Given $A \in \text{Sp}(2n)$ let $E = \text{Spec}(A) \cap (S^1 \cup \mathbb{R})$ be the collection (with repetition) of the 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 generalized eigenspace E_λ where,

$$\begin{aligned} Q_\lambda: E_\lambda \times E_\lambda &\rightarrow \mathbb{R} \\ (z, z') &\mapsto \text{im}(\omega(z, \bar{z}')). \end{aligned}$$

Letting m^- denote the sum of the algebraic multiplicities for the real negative eigenvalues we let

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

For our purposes we will be assuming our eigenvalues are unique²⁸ so that the corresponding symmetric forms Q_λ will all have two positive eigenvalues or zero, as there is no possibility of overlapping elliptic eigenvalue pairs (each of potentially varying Krein type). In particular we may choose a collection $(\lambda_i)_{i=1}^l \subset S^1 \setminus \{\pm 1\}$ for which each $\lambda_i \neq \bar{\lambda}_j$ for any $i, j \leq l$ and each Q_{λ_i} has two positive eigenvalues. Additionally we require that every elliptic eigenvalue of A is either (a) contained in the list, or (b) the conjugate of some λ_i in the list (i.e. the ordered tuple $(\lambda_i)_{i=1}^l$ is a certain collection of representatives from each elliptic eigenvalue pair). Additionally, since we will be working with ρ^2 we may ignore the factor

²⁸In theorem V.2.1, we have a sequence $\{A_i\}_{i=1}^\infty \subseteq \text{Sp}(V)$ of semi-simple symplectic maps, i.e. the minimal polynomial $\mathcal{M}(A_i) \in \mathbb{R}(x)$ has distinct irreducible factors over \mathbb{R} and in turn, unique eigenvalues over \mathbb{C} .

counting negative eigenvalue multiplicity and write the following simplified formula in terms of the above choice of elliptic eigenvalues;

$$\rho^2(A) = \prod_{1 \leq i \leq l} \lambda_i^2.$$

Proposition IV.1.2. [82] The map $\rho: \mathrm{Sp}(2n) \rightarrow S^1$ has the following properties:

1. (*determinant*) For $A \in U(n) \subset \mathrm{Sp}(2n)$ we have $\rho(A) = \mathrm{Det}_{\mathbb{C}}(A)$.
2. (*invariance*) ρ is invariant under conjugation,

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

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

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

$$\text{then } \rho(A \oplus B) = \rho(A)\rho(B).$$

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 $L \in \Lambda_{2n}^0$ (with the expected power of two factored in) but for Lagrangian subspaces contained in $\mathcal{L}_{2n}^{\geq 1}$, most properties no longer have an analog. In particular, example VII.1.2 exhibits that $\hat{\rho}$ is not a circle map over \mathcal{L}_2 (see definition I.3.7). Of course, this statement is trivial as each \mathcal{L}_{2n} is not a Lie group, but even if we ignore the Lie group assumption and simply require $\hat{\rho}$ to satisfy the remaining condition (that $\hat{\rho}_*$ is an isomorphism of fundamental groups) it still fails to satisfy this. Namely, on \mathcal{L}_2 we show that $\hat{\rho}$ induces a non-injective homomorphism of fundamental groups $\hat{\rho}_* : \pi_1(\mathcal{L}_2) \rightarrow \pi_1(S^1)$, as there exists a non-contractible loop $\gamma : I \rightarrow \mathcal{L}_2$ for which $\hat{\Delta}(\gamma) = 0$.

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

As described in the previous section, $\hat{\rho}$ does inherit some of the above properties with the caveat that they are generally defined only for $L \in \Lambda_{2n}^0$ and in turn are essentially trivial restatements of proposition IV.1.2. Two non-trivial properties may be found in the following lemma.

Lemma IV.2.1.

1. The operation $(*)^l$ (equation (I.3.1)) is a continuous map for any $l \geq 0$.
2. (Homogeneity) Assuming theorem I.4.3 is true, that is $\hat{\rho}$ is continuous and $\hat{\rho}(L) = \rho^2(\phi)$, then given any $L \in \mathcal{L}_{2n}$ with graph part $\text{Gr}(\phi)$ we claim $\hat{\rho}(L^l) = 2l \cdot \hat{\rho}(L)$ 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 $\text{Sp}(2n-2k) \hookrightarrow \mathcal{L}_{2n}^k \rightarrow \mathcal{I}_k$ for each $1 \leq k \leq n-1$. Since lemma II.2.2 shows that L and L^l share the same domain and range for all $l \in \mathbb{N}_{\geq 1}$, then each $(*)^l$ preserves fibers on \mathcal{L}_{2n}^k for all $0 \leq k \leq n$ (which is a trivial statement in the $k = 0$ case being that the entire stratum is the fiber). Since the fibers vary smoothly over the base we see that $(*)^l$ is continuous as it inherits continuity on each fiber from the group operation on $\text{Sp}(2n-2k)$ via the diffeomorphism $\text{Sp}(B_1^\omega \cap B_2^\omega) \cong (\mathcal{L}_{2n}^k)_{(B_1, B_2)}$ shown in theorem III.2.6. When $k = 0$ this operation corresponds to the group operation in $\text{Sp}(2n)$ 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 $L^0 = \text{Gr}(\text{Id}_V)$. Regardless, since the image of each stratum of $(*)^0$ is the single point $\text{Gr}(\text{Id}_V) = \Delta_V$, 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. We compute L^2 in coordinates using the decomposition shown

in theorem III.2.6 in which we write $L \in \mathcal{L}_{2n}$ as,

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

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

Let $\langle v_i \rangle_{i=1}^k = \ker(L)$ and $\langle w_i \rangle_{i=1}^k = \text{halo}(L)$. 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, $\ker(L) \times \{0\} \oplus \{0\} \times \text{halo}(L) \leq L \circ L$. Next we consider any $(v, w) \in \text{Gr}(\phi)$, so that as $\text{Im}(\phi) = V_g$, we must have $w = \phi(v) \in V_g$. Then since ϕ is non-singular there must be some unique $z \in V_g$ for which $(w, z) \in \text{Gr}(\phi)$. It follows from definition I.3.2 that $(v, z) \in L \circ L$ so that we might observe from $w = \phi(v)$ and $z = \phi(w)$ that a coordinate free²⁹ representation $\text{Gr}(\phi) \circ \text{Gr}(\phi) = \text{Gr}(\phi^2)$ may be written. The remaining basis elements not considered are those of the $\text{halo}(L)$ composed with $\ker(L)$, yet given any $(0, w_i) \in \text{halo}(L)$ in the first L and $(v_i, 0) \in \ker(L)$ belonging to the second, the only resulting vector derived from these is $(0, 0) \in L \circ L$ regardless of whether $L \in H$ or not (more precisely regardless of whether $r > 0$ or not). Consequently we may write,

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

for any $l \geq 1$ so that $\hat{\rho}(L^l) = \rho^2(\phi^l) = 2l \cdot \rho^2(\phi) = 2l \cdot \hat{\rho}(L)$, inheriting homogeneity (scaled by a factor of two) from ρ on $\text{Sp}(V_g)$. When $l = 0$ we have $\hat{\rho}(L^0) = \hat{\rho}(\text{Gr}(\text{Id})) = 0 \in \mathbb{R}/\mathbb{Z}$, so homogeneity holds in this case too. \square

In the following proposition we enumerate some of the inherited properties of $\hat{\rho}$, of which 1 – 3 are entirely trivial (as they only hold over $\mathcal{L}_{2n}^0 \cong \text{Sp}(V)$) whereas 4 – 5 are essentially trivial, being immediate consequences of the definition.

²⁹Free from some fixed Darboux coordinates. To write the above though we must still be working over the fixed decomposition $V = V_s \oplus V_g$.

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

1. (*determinant*): For $A \in U(n) \subset \text{Sp}(2n)$ we have that $\hat{\rho}(\text{Gr}(A)) = (\det_{\mathbb{C}}(A))^2$.
2. (*invariance*): $\hat{\rho}$ is invariant under conjugation on Λ_{2n}^0 ,

$$\hat{\rho}(\text{Gr}(B^{-1}) \circ \text{Gr}(A) \circ \text{Gr}(B)) = \hat{\rho}(\text{Gr}(A)) \quad \forall B \in \text{Sp}(2n).$$

3. (*normalization I*): $\hat{\rho}(\text{Gr}(A)) = 1$ if A has no elliptic eigenvalues.
4. (*normalization II*): $\hat{\rho}(L) = 1$ for all $L \in \mathcal{L}_{2n}^n$.
5. (*multiplicativity*): If $L = K_1 \oplus K_2 \in \mathcal{L}_{2n} \times \mathcal{L}_{2m} \subset \mathcal{L}_{2(n+m)}$ then $\hat{\rho}(K_1 \oplus K_2) = \hat{\rho}(K_1)\hat{\rho}(K_2)$.

Proof: (1-4). These follow from the defining equation $\hat{\rho}(\text{Gr}(A)) = \rho^2(A)$;

1. $\hat{\rho}(\text{Gr}(A)) = \rho^2(A) = \det_{\mathbb{C}}(A)^2$.
2. Recalling that Lagrangian composition is identical to multiplication on $Sp(V)$ (i.e. $\text{Gr}(A) \circ \text{Gr}(B) = \text{Gr}(AB)$) we see from proposition IV.1.2 that,

$$\hat{\rho}(\text{Gr}(B^{-1}AB)) = \rho^2(B^{-1}AB) = \rho^2(A) = \hat{\rho}(\text{Gr}(A)).$$

3. $\hat{\rho}(\text{Gr}(A)) = \rho^2(A) = (\pm 1)^2 = 1$.
4. Assuming theorem V.1.1 to be true, then this is an immediate corollary.

□

Proof: (5). 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

$$L = (\ker(K_1) \times \{0\} \oplus \{0\} \times \text{halo}(K_1) \oplus \text{Gr}(\phi_{K_1})) \\ \oplus \\ (\ker(K_2) \times \{0\} \oplus \{0\} \times \text{halo}(K_2) \oplus \text{Gr}(\phi_{K_2})).$$

Then as $\text{dom}(K_1), \text{ran}(K_1) \leq \langle x_i, y_i \rangle_{i=1}^n$ and $\text{dom}(K_2), \text{ran}(K_2) \leq \langle u_i, v_i \rangle_{i=1}^m$ then theorem III.2.6 implies the two subspaces V_s^j and V_g^j , $j = 1, 2$ determined by K_1 and K_2 have pair-wise 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 $\ker(L) = \ker(K_1) \oplus \ker(K_2)$, $\text{halo}(L) = \text{halo}(K_1) \oplus \text{halo}(K_2)$ and $\text{Gr}(\phi_L) = \text{Gr}(\phi_{K_1}) \oplus \text{Gr}(\phi_{K_2})$. It follows that we may write $\phi_L = \phi_{K_1} \oplus \phi_{K_2} \in \text{Sp}(2n) \times \text{Sp}(2m) \subset \text{Sp}(2(n+m))$. To conclude the proof we assume theorem I.4.3 is true, in which case $\hat{\rho}$ is indeed multiplicative;

$$\hat{\rho}(L) = \rho^2(\phi_L) = \rho^2(\phi_{K_1})\rho^2(\phi_{K_2}) = \hat{\rho}(K_1)\hat{\rho}(K_2).$$

□

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 no elliptic eigenvalues.

Theorem V.1.1. Suppose $\{A_i\}_{i=1}^\infty \subset \mathrm{Sp}(2n)$ is a sequence of semi-simple linear symplectomorphisms for which $\mathrm{Gr}(A_i) \xrightarrow{i \rightarrow \infty} L_1 \times \{0\} \oplus \{0\} \times L_2 \in \mathcal{L}_{2n}$ where $L_1, L_2 \in \Lambda_n$ and $L_1 \cap L_2 = \{0\}$. Then there exists $K \in \mathbb{N}$ such that

$$\mathrm{Spec}(A_i) \cap (S^1 \setminus \{\pm 1\}) = \emptyset, \quad (\text{V.1.1})$$

for all $i \geq K$.

Proof. ³⁰ Suppose there exists some $\lambda_i \in \mathrm{Spec}(A_i)$ for which $\lambda_i \in S^1 \setminus \{\pm 1\}$ for all $i \geq K \in \mathbb{N}$. Then as there are no multiple roots (each A_i is semi-simple) such a λ_i belongs to a unique symplectic eigenvalue pair $\{\lambda_i, \bar{\lambda}_i\}$, $|\lambda_i| = 1$. 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.

Remark V.1.2. As our eigenvalue pair is unique and elliptic for all $i \in \mathbb{N}$, the eigenvalues are stable in the sense that the limit point λ of the λ_i (under some subsequence) is restricted to $\lambda \in S^1$. Namely, any unique pair of elliptic eigenvalues remain elliptic away from the exceptional points ± 1 , being 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, e.g. a quadruple converging to an elliptic pair already inhabited by another pair of eigenvalues). In other words we know that our

³⁰We will proceed in four steps, see lemmas V.1.3, V.1.4, V.1.5, V.1.6 below.

unique pair of elliptic eigenvalues (which persist for all but finitely many $i \in \mathbb{N}$) guarantees that their limit is confined to S^1 .

This pair of eigenvalues has eigenvectors $x_i \pm iy_i \in \mathbb{C}^{2n}$ with a convergent subsequence of corresponding real eigenspaces $E_i = \langle x_i, y_i \rangle$ with some subsequence yielding $E_i \rightarrow E$ as $i \rightarrow \infty$ (due to the compactness of the symplectic Grassmannian $\text{SG}_2(V)$) such that each A_i restricted to E_i are conjugate to a rotation for all $i \in \mathbb{N}$. Then we fix a pair of normalized vectors $x := \lim_{i \rightarrow \infty} \frac{x_i}{|x_i|}$ and $y = \lim_{i \rightarrow \infty} \frac{y_i}{|y_i|}$ for which $\langle x, y \rangle = E$.

As the E_i are symplectic and the A_i semi-simple, there exists a decomposition $V = E_i \oplus F_i$ where the $F_i \in \text{SG}_{2n-2}(V)$ form a sequence of some symplectic complements to each E_i and $A_i = \psi_i \oplus \phi_i: E_i \oplus F_i \rightarrow E_i \oplus F_i$ where ψ_i and ϕ_i are symplectic for each $i \in \mathbb{N}$.

Lemma V.1.3. Let $\{A_i\}_{i=1}^\infty \subset \text{Sp}(2n)$ denote a sequence of symplectic maps and consider the following;

1. For $\{A_i\}_{i=1}^\infty$ such that $\text{Gr}(A_i) \rightarrow L \in \Lambda_{2n}^{\geq 1}$ then $v \in \ker(L) \Leftrightarrow A_i v \rightarrow 0$ and $v \in \text{halo}(L) \Leftrightarrow A_i^{-1} v \rightarrow 0$ as $i \rightarrow \infty$.
2. For $\{A_i\}_{i=1}^\infty$ with graph limit $\text{Gr}(A_i) \rightarrow L = L_1 \times \{0\} \oplus \{0\} \times L_2$ it is true that $L_1 \cap A_i^{-1} L_2 = \{0\}$ for sufficiently large i .

Note that both claims of lemma V.1.3 hold regardless of whether $L \in H$ or not.

Proof 1: We prove the first part of the lemma, after which the second will follow via contradiction, so suppose $\{A_i\}_{i=1}^\infty$ such that $\text{Gr}(A_i) \rightarrow L \in \Lambda_{2n}^{\geq 1}$.

1. $v \in \ker(L) \Leftrightarrow A_i v \rightarrow 0$.

We observe $v \in \ker(L)$ if and only if $(v, 0) \in L$ so that since $\text{Gr}(A_i) \rightarrow L$ then we see that $(v, 0) \in L$ if and only if $(v, A_i v) \rightarrow (v, 0)$ as $i \rightarrow \infty$.

2. $v \in \text{halo}(L) \Leftrightarrow A_i^{-1}v \rightarrow 0$.

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

□

Proof 2: We now prove the second claim of the lemma: given some sequence $\{A_i\}_{i=1}^\infty \subset \text{Sp}(2n)$ for which $\text{Gr}(A_i) \rightarrow L \in \Lambda_{2n}^n$ and suppose there exists a sequence $\{v_i\}_{i=1}^\infty \subset V$ and some $N \in \mathbb{N}$ for which $v_i \rightarrow v \neq 0$ 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 $|A_i v_i| \rightarrow 0$ and $v_i \in A_i^{-1}L_2$ implies that $w_i := \frac{A_i v_i}{|A_i v_i|} \in L_2$ (note that $\lim_{i \rightarrow \infty} w_i$ exists as each w_i is normalized). We see that $|A_i^{-1}w_i| \rightarrow 0$ so that,

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

so that since $|v_i| \rightarrow |v| \neq 0$ we have the contradiction $|A_i v_i| \rightarrow \infty$.

□

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 (regardless of where the limit winds up in Λ_{2n}). 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 $|A_i v_i| \rightarrow 0$ and $|A_i w_i| \rightarrow \infty$. Denote $\psi_i := A_i|_{E_i}$ so that $|\psi_i v_i| \rightarrow 0$ and $|\psi_i w_i| \rightarrow \infty$ and consider the following lemma.

Lemma V.1.4. Given a sequence of elliptic eigenspaces $E_i \rightarrow E$ and symplectic maps $\{\psi_i : E_i \rightarrow E_i\}_{i=1}^\infty$ 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. Suppose $E \cap L_1 = \{0\}$, that is every sequence $u_i = v_i + w_i \in E_i$ has $w_i \not\rightarrow 0$. We would have for every $u_i = v_i + w_i \in E_i$ that

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

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 neighborhoods³¹ $U_i \subset E_i$ about zero on which

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

Then for any $M > 1$ there exists a K for which any normalized sequence $u_i \in E_i$, $u_i \rightarrow u \neq 0$ 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 $i \geq K$. As each ψ_i is a symplectomorphism we may reach a contradiction by choosing any $M > 1$, in which case some K exists for which

$$1 = \int_{U_K} \omega_K = \int_{\psi_K U_K} \omega_K \geq \int_{MU_K} \omega_K > 1. \quad (\text{V.1.2})$$

□

This lemma leads us to the next lemma needed to prove theorem [V.1.1](#).

Lemma V.1.5. Given the eigenspaces $E_i \rightarrow E$ and $\{\psi_i\}_{i=1}^\infty$ as above then $E \cap L_2 \neq \{0\}$.

Proof. We have already established there exists a sequence $v_i \rightarrow v \neq 0$ of $v_i \in E_i$ for which $A_i v_i \rightarrow 0$. We consider any $u_i = v_i + w_i \in E_i$ where $v_i \in L_1$ and $w_i \in A_i^{-1} L_2$ and suppose $w_i \rightarrow w \neq 0$ (such a decomposition exists for sufficient large i by lemma [V.1.3](#)). 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 $|A_i v_i| \rightarrow 0$ as $i \rightarrow \infty$ we see $\lim_{i \rightarrow \infty} \frac{A_i u_i}{|A_i w_i|} = \lim_{i \rightarrow \infty} \frac{A_i w_i}{|A_i w_i|} \rightarrow w \in L_2 \cap E$. □

³¹These U_i may grow without bound but it is of no consequence since we will not pass to the limit.

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 $v_i \rightarrow v \in L_1$ and $w_i \rightarrow w \in L_2$.

Lemma V.1.6. 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 $\{v_i\}_{i=1}^\infty$ with each $v_i \in E_i$ and $\frac{v_i}{|v_i|} \rightarrow v \neq 0$ 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 $(x'_i, y'_i) = (\sigma_i x_i, \sigma_i y_i)$ plane³² E_i by $\lambda_i = e^{i\theta_i}$ so that we may compute for any sequence $\{v_i\}_{i=1}^\infty$,

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

We observe that 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,

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

³²We assume the (x_i, y_i) and (x'_i, y'_i) to be normalized for all $i \in \mathbb{N}$.

With squared norm,

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

On the other hand,

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

so that

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

It's evident that $|\psi_i v_i|^2 = |\psi_i^{-1} v_i|^2$. Since the last conclusions above follow from lemma [V.1.3](#) we see when E_i is a sequence of two dimensional elliptic eigenspace which persists for arbitrarily large i then any sequence $v_i \in E_i$ where $v = \lim_{i \rightarrow \infty} \frac{v_i}{|v_i|}$ we see that $v \in \ker(L) \Leftrightarrow v \in \text{halo}(L)$. \square

Since our hypothesis guarantees that such a sequence of elliptic E_i eigenspaces exist for all but finitely many $i \in \mathbb{N}$, we see that any sequence $\{A_i\}_{i=1}^{\infty} \subset \text{Sp}(2n)$ for which $\text{Gr}(A_i) \rightarrow L \in \mathcal{L}_{2n}^n$ has only real pairs or the usual symplectic quadruples away from the unit circle for sufficiently large i . \square

Remark V.1.7. Note that the above proof is independent of where the convergent subsequence of the λ_i might converge to, so that in particular it holds even for those elliptic λ_i which converge along a subsequence to some $\lambda = \pm 1$.

V.2 A Decomposition for Certain Unbounded Sequences of Symplectic Maps

Theorem V.2.1. Consider any sequence $\{A_i\}_{i=1}^{\infty}$ of semi-simple symplectic maps for which

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

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

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

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

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

uniquely determined by L such that each $\beta_i : E_g^i \rightarrow E_g^i$ is conjugate via I_i to some $\phi_i \in \text{Sp}(V_g)$ for all $i \geq N$ where $\phi_i \rightarrow \phi$. 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.2.6 that L induces the symplectic splitting $V_s \overset{\omega}{\oplus} V_g = V$ for which $L|_{V_s \times V_s} = \ker(L) \times \{0\} \oplus \{0\} \times \text{halo}(L)$ and $L|_{V_g \times V_g} = \text{Gr}(\phi)$. We will be working over a subsequence of the A_i but for brevity's sake will notate this subsequence as A_i as we have been doing throughout the thesis.

Lemma V.2.2. For some $N \in \mathbb{N}$ there exists a subsequence of splittings $\{E_s^i, E_g^i\}_{i=N}^{\infty}$ of V by symplectic 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;

$$E_s^i \xrightarrow{i \rightarrow \infty} E_s \text{ and } E_g^i \xrightarrow{i \rightarrow \infty} E_g,$$

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

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

Lemma V.2.3. Given the above decomposition, we claim that $\text{Gr}(\alpha_i) \rightarrow \ker(L) \times \{0\} \oplus \{0\} \times \text{halo}(L)$ and $\text{Gr}(\beta_i) \rightarrow \text{Gr}(\beta)$ for some $\beta \in \text{Sp}(E_g)$ in the appropriate dimension isotropic Grassmannian, namely $I_{4k}(V \times \bar{V})$ and $I_{4n-4k}(V \times \bar{V})$ respectively.

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:

- $\ker(L) \cap E_g^i = \{0\}$
- $\text{Proj}_{\text{halo}(L)}(E_g^i) = \{0\}$.

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

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

Lemma V.2.5. For all $i \geq N$ there exists a unique sequence of symplectic isomorphisms $I_i: E_g^i \rightarrow V_g$ such that $I_i = \pi_1|_{E_g^i}$ and $I_i \xrightarrow{i \rightarrow \infty} I: E_g \rightarrow V_g$ where the function $\pi_1: \text{dom}(L) \rightarrow \text{dom}(L)/\ker(L) \cong V_g$ is the coisotropic projection map with image the set $\text{dom}(L)/\ker(L)$ uniquely identified by theorem III.2.6 with V_g .

Lemma V.2.6. Defining $\phi_i := I_i \circ \beta_i \circ I_i^{-1} \in \text{Sp}(V_g)$, then $\phi_i \rightarrow \phi: V_g \rightarrow V_g$ where $\text{Gr}(\phi)$ 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 Preparation for the Proof of Theorem V.2.1

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

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

where $E_{[\lambda]}$ is the real eigenspace associated to the quadruple $(\lambda, \lambda^{-1}, \bar{\lambda}, \bar{\lambda}^{-1})$ satisfying $E_{[\lambda]}^{\mathbb{C}} = E_{\lambda} \oplus E_{\lambda^{-1}} \oplus E_{\bar{\lambda}} \oplus E_{\bar{\lambda}^{-1}} \leq V^{\mathbb{C}}$ and E_{λ} denotes the (generalized³³) complex eigenspace associated to λ . Note that our requirement that $\lambda \in D_2^+$, the upper half unit disk, is simply a convenient method for picking a candidate from each quadruple, as well as providing a unique limit point for any eigenvalues quadruples which diverge (in which case the representative chosen from that quadruple always 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 $\text{im}(\lambda) \neq 0$ 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 $A \in \text{Sp}(2n)$ 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 of $E_{[\lambda]} \oplus E_{[\frac{1}{\lambda}]}$;

³³For arbitrary symplectic A_i this is a necessary specification, but as our eigenvalues are unique this is present only for completeness.

$$A|_{E_{[\lambda]} \oplus E_{[\frac{1}{\lambda}]}} = \begin{pmatrix} R(re^{-i\theta})^{-1} & 0 \\ 0 & R(re^{-i\theta})^t \end{pmatrix} = \begin{pmatrix} \frac{1}{r}R(e^{i\theta}) & 0 \\ 0 & rR(e^{i\theta}) \end{pmatrix},$$

where

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

It has been shown that each A_i (having unique eigenvalues) is the direct sum of a combination of the above symplectic eigenspaces [38] although in general the normal form problem is complicated quite a bit by the presence of repeat eigenvalues (particularly those equal to ± 1).

We proceed by considering the eigenvalues of each A_i for all $i \in \mathbb{N}$ as a sequence of tuples $(\lambda^i) := (\lambda_1^i, \dots, \lambda_{2n}^i) \in \mathbb{C}^{2n}$ 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 $\sigma \cdot (\lambda_1, \dots, \lambda_{2n}) = (\lambda_{\sigma(1)}, \dots, \lambda_{\sigma(2n)})$ for any permutation $\sigma \in S_{2n}$. Following [49] 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,

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

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 yield a minimum in the above metric). Given any order for (λ^1) we choose the order of the i^{th} eigenvalue list $(\lambda^i) = (\lambda_1^i, \dots, \lambda_{2n}^i)$ for any $i \geq 2$ by choosing a permutation which minimizes the above metric with respect to the previous element, i.e. $(\lambda_{\tau(j)}^i) = (\lambda_j^{i-1})$ where τ is the minimizing permutation chosen by the above

metric (V.2.1). This allows us to treat the tuple as ordered given the order of the first.

With this notion we let $\{\lambda_j^i\}_{j=1}^l$ denote the l representatives from each eigenvalue quadruple of A_i which lies in the closed upper half disc so that $V = \bigoplus_{j=1}^l E_{[\lambda_j^i]}$ for all $i \in \mathbb{N}$. As mentioned above, each sequence $E_{[\lambda_j^i]}$ 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 distinct eigenvalues/eigenspaces (which do not converge to 0 or diverge) will not converge to a single value/subspace (as ϕ too is semi-simple) so that the dimension of each eigenspace on which $\|A_i|_{E_{[\lambda_j^i]}}\|$ is bounded will be preserved due to the requirement that ϕ remains semi-simple, precluding this possibility [49]. That said, the details of this argument will need to wait until we have split each A_i into its divergent and convergent parts, α_i and β_i respectively. Now that we have established the prerequisite notions, we are ready to begin the proofs.

V.2.2 Proof of Theorem V.2.1

Lemma V.2.2. For some $N \in \mathbb{N}$ there exists a subsequence of splittings $\{E_s^i, E_g^i\}_{i=N}^\infty$ of V by symplectic 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;

$$E_s^i \xrightarrow{i \rightarrow \infty} E_s = \ker(L) \oplus \text{halo}(L) = V_s \text{ and } E_g^i \xrightarrow{i \rightarrow \infty} E_g.$$

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 $\|A_i|_{E_{[\lambda_j^i]}}\| = \sup_{v \in E_{[\lambda_j^i]}} \frac{|A_i v|}{|v|} \rightarrow \infty$ or $\|A_i|_{E_{[\lambda_j^i]}}\| = \sup_{v \in E_{[\lambda_j^i]}} \frac{|A_i v|}{|v|} \rightarrow c_j \in \mathbb{R}$. 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 $v_i \rightarrow v \neq 0$ yet $\frac{|A_i v_i|}{|v_i|} \rightarrow \infty$ so that $|A_i v_i| \rightarrow \infty$. Then

by setting $v'_i = \frac{v_i}{|A_i v_i|}$ then $v'_i \rightarrow 0$ and we see

$$\frac{|A_i v'_i|}{|v'_i|} = \frac{|A_i v_i|}{|v_i|} \rightarrow \infty, \quad (\text{V.2.2})$$

and since $(v'_i, A_i v'_i) \in \text{Gr}(A_i|_{E_{[\lambda_j^i]}})$ for each i we see upon passing to the limit that $(v'_i, A_i v'_i) = (v'_i, \frac{A_i v_i}{|A_i v_i|}) \rightarrow (0, w) \in L$ where $\lim_{i \rightarrow \infty} \frac{A_i v_i}{|A_i v_i|} = w \in \text{halo}(L)$ (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, given 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; $\{\lambda_j^i\}_{i=1}^\infty \xrightarrow{i \rightarrow \infty} \lambda_j$ and the associated quadruple eigenspaces $\{E_{[\lambda_j^i]}\}_{i=1}^\infty$ (note that in the limit the eigenvalue is allowed to vanish). We define

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

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

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

We know that $V = E_s^i \overset{\omega}{\oplus} E_g^i$ for each i and so $A_i = \alpha_i \oplus \beta_i \in \text{Sp}(E_s^i) \times \text{Sp}(E_g^i)$. 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 of constant dimension). Since they reside in a compact space and there are no eigenvectors converging to some single vector in the limit.

Now we state the aforementioned proposition required for lemma [V.2.2](#).

Proposition V.2.2.1. Equations [\(V.1.2\)](#), [\(V.2.2\)](#) together imply $\ker(L) \oplus \text{halo}(L) = E_s$.

Proof. 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

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

since each $w \in \text{halo}(L)$ 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 $|A_i^{-1}w_j^i| \rightarrow 0$. Using the argument found in lemma V.1.4 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 $|A_i v_j^i| \rightarrow \infty$. Additionally 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 $z_j^i \rightarrow z_j \neq 0$ such that $|A_i z_j^i| \rightarrow 0$ for all $j \leq k$. Thus $\langle z_j \rangle_{j=1}^k = \ker(L)$ implies that $\ker(L) \leq E_s$ so that since $\ker(L) \cap \text{halo}(L) = \emptyset$ we see that $\ker(L) \oplus \text{halo}(L) \leq E_s$. Furthermore, since $A_i|_{E_g^i}$ converges we see for dimensional reasons that $E_s = \ker(L) \oplus \text{halo}(L)$. \square

By proposition V.2.2.1, there exists a subsequence for which both E_s^i and E_g^i converge to symplectic subspaces $E_s, E_g \leq V$ (this, as above, is a consequence of identifying $E_s^i \in \text{SG}_{2k}(V)$ and $E_g^i \in \text{SG}_{2n-2k}(V)$ as each is a compact space). For $v \in E_s$ it's true that $\lim_{i \rightarrow \infty} |A_i v| \rightarrow 0, \infty$ whereas for $v \in E_g$ we have $\lim_{i \rightarrow \infty} |A_i v| \rightarrow d \in (0, \infty)$ so that certainly $E_s \cap E_g = \{0\}$. Since these exhaust the set of eigenspaces (which span V as each A_i has unique eigenvalues) we must have $V = E_s \oplus E_g$. We also see that since $\dim(E_s) = 2k$ we must have $\dim(E_g) = 2n - 2k$ and the first lemma of the proof is finished. \square

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

Proof. Beginning by showing that $\text{Gr}(\beta) = \lim_{i \rightarrow \infty} A_i|_{E_g^i}$ we recall the subsequence for which

$E_g^i \rightarrow E_g$ so that since $\text{Gr}(\alpha_i) \oplus \text{Gr}(\beta_i) = \text{Gr}(A_i) \rightarrow L$ we see that $\text{Gr}(\beta_i) \rightarrow K_g \leq E_g \times E_g$. If $K_g \neq \text{Gr}(\beta)$ for any $\beta \in \text{Sp}(E_g)$ 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 $K_g = \text{Gr}(\beta)$ for some $\beta \in \text{Sp}(E_g)$.

As for $\text{Gr}(\alpha_i) = \{(v, A_i v) \mid v \in E_s^i\}$ 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 $\langle v_j^i \rangle_{j=1}^k \rightarrow \ker(L)$ and $\langle w_j^i \rangle_{j=1}^k \rightarrow \text{halo}(L)$. Then we see that $(v_j^i, A_i v_j^i) \rightarrow (v_j, 0) \in \ker(L) \times \{0\}$ and $(A_i^{-1} w_j^i, w_j^i) \rightarrow (0, w_j) \in \{0\} \times \text{halo}(L)$ for each $1 \leq j \leq k$ so that $\text{Gr}(\alpha_i) \rightarrow K_s = \ker(L) \times \{0\} \oplus \{0\} \times \text{halo}(L)$ as desired. \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:

- $\ker(L) \cap E_g^i = \{0\}$
- $\text{Proj}_{\text{halo}(L)}(E_g^i) = \{0\}$.

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

Proof. We begin by proving the following proposition for lemma V.2.4,

Proposition 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 $i \rightarrow \infty$ if and only if there exists an $N \in \mathbb{N}$ for which $\text{Proj}_{\text{halo}(L)}(E_{[\lambda_j^i]}) = \{0\}$ for all $i \geq N$.

Proof. First we clarify that $\text{Proj}_{\text{halo}(L)} : V \rightarrow \text{halo}(L)$ with $\ker(\text{Proj}_{\text{halo}(L)}) = \text{dom}(L)$. Now suppose $F_i := \text{Proj}_{\text{halo}(L)}(E_{[\lambda_j^i]})$ 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 $\{w_i\}_{i=1}^\infty \in E_{[\lambda_j^i]}$ with $w_i \rightarrow w \neq 0$ so that since $V = \text{halo}(L) \oplus \text{dom}(L)$ we may write $w_i = f_i + g_i \rightarrow f + g = w$ where each $f_i \in \text{halo}(L)$

and $g_i \in \text{dom}(L)$. Then since $|A_i^{-1}f_i| \rightarrow 0$ 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 $g_i \rightarrow g$ with $|g| < \infty$, we suppose that the same holds for f_i , in which case it must be $|A_i v_i| \rightarrow \infty$ and $\|A_i|_{E_{[\lambda_j^i]}}\|$ is unbounded.

Alternatively if for some $N \in \mathbb{N}$ for which $\text{Proj}_{\text{halo}(L)}(E_{[\lambda_j^i]}) = \{0\}$ 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 $v_i = k_i + g_i \in \ker(L) \oplus V_g = \text{dom}(L)$, so that any $i \geq N$ has $k_i + g_i \rightarrow k + g = v$. Then since $|A_i k_i| \rightarrow 0$ and $A_i g_i \rightarrow \phi(g) \in V_g$ 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 proposition V.2.4.1 shows for sufficiently large i that $E_g^i \leq \ker(L) \oplus V_g$. We conclude this proof with a corresponding proposition regarding the second claim of lemma V.2.4.

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

Proof. Assume $\|A_i\|_{E_{[\lambda_j^i]}} \rightarrow c \in \mathbb{R}$. Then as $\text{Gr}(A_i) \rightarrow L$ this implies that the following sequence converges,

$$\text{Gr}(A_i|_{E_{[\lambda_j^i]}}) \rightarrow K \leq E_{[\lambda_j]} \times E_{[\lambda_j]},$$

where $K \leq L$ is some isotropic subspace in $V \times V$. As established with respect to E_g^i in lemma V.2.3, since the norm remains bounded the limit is the graph of a symplectic map; $\text{Gr}(A_i|_{E_{[\lambda_j^i]}}) \rightarrow \text{Gr}(f) \leq E_{[\lambda_j]} \times E_{[\lambda_j]}$ for some $f \in \text{Sp}(E_{[\lambda_j]})$. 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 $N_j \in \mathbb{N}$ that $\ker(L) \cap E_{[\lambda_j^i]} = \{0\}$ for all $i \geq N_j$ (and all $E_{[\lambda_j^i]} \leq E_g^i$) as a consequence. \square

It's clear that proposition [V.2.4.2](#) concludes the proof of lemma [V.2.4](#), as each E_g^i is by definition just the direct sum of the aforementioned eigenspaces, so as each of the above eigenspaces have trivial image under projection onto $\text{halo}(L)$ for sufficiently large i , and each is a summand of each E_g^i , we see the same holds on E_g^i for $i \geq N$ (taking N to be the maximum of the N_j above). \square

Lemma V.2.5. For all $i \geq N$ there exists a unique sequence of symplectic isomorphisms $I_i: E_g^i \rightarrow V_g$ such that $I_i = \pi_1|_{E_g^i}$ and $I_i \xrightarrow{i \rightarrow \infty} I: E_g \rightarrow V_g$, where the function $\pi_1: \text{dom}(L) \rightarrow \text{dom}(L)/\ker(L) \cong V_g$ is the coisotropic reduced space of $\text{dom}(L)$ uniquely identified by theorem [III.2.6](#) with V_g .

Proof. Since E_g^i is a $2n - 2k$ dimensional subspace of $\text{dom}(L)$ with $\ker(L) \cap E_g^i = \{0\}$, for sufficiently large i there exists a unique symplectic map $\text{dom}(L)/\ker(L) \cong E_g^i$ (refer to theorem [III.2.6](#)) for each i . We denote the above isomorphisms (The co-isotropic reduction of $\text{dom}(L)$ restricted to E_g^i) by $I_i: E_g^i \rightarrow \text{dom}(L)/\ker(L)$ and thus $V_g \cong \text{dom}(L)/\ker(L) \cong E_g^i$ 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 converging subspace of dimension $2n - 2k$ for sufficiently large i (which by lemma [V.2.4](#) eventually has trivial intersection with $\ker(L)$ and is contained in $\text{dom}(L)$, the latter implied by $\text{Proj}_{\text{halo}(L)}(E_g^i) = \{0\}$), implies that the sequence of isomorphisms converge; $I_i \rightarrow I: E_g \rightarrow V_g$. \square

Lemma V.2.6. Defining $\phi_i := I_i \circ \beta_i \circ I_i^{-1} \in \text{Sp}(V_g)$ then $\phi_i \rightarrow \phi: V_g \rightarrow V_g$ where $\text{Gr}(\phi)$ 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.2.6](#) between V_g and the reduced domain $\text{dom}(L)/\ker(L)$ we may define $\phi_i := I_i \circ \beta_i \circ I_i^{-1} \in \text{Sp}(V_g)$ so that the ϕ_i and β_i are

conjugate. Then since $\text{Gr}(A_i) = \text{Gr}(\alpha_i) \oplus \text{Gr}(\beta_i) \rightarrow L$ and $\beta_i \rightarrow \beta$ as $i \rightarrow \infty$ we see from the continuity of the projection and inclusion maps that

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

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 must verify manually that $\rho^2(\phi_i) = \rho^2(\beta_i)$ for $i \geq N$. First if $\lambda_i \in \text{Spec}(\beta_i)$ with eigenvector v_λ^i we let $w_\lambda^i = I_i(v_\lambda^i) \in V_g$ so that $(I_i \circ \beta_i \circ I_i^{-1})(w_\lambda^i) = (I_i \circ \beta_i)(v_\lambda^i) = I_i(\lambda_i v_\lambda^i) = \lambda_i w_\lambda^i$, 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

$$\begin{aligned} Q_i : E_{\lambda_i} \times E_{\lambda_i} &\rightarrow \mathbb{R} \\ (z, z') &\mapsto \text{im}(\omega(z, \bar{z}')). \end{aligned}$$

So as the corresponding bilinear form for ϕ_i will be given by $Q_i \circ (I_i^{-1} \times I_i^{-1})$ we may write it defined over the eigenspace $F_{\lambda_i} = I_i(E_{\lambda_i})$,

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

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

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

Then since each I_i is simply the co-isotropic reduction of $\text{dom}(L)$ 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 $\lim_{i \rightarrow \infty} \rho^2(\beta_i) = \rho^2(\phi)$. \square

V.3 Extending ρ^2 Continuously to \mathcal{L}_{2n}

Theorem I.4.3. Define the map $\hat{\rho}: \mathcal{L}_{2n} \rightarrow S^1$ by $\hat{\rho}(L) := \rho^2(\phi)$ for any $L \in \mathcal{L}_{2n}$ possessing the unique decomposition,

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

Then the map $\hat{\rho}$ is continuous and $\hat{\rho}(\text{Gr}(A)) = \rho^2(A)$ for all $A \in \text{Sp}(V)$.

$$\begin{array}{ccc} & \mathcal{L}_{2n} & \\ \text{Gr} \uparrow & \searrow \hat{\rho} & \\ \text{Sp}(2n) & \xrightarrow{\rho^2} & S^1 \end{array}$$

Figure V.1: The diagram above commutes.

Proof. We first refer above to our implicit use of theorem III.2.6 decomposing L since $\text{dom}(L) \cap \text{halo}(L) = \{0\}$. For every $L = \text{Gr}(A) \in \mathcal{L}_{2n}^0 \cong \text{Sp}(2n)$ we know $\phi = A: V \circlearrowleft$ 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 (see theorem V.1.1) so that $\hat{\rho}(L) = \lim_{i \rightarrow \infty} \rho^2(A_i) = 1$ for any sequence A_i such that $\text{Gr}(A_i) \rightarrow L \in \mathcal{L}_{2n}^n$ and ρ^2 may be (rather trivially) continuously extended to \mathcal{L}_{2n}^n . For $L \in \mathcal{L}_{2n}^k$ with $1 \leq k \leq n-1$ from

theorem [V.2.1](#) we see there exists some $N \in \mathbb{N}$ for which $A_i = \alpha_i \oplus \beta_i \in \text{Sp}(E_s^i) \times \text{Sp}(E_g^i) \cong \text{Sp}(V_s) \times \text{Sp}(V_g)$ (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 lemma [V.2.6](#)) such that

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

For any $(v, w) \in \text{Gr}(\beta)$ 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 $\text{Proj}_{\text{halo}(L)}(E_g^i) = \{0\}$ for $i \geq N$ then $w = w_h + w_g \in \text{halo}(L) \oplus V_g$ with $w_h = 0$. Consequently we see $(v_g, w_g) \in \text{Gr}(\beta) \leq L$ since $v_g \in V_g$ implies that $w_g = \phi(v_g)$, i.e. the two graphs are seen to coincide after removing $\ker(L)$ components from the domain in $\text{Gr}(\beta)$. This is simply an excessive confirmation that the normal form given in theorem [III.2.6](#) is identical to the limit of $\text{Gr}(\alpha_i) \oplus \text{Gr}(\beta_i)$ after what amounts to some column operations on $\text{Gr}(\beta)$, harmless in that they only remove vectors contained in the kernel from $\text{Gr}(\beta)$, which have no effect on the image.

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

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

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

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 $\gamma: I \rightarrow \text{Sp}(2n)$ we have that $\hat{\Delta}(\text{Gr}(\gamma)) = 2\Delta(\gamma)$.

Proof. We begin by restating the construction of the extended mean index for arbitrary paths $\gamma: I \rightarrow \mathcal{L}_{2n}$ using the unique continuous map $\hat{\theta}: I \rightarrow \mathbb{R}$ 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 $\Lambda_{2n}^0 \cong \text{Sp}(2n)$ its clear that any $\gamma \subset \text{Sp}(2n)$ has $\hat{\theta}(t) = 2\theta(t)$ and thus,

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

□

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.4.10 to be an open and dense subset of $C^1([0, 1], \mathcal{L}_{2n})$). Recalling definition II.1.2 we note that $\text{Pr}_I : \mathcal{L}_{2n} \rightarrow \bigsqcup_{k=0}^n \mathcal{I}_k$ is the map sending each $L \in \mathcal{L}_{2n}^k$ to the associated isotropic pair $(\ker(L), \text{halo}(L)) \in \mathcal{I}_k$ 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 $\text{Im}(\text{Gr}) = \mathcal{L}_{2n}^0 \subset \mathcal{L}_{2n}$, 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 $L \in \mathcal{L}_{2n}^0$ if and only if there exists $A \in \text{Sp}(V)$ for which $L = \text{Gr}(A)$, then for all $A \in \text{Sp}(V)$ we have $\text{Pr}_I(\text{Gr}(A)) = \{0\} =: \mathcal{I}_0$.

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

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

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 \mathcal{L}_{2n}^0 for all but finitely many $t \in [0, 1]$ by proposition II.4.18. 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 $\gamma \in \mathcal{P}_{reg}(\mathcal{L}_{2n})$ 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.4.9.

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,

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

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,

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

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

In other words, letting $\gamma_{\mathcal{K}} = \dim(\ker(\gamma))(t)$, this partition marks each time t_i at which $\gamma_{\mathcal{K}}$ has a discontinuity (each of which are isolated jump discontinuities from the 0^{th} to the 1^{st} stratum), which happen finitely often due to our regularity condition (see proposition II.4.18). More precisely, 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.4.7 and theorems I.4.6, I.4.8. This lemma's proof has been relegated to section VI.3 in addition to proof VI.3 for proposition VI.1.3 (that compatibility of stratum-regular paths is an equivalence relation) 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 $\phi_i := \gamma|_{U_i} : U_i \rightarrow \text{Sp}(V)$ then there exists a sequence of bounded paths $\tilde{\phi}_i : [t_i, t_{i+1}] \rightarrow \text{Sp}(V)$ (so that $\Delta(\tilde{\phi}_i)$ is well defined) for which $\Delta(\tilde{\phi}_i) = \hat{\Delta}(\phi_i)$ ³⁴ for each $0 \leq i \leq M$. Consequently $\hat{\Delta}$ may be decomposed via concatenation with respect to the partition,

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

³⁴This equality is slightly incorrect in that the scaling factor of 2 that is an artifact of squaring ρ is absent. We have decided to omit it as it's entirely inconsequential for the proof, as the scaling is uniform over all ϕ_i and over all intervals.

Theorem I.4.6. 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 $i \geq 1$.

Proof. We proceed given any stratum-regular γ and any γ -compatible τ by 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 $\text{Pr}_I(\gamma(t)) = \text{Pr}_I(\tau(t))$ for all $t \in [0, 1]$ which is trivially satisfied on the U_i where $\text{Pr}_I(\gamma(t)) = \text{Pr}_I(\tau(t)) = \{0\}$. More importantly we have $\text{Pr}_I(\gamma(t_i)) = \text{Pr}_I(\tau(t_i)) = (\ker(\gamma(t_i)), \text{halo}(\gamma(t_i))) \in \mathcal{I}_1 \cong \mathbb{R}\mathbb{P}^{2n-1} \times \mathbb{R}\mathbb{P}^{2n-1} \setminus \hat{H}_1$. We will proceed referring to these isotropic pairs in terms of the representative γ .

Then recalling theorem III.2.6, we see any $L, L' \in \mathcal{L}_{2n}^1$ for which $\text{Pr}_I(L) = \text{Pr}_I(L')$ yield a pair of unique decompositions,

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

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

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

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

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

Now in terms of the 0^{th} stratum intersections, lemma VI.1.5 produces a collection of symplectic paths $\phi_i, \eta_i : U_i \rightarrow \text{Sp}(V)$ 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 $\mathcal{L}_{2n}^0 \cong \text{Sp}(V, \omega)$ 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).

All that remains to be shown is that the composite path $\tau \circ \gamma$ is continuous at each t_i . 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. Similarly we will also only consider the question of continuity from one side of t_i , i.e. $t \rightarrow t_i^-$ for $t \in U_i$, as the two-sided limit is a consequence of the aforementioned claim after reversing the time of the path in a neighborhood about t_i and applying the above again to $(\tau \circ \gamma)|_{U_i}$ as $t \rightarrow t_i^-$, effectively proving the desired result for $t \rightarrow t_i^+$.

So 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 $a_l \rightarrow t_i^-$ as $l \rightarrow \infty$. Recall that $\gamma(t_i)$ and $\tau(t_i)$ share the symplectic splitting $V = V_s(t_i) \overset{\omega}{\oplus} V_g(t_i)$ where $V_g(t_i) := \text{dom}(\gamma(t_i)) \cap \text{ran}(\gamma(t_i))$ 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 $l \rightarrow \infty$.

Then since

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

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

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

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 to $\gamma_{\ker}(t_i)$ and $\gamma_{\text{halo}}(t_i)$

respectively as $l \rightarrow \infty$. In particular we desire,

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

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

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

we see that $\text{Gr}(\eta_l \phi_l)$ converges to $\tau(t_i) \circ \gamma(t_i)$ on U_i provided proposition VI.1.7 below holds.

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

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

Proof. We suppose $K \in H$, then there exists a sequence of normalized vectors $v_l \rightarrow v \neq 0$ in V such that $v \in \ker(K) \cap \text{ran}(K)$. It follows from the kernel membership that $|\eta_l \phi_l(v_l)| \rightarrow 0$ as $l \rightarrow \infty$ 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) $|\eta_l(w_l)| \rightarrow 0$ or (b) $|\phi_l(v_l)| \rightarrow 0$ and $|\eta_l(w_l)| < \infty$. In both cases $|\eta_l \phi_l(v_l)| < \infty$ which ensures that $\lim_{l \rightarrow \infty} v_l = v \notin \text{halo}(K)$ (Recall that the v_l are normalized).

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

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

is guaranteed to be symplectic and in particular, a well defined function so that $K_{\text{Gr}}([0]_{\text{dom}}) \neq [0]_{\text{ran}}$ 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.

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

Proof. Since $K := \lim_{l \rightarrow \infty} \text{Gr}(\eta_l \gamma_l) \notin H$ we may invoke theorem V.2.1 to obtain a sequence of symplectic splittings $E_s^l \overset{\omega}{\oplus} E_g^l = V$ for which E_s^l converges to some \tilde{V}_s . Additionally equation (VI.1.1) shows that $E_g^l \rightarrow V_g$ as $l \rightarrow \infty$. Then since $V_g = V_s^\omega$ by theorem III.2.6 (which we are able to use thanks to proposition VI.1.6) we see that it must be that $\tilde{V}_s = V_s$ so that $\tilde{V}_s = \ker(K) \oplus \text{halo}(K) = V_s$ 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 $\text{halo}(K) = \text{halo}((\tau \circ \gamma)(t_i))$.

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 $|(\eta_l \phi_l)(v)| \not\rightarrow 0$ if and only if $v \in \text{halo}(\tau(t_i))$ since $|\phi_l(v)| \rightarrow 0$. Then since $\ker(\gamma(t_i)) \cap \text{halo}(\gamma(t_i)) = \{0\}$ and $\text{halo}(\gamma(t_i)) = \text{halo}(\tau(t_i))$ this can't be possible and the only alternative is that $|(\eta_l \phi_l)(v)| \rightarrow 0$ and $\ker((\tau \circ \gamma)(t_i)) = \ker(K)$.

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

As detailed above, the theorem follows from equation (VI.1.1) as proposition VI.1.7 is true. That is, $(\tau \circ \gamma)(t_i) = K = \ker(\gamma) \times \{0\} \oplus \{0\} \times \text{halo}(\gamma) \oplus \text{Gr}(\phi_i \eta_i)$ and the path $\tau \circ \gamma$ converges to the appropriate $K = (\tau \circ \gamma)(t_i) \in \mathcal{L}_{2n}^1$ and therefore is continuous at each of the t_i , ergo continuous on $[0, 1]$. Since $\phi_i \circ \eta_i \in C^1(U_i, \text{Sp}(V))$ for all $0 \leq i \leq M$ we see that $\gamma \circ \tau$ is piece-wise differentiable on $[0, 1]$, with the only possible discontinuities in the derivative occurring at the t_i . \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.4.7. 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.4.6 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

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

□

Now we recover the quasimorphism bound on the compatibility classes.

Theorem I.4.8. 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 class³⁵.

Proof. Consider two compatible paths $a, b : [0, 1] \rightarrow \mathcal{L}_{2n}$. Then by definition they share a unit interval partition, $0 \leq t_1 < \dots < t_N \leq 1$ and thus also share intervening intervals $U_i \subseteq [0, 1]$

over which we will denote the restrictions $\alpha_i, \beta_i : U_i \rightarrow \text{Sp}(V)$ where $a|_{U_i} = \text{Gr}(\alpha_i)$ and

³⁵Whether there exists a bound which is uniform over all compatibility classes is unlikely to be shown via this method, as arbitrary numbers of intersections with higher strata are possible meaning any attempt in taking the supremum over all compatibility classes without somehow normalizing with respect to this intersection count would lead to $C \rightarrow \infty$ (as C is a linear function of $M = |\text{Im}(\gamma) \cap \mathcal{L}_{2n}^1|$.)

$b|_{U_i} = \text{Gr}(\beta_i)$ for all $i \leq M$. These two maps, being unbounded paths in the symplectic group with $\text{Gr}(\alpha_i(t)) \rightarrow L \in \mathcal{L}_{2n}$ possess ‘bounded approximations’ $\tilde{\alpha}_i, \tilde{\beta}_i : K_i \rightarrow \text{Sp}(V)$ where $K_i = \text{cl}(U_i)$. As the construction of these bounded paths is purely technical, we postpone the details until the proof of lemma VI.1.5 below and simply define them here.

Definition VI.2.1. Given a path $\phi : (0, 1) \rightarrow \text{Sp}(V)$ for which $\|\phi(t)\| \rightarrow \infty$ and $\text{Gr}(\phi(t)) \rightarrow L_0, L_1 \in \mathcal{L}_{2n}$ as $t \rightarrow 0, 1$, we define the bounded approximation $\tilde{\phi} : [0, 1] \rightarrow \text{Sp}(V)$ of ϕ as a path for which there exists some sufficiently small $\epsilon > 0$ such that $\tilde{\phi}|_{[\epsilon, 1-\epsilon]} = \phi|_{[\epsilon, 1-\epsilon]}$ and $\hat{\Delta}(\text{Gr}(\phi)) = \Delta(\tilde{\phi})$. Here $\hat{\Delta}(\text{Gr}(\phi))$ is defined over the closure of the path $\text{Gr}(\phi) : [0, 1] \rightarrow \mathcal{L}_{2n}$ where $\text{Gr}(\phi(t)) = L_0, L_1$ for $t = 0, 1$ respectively. For details regarding the existence and construction of $\tilde{\phi}$ see proof VI.3 below.

Now recalling lemma VI.1.5, we have

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

where the product $\tilde{\alpha}_i \tilde{\beta}_i$ is the usual group operation on $\text{Sp}(V)$ identified with \mathcal{L}_{2n}^0 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 $p_{\tilde{\alpha}_i}, p_{\tilde{\beta}_i}, p_{\tilde{\alpha}_i \tilde{\beta}_i} : I \rightarrow \text{Sp}(V)$ 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 $\text{Sp}(2n)$ may be written as

$$\widetilde{\text{Sp}}(2n) = \{(g, c) \in G \times \mathbb{R} \mid \rho(g) = e^{ic}\},$$

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 $\tilde{\mathcal{P}}_{\text{Id}, \tilde{\alpha}_i(0)}(\text{Sp}(2n))$, 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 $\tilde{\mathcal{P}}_{\text{Id}, \tilde{\alpha}_i(0)}(\text{Sp}(2n))$ then the group action of $\pi_1(\text{Sp}(2n))$ on the universal cover along with lemma 1.3.9 implies there exists a deck transformation (identified with concatenation) induced by some nonzero $[l] \in \pi_1(\text{Sp}(2n))$ so that for some $[p'_{\tilde{\alpha}_i}] \in \tilde{\mathcal{P}}_{\text{Id}, \tilde{\alpha}_i(0)}(\text{Sp}(2n))$ we may write $[p_{\tilde{\alpha}_i}] = [p'_{\tilde{\alpha}_i}] * [l]$. 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 identical S^1 eigenvalues each oriented such that the quadratic form $\text{im}(\omega(z, \bar{z}'))$ 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 $\tilde{\alpha}_i(0)$ 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 so that we may choose each connecting path such that,

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

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,

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

and 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,

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

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

Remark VI.2.2. 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 $\text{Sp}(V)$ 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, $\Pr_I(\tau(t)) = \Pr_I(\tau^\epsilon(t))$ (or in our case the sequence $\{\Pr_I(\tau(t_i))\}_{i=1}^M$) 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 by compatibility classes, we note that every stratum-regular path is contained within its own compatibility class implying that the above collection indeed covers $\mathcal{P}_{reg}(\mathcal{L}_{2n})$ ³⁶. 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 $\Pr_I(\xi(t_0)) = \Pr_I(\tau(t_0)) \neq \Pr_I(\gamma(t_0)) = \Pr_I(\xi(t_0))$. No such ξ exists and we have,

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

³⁶Note that theorem II.4.15 shows this is in fact an equivalence relation on an open and dense subset of $C^1([0, 1], \mathcal{L}_{2n})$.

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

□

Remark VI.3.1. 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,

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

If we had used the definition as discussed in the remark I.4.9 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 intersections.

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 $\phi_i := \gamma|_{U_i} : U_i \rightarrow \text{Sp}(V)$ then there exists a sequence of bounded paths $\tilde{\phi}_i : [t_i, t_{i+1}] \rightarrow \text{Sp}(V)$ (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 $\text{Im}(\gamma)$ and \mathcal{L}_{2n}^1 at each t_i (recall definition VI.1.4) by $\Phi_i = \text{Pr}_{\text{Sp}}(\gamma(t_i)) \in \text{Sp}(V_g(t_i)) \cong \text{Sp}(2n-2)$ for each $0 \leq i \leq M$ recalling that

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

where $\text{SG}_{2k}(V) = \{W \in \text{Gr}_{2k}(V) \mid \omega|_W \text{ is nondegenerate}\}$ is the Grassmannian of $2k$ dimensional symplectic subspaces of V for $1 \leq k \leq n$. Assuming for some $\epsilon > 0$ that both Φ_i and $\phi_i(t)$ (the graph part of $\gamma|_{U_i(t)} = \text{Gr}(\phi_i(t))$) have semi-simple eigenvalues for all $0 \leq i \leq M$ and $t \in (t_i, t_i + \epsilon] \cup [t_{i+1} - \epsilon, t_{i+1})$, then we see from [49] that the eigenvalues of each $\phi_i(t)$ (with repetition) form a continuous path³⁷ approaching t_i, t_{i+1} in the space of unordered \mathbb{C} -tuples (as in the proof of theorem V.2.1), $(\lambda)_i^- : (t_i, t_i + \epsilon] \rightarrow \tilde{\mathbb{C}}^{2n} := \mathbb{C}^{2n}/S_{2n}$ and $(\lambda)_i^+ : [t_{i+1} - \epsilon, t_{i+1}) \rightarrow \tilde{\mathbb{C}}^{2n}$. As opposed to the proof for theorem V.2.1, we will be assuming each $(\lambda)_i(t)$ is an unordered list of all eigenvalues (obviously with repetition if necessary) and not concerning ourselves with the extra symplectic structure associating certain quadruples and pairs.

Without loss of generality we will be working with the connected component $(t_i, t_i + \epsilon]$ and in doing so will use (λ) (omitting the i -index altogether) to refer to the path $(\lambda)_i^-$ as defined above. We let $(\hat{\lambda}) \in \tilde{\mathbb{C}}^{2n-2}$ denote the eigenvalues of Φ_i (the reduced dimension is a consequence of $\Phi_i \in \text{Sp}(2n-2)$) and recall that the topology induced on the space $\tilde{\mathbb{C}}^{2n}$ 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 $\max_{j \leq 2n} |\lambda_j| \leq |(\lambda)| = |\sigma(\lambda)|$ for all $(\lambda) \in \tilde{\mathbb{C}}^{2n}$ (where the maximum runs over each eigenvalue $\lambda_j \in (\lambda)$ and $|(\lambda)|$ denotes the usual norm on \mathbb{C}^{2n} applied to the

³⁷Note that these are only continuous on $(t_i, t_i + \epsilon]$, as by design a pair of these eigenvalues will diverge as $t \rightarrow t_i^+$.

unordered tuple), we see that $|\ast| : \widetilde{\mathbb{C}}^{2n} \rightarrow \mathbb{R}$ is a well defined continuous function. This implies that the unordered tuple (λ) is bounded (with respect to the above minimax metric) as $t \rightarrow t_i^+$ 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 two distinct subsets of eigenvalues based on their asymptotic behavior; the following map is well defined for all $t \in (t_i, t_i + \epsilon]$ given sufficiently small ϵ ,

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

where the $\lambda_s(t)$ is defined as the subset $\{\lambda_j\}_{j=1}^2 \subset (\lambda)$ for which given any $\delta > 0$ there exists $s \in (t_i, t_i + \epsilon]$ such that either $|\lambda_j(t)| < \delta$ or $|\lambda_j(t)| > \delta^{-1}$ for each $j = 1, 2$ and for all $s \in (t_i, t)$. In other words λ_s consists of those eigenvalues quadruples/pairs which are unbounded³⁸ (alternatively vanishing) as $t \rightarrow t_i^+$; i.e. $|\lambda_s(t)| \rightarrow \infty$ as $t \rightarrow t_i^+$. On the other hand $\lambda_g(t)$ consists of those eigenvalues which remain bounded and non-zero as $t \rightarrow t_i^+$ and therefore $\lambda_g(t)$ converges to a unique set of (unordered) eigenvalues $\lambda_g(t_i) := \lim_{t \rightarrow t_i^+} \lambda_g(t)$ which must be identical to $(\hat{\lambda}) = \text{Spec}(\Phi(t_i)) \in \mathbb{C}^{2n-2}/S_{2n-2}$ due to continuity.

Now consider a continuous bounded map $r : [t_i, t_i + \epsilon] \rightarrow [1, \infty)$ for which $\lambda_g(t) \subset D_{r(t)}^2$ for each $t \in [t_i, t_i + \epsilon]$ (where D_r^2 is the open disk of radius r). The boundedness assumption is possible precisely since $|\lambda_g(t)|$ is bounded as $t \rightarrow t_i^+$. 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 $t \in [t_i, t_i + \epsilon]$ (thereby implying that $\lambda_s(t)$ is contained in the complement) where $A_r = \left\{ z \in \widetilde{\mathbb{C}}^{2n} \mid r^{-1} < z < r, \right\}$ for any $r > 1$. Then since the eigenvalue path $\lambda_g(t)$ consists of distinct eigenvalues which don't vanish and are bounded as $t \rightarrow t_i^+$ then as shown in [49] for sufficiently small ϵ there exists a continuous path of sym-

³⁸This could also suffice as a definition by letting λ_s be the smallest subset of eigenvalues in (λ) for which every $|\lambda_s|$ diverges or vanishes.

plectic eigenspaces $E_g : (t_i, t_i + \epsilon] \rightarrow \text{SG}_{2n-2}(V)$ corresponding to the λ_g (symplectic as each eigenspace of $\phi_i(t)$ is symplectic) which converge as $t \rightarrow t_i^+$. Setting $E_g(t_i) := \lim_{t \rightarrow t_i^+} E_g(t)$ we assume now that the domain of the E_g is the closed interval $[t_i, t_{i+1}]$. Then continuity implies $E_g(t_i) \leq \text{dom}(\Phi_i)$ as an application of lemma V.2.4 guarantees that the image of the co-isotropic projection $E_g(t) \rightarrow \text{halo}(\gamma(t_i))$ is eventually zero on a sequence approaching t_i from the right, and a simple proof by contradiction shows this holds in the continuous case on $[t_i, t_i + \epsilon]$. We note these are the image of a contour integral of the continuous resolvent, $P_g : [t_i, t_i + \epsilon] \times V \rightarrow E_g(t)$, for which $E_g(t_i) \cong V_g(t_i)$ by restricting the co-isotropic projection $\text{dom}(\gamma(t_i)) \rightarrow \text{dom}(\gamma(t_i))/\ker(\gamma(t_i)) \cong V_g(t_i)$ to $E_g(t_i)$ as in lemma V.2.5.

Remark VI.3.2. 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 a more explicit version of the above argument using the graph resolvent, denoting the circle of radius r as S_r^1 , we may give each $P_g(t) : V \rightarrow E_g(t)$ as,

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

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

If we let $j : V \rightarrow V_g(t_i)$ denote the projection along $V_s(t_i)$ onto $V_g(t_i)$ (where $V_s(t_i) = \ker(\gamma(t_i)) \oplus \text{halo}(\gamma(t_i)) \cong V_g^\omega$) then the subset of $W \in \text{SG}_{2n-2}(V)$ for which $(j|_W, \omega|_W)$ is a symplectic isomorphism is open and dense in $\text{SG}_{2n-2}(V)$; $j|_W$ is an isomorphism if and only if $W \cap V_s(t_i)$. It follows that for sufficiently small ϵ that the path of symplectic isomorphisms $j|_{E_g(t)} =: I_i(t) : E_g(t) \rightarrow V_g(t_i)$ induces a unique continuous path of symplectic maps $\beta_i : [t_i, t_i + \epsilon] \rightarrow \text{Sp}(V_g(t_i))$ such that $I(t_i) = \text{Id} \in \text{Sp}(V_g(t_i))$ (implying that $\beta_i(t_i) = \Phi_i$). Additionally we see that the following holds for all $t \in [t_i, t_i + \epsilon]$,

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

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




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

We are about ready to construct our bounded approximation $\tilde{\phi}_i$ but first we must set up the necessary notation for the other side of U_i , namely the boundary neighborhood $[t_{i+1} - \epsilon, t_{i+1}]$. We let $F_g : [t_{i+1} - \epsilon, t_{i+1}] \rightarrow \text{SG}_{2n-2k}(V)$ denote the continuous path of eigenspaces, $\hat{I}_i(t) : F_g(t) \rightarrow V_g(t_{i+1})$ the family of isomorphisms and $\hat{\beta}_i : [t_{i+1} - \epsilon, t_{i+1}] \rightarrow \text{Sp}(V_g(t_{i+1}))$ the associated symplectic path. It follows that $\beta_i|_{E_g(t)}$ and $\hat{\beta}_i|_{F_g(t)}$ converge to $\Phi_i \in \text{Sp}(V_g(t_i))$ and $\Phi_{i+1} \in \text{Sp}(V_g(t_{i+1}))$ as $t \rightarrow t_i^+, t_{i+1}^-$ 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 may choose a pair of sequences $\{a_i\}_{i=1}^\infty, \{b_i\}_{i=1}^\infty$ for which $a_i, b_i \in U_i$ for all $i \in \mathbb{N}$ and that $a_i \rightarrow t_i^+$ and $b_i \rightarrow t_{i+1}^-$ as $i \rightarrow \infty$. Then for all $\epsilon > 0$ there exists some $N \in \mathbb{N}$ large enough such that

- $\rho^2(\alpha(a_i)) = \rho^2(\hat{\alpha}(b_i)) = 1$ for all $i > N$.
- $a_N \in (t_i, t_i + \epsilon]$ and $b_N \in [t_{i+1} - \epsilon, t_{i+1})$.

Now we may fix the pair of elements $\alpha_N \in \text{Sp}(E_s^N)$ and $\hat{\alpha}_N \in \text{Sp}(F_s^N)$ 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 $\rho(\alpha_N) = \rho(\hat{\alpha}_N) = 1$ with the plan of effectively stopping the divergent terms in their tracks (although not before they’re safely hyperbolic as shown in theorem V.1.1) near t_i, t_{i+1} while letting their graph portions continue on as usual (so as to preserve the mean index).

Below we define the ‘bounded approximations’ on each U_i ,

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

Since $\rho^2(\alpha_N) = \rho^2(\hat{\alpha}_N) = 1$ we see that for all $t \in [t_i, t_{i+1}]$ that

$$\rho^2(\tilde{\phi}_i(t)) = \hat{\rho}(\gamma(t)),$$

so that in particular, $\rho^2(\tilde{\phi}_i(t)) = \rho^2(\phi_i(t))$ for each $t \in U_i$.

Then any pair of lifts, $\hat{\theta}$ of $\hat{\rho} \circ \gamma : [t_i, t_{i+1}] \rightarrow \mathbb{R}$ (lifting the path in \mathcal{L}_{2n}) and $\tilde{\theta}$ of $\rho \circ \tilde{\phi}_i : [t_i, t_{i+1}] \rightarrow \mathbb{R}$ (lifting the bounded approximation in $\text{Sp}(V)$) for which

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

must be identical everywhere from the above identity; $\rho^2(\tilde{\phi}_i(t)) = \hat{\rho}(\gamma(t))$. It follows from $\hat{\theta}(t) = \tilde{\theta}(t)$ for all $t \in [t_i, t_{i+1}]$ that we have

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

and thus $\hat{\Delta}(\phi_i) = \Delta(\tilde{\phi}_i)$ for each $i \leq M$, which concludes the lemma after noting $\hat{\Delta}$ ’s additivity with respect to concatenation. \square

Part VII

Concluding Remarks

VII.1 The Extended Mean Index on \mathcal{L}_2

As mentioned above, when Δ is restricted to paths in $\mathrm{Sp}(2n)$ originating at the identity then lemma I.3.8 may be applied to write $\Delta : \widetilde{\mathrm{Sp}}(2n) \rightarrow \mathbb{R}$. This is defined via the association of any γ where $\gamma(0) = \mathrm{Id}$ with some $g \in \widetilde{\mathrm{Sp}}(2n)$ by setting $g := \tilde{\gamma}(1)$ (where $\tilde{\gamma}$ is the lifted path). This map is clearly onto as $\widetilde{\mathrm{Sp}}(2n)$ 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 $\hat{\Delta}$ so that the corresponding covering space E has the property that any $[\gamma] \in \pi_1(\mathcal{L}_2)$ which lifts to a non-trivial loop $[\gamma_l] \in \pi_1(E)$ has $\hat{\Delta}(\gamma) = 0$. In other words, when defined on E the mean index is zero on a loop in \mathcal{L}_2 only if that loop, when lifted to E , is non-trivial. It remains unclear whether these non-trivial loops arising from the removal of H should be tossed out or if they carry important information which should not simply be ignored.

Example VII.1.2. For $\mathrm{Sp}(2, \mathbb{R}) \hookrightarrow \mathcal{L}_2$ we have the nice geometric interpretation; $\mathrm{Sp}(2, \mathbb{R}) \cong D^2 \times S^1$ so that we may write

$$\Lambda_2 = S^2 \times S^1 / \sim_1 \cong \overline{D}^2 \times S^1 / \sim_2,$$

where \sim_1 identifies $(x, t) \sim_1 (-x, -t)$ for all $(x, t) \in S^2 \times S^1$. On the other hand \sim_2 has one element equivalence classes for all interior points $(r, \theta, t) \in [0, 1) \times S^1 \times S^1$ with non-trivial identifications occurring only in the boundary on which antipodal points are identified with respect to the second S^1 term while the boundary of \overline{D}^2 is left fixed; $(1, \theta, t) \sim (1, \theta, t + \pi)$. The two models may be reconciled by seeing each $D_2 \times \{t\}$ as the identified pair of the upper hemisphere of $S^2 \times \{t\}$ and the lower hemisphere of $S^2 \times \{t + \pi\}$, explaining why the quotient maps are so different. We write $p \in S^2 \times S^1$ using cylindrical coordinates on the S^2 term;

$p = (z, \theta, t) \in [1, 1] \times S^1 \times S^1$ and see that

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

i.e. we have removed the (now identified) equators from the spheres at $t = \pi/2, 3\pi/2$. We observe that \mathcal{L}_2 may be identified as a subset of the latter quotient in a nearly identical manner.

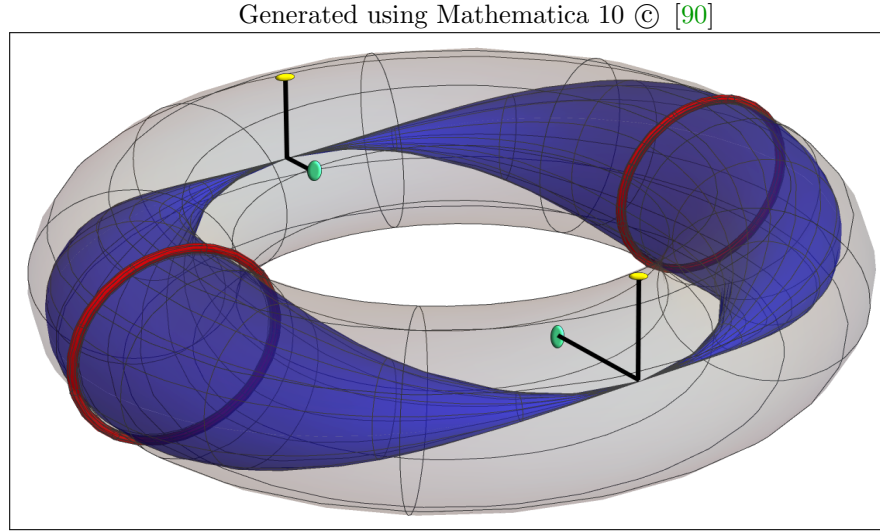


Figure VII.1: The interior of the solid torus corresponds to $\Lambda_2^0 \cong \text{Sp}(2)$ while the rest of Λ_2 , namely Λ_2^1 , is the quotient of the boundary torus using the aforementioned identification, depicted here by two pair of like colored ellipses (note the lack of rotation in the minor 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’.

We now consider the fundamental group as promised in remark I.3.13. It’s an easy observation that $\pi_1(\mathcal{L}_2) = \mathbb{Z}[\eta] * \mathbb{Z}[\tau]$ where $[\eta]$ is the push-forward of the generator for $\pi_1(\text{Sp}(2))$ and τ corresponds to a loop about the missing circle, i.e. the fundamental group of \mathcal{L}_2 is the free group on two generators;

Consider the minimal covering space $\text{Pr} : E \rightarrow \mathcal{L}_2$ satisfying the following condition; that any non-trivial loop in \mathcal{L}_2 with zero mean index has a lift homotopic to some non-trivial




Figure VII.2: The simplest model space $X = S^1 \vee S^1$ for which $\pi_1(\mathcal{L}_{2n}) \cong \pi_1(X)$.

element in $\pi_1(E)$. In terms of group theory we seek the minimal cover E for which the induced subgroup $\text{Pr}_*(\pi_1(E))$ of $\mathbb{Z}[\eta] * \mathbb{Z}[\tau]$ consists of those fully reduced homotopy words of the type $\eta^l \tau \eta^{-l}$, which we may characterize formally as the following subgroup of infinitely many generators;

$$\text{Pr}_*(\pi_1(E)) = \langle \eta^l \tau \eta^{-l} \mid \forall l \in \mathbb{Z} \rangle.$$

This particular covering space is depicted below in figure VII.3 in terms of the analogous cover over $S^1 \vee S^1$.




Figure VII.3: A covering space for $S^1 \vee S^1$ analogous to the covering space E , note the similarities with figure VII.4 below. Such an identification is possible after establishing some isomorphism $\pi_1(S^1 \vee S^1) \cong \pi_1(\mathcal{L}_2) \cong \mathbb{Z} * \mathbb{Z}$.

Remark VII.1.3. We first note that $\text{Pr}_*(\pi_1(E)) = \langle g_i \rangle_{i \in \mathbb{N}}$ is isomorphic to the infinite rank free group where $e_i := \eta^i \tau \eta^{-i}$. Despite this we note that the above graph depicted in figure VII.3 is maximally symmetric³⁹ which is equivalent (see [40]) to $\text{Pr}_*(\pi_1(E))$ being a normal subgroup of $\pi_1(\mathcal{L}_2)$ where the symmetries referenced above correspond to the quotient group $\pi_1(\mathcal{L}_2) / \text{Pr}_*(\pi_1(E)) \cong \mathbb{Z}$ of deck transformations on E (namely, arbitrary left and right shifts). We also observe that for each $L \in \mathcal{L}_2$, the group of deck transformations acts

³⁹In the sense that any vertex may be mapped to any other vertex in a manner preserving the edge labels (such a cover is often referred to as normal, regular or Galois).

transitively on the fiber $\text{Pr}^{-1}(L)$ and $E \rightarrow \mathcal{L}_2$ has the structure of a principal \mathbb{Z} bundle.

Generated using Mathematica 10 © [90]

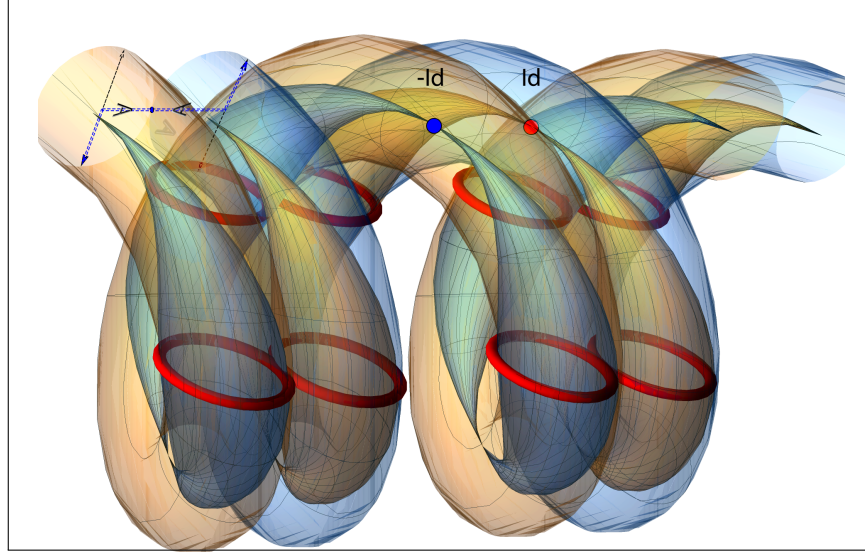


Figure VII.4: The covering space E is depicted here as two attached copies of $\overline{D}^2 \times \mathbb{R}$, the blue and orange helical cylinders (see below for an explicit construction) while the red rings represent the connected components of $\text{Pr}^{-1}(H)$. Note that each adjacent pair is in fact just one circle (as above in figure VII.1), their single point of contact is due to the boundary identifications only being shown along a single line. A significant alteration was required in that the orange helical cylinder is involuted about the minor axis so as to ensure the identifications which fix the minor radius are possible to depict in three dimensions (see the boundary identifications on the upper left). More importantly, the orange copy has been rotated about its major axis by π (recall $t \mapsto t + \pi$), signaled in the figure by the red and blue dots, the first of which lies in $\text{Pr}^{-1}(\text{Id})$ and the latter in $\text{Pr}^{-1}(-\text{Id})$.

We may alternatively construct E in a geometric fashion (see figure VII.4 above); let $K = \partial D^2 \times (\frac{\pi}{2} + \pi\mathbb{Z}) = \text{Pr}^{-1}(H)$ and $\phi : \partial D^2 \times \mathbb{R} \rightarrow \partial D^2 \times \mathbb{R}$ be a map identifying the boundary of a solid cylinder with another via $(\theta, t) \sim_{\phi} (\theta, t + \pi)$ for all $t \in \mathbb{R}$. Then it follows we may realize E as a pair of infinite solid cylinders (helical cylinders above) identified along their boundaries via ϕ ,

$$E = (\overline{D}^2 \times \mathbb{R}) \sqcup (\overline{D}^2 \times \mathbb{R}) / \sim_{\phi} \setminus K.$$

Remark VII.1.4. One might recall our alternate definition I.3.10 in which the extended

mean index is characterized as the lift $\hat{\Delta} : \tilde{\mathcal{L}}_{2n} \rightarrow \mathbb{R}$ of the map $j^* \hat{\rho} : \tilde{\mathcal{L}}_{2n} \rightarrow S^1$ where $j : \tilde{\mathcal{L}}_{2n} \rightarrow \mathcal{L}_{2n}$ is the covering map. This lift is well behaved in the classic case of $\text{Sp}(V)$ in part because of the isomorphic fundamental groups $\pi_1(\Lambda_{2n}) \cong S^1$ whereas in our case some redundancy is likely to exist in $\tilde{\mathcal{L}}_{2n}$ for many $n \geq 2$ (like it does when $n = 1$ as seen in example VII.1.2 above). An analogous procedure for larger n may exist⁴⁰ which could use the correspondence between inclusion in the subgroup lattice of $\pi_1(\mathcal{L}_{2n})$ and the presence of non-trivial coverings generating a lattice of covering spaces (up to isomorphism) to identify an analogous covering space E for $n \geq 2$.

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 $B \rightarrow M$ and $C \rightarrow M$ 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.

⁴⁰This is assuming that the codimension of H leads to some form of regularity in the fundamental group sequence $(\pi_i(\mathcal{L}_{2n}))_{i \in \mathbb{N}}$ in the sense that they're all (or mostly all) large free groups, in which case the rich theory for free groups of countable rank [52] may prove useful.

2. Next we would need that $T_x L \notin \text{LagGr}(T_x M) \setminus H_x$ 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 $H_x = \{L \in \text{LagGr}(T_x M) \mid \pi_{B_x}(L) \cap (L \cap (\{0\}_x \times C_x)) \neq \{0\}_x\}$ as well as some further transversality conditions as detailed in [88] (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. free 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 [87] and [53]. 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’ $\text{LagGr}_\bullet(V)$. As a set $\text{LagGr}_\bullet(V) = \Lambda_{2n} \times \mathbb{N}$ albeit equipped with a topology quite different from the product topology (which they call the Sabot topology). They use a discrete metric $d(L, L') = \text{codim}(L \cap L')$ (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 losing 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, \dots, 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 results of part 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

Proposition VII.3.1. Such a V_g as described in proposition III.2.6.2 with an associated unique $\phi \in \text{Sp}(V_g)$ exists 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.2.6 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 and example VII.3.2 below for such a Lagrangian in \mathbb{R}^{2n} .

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.2.6, 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 $V_g \cong_{\pi_{1,2}} \hat{V}_g$. 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 $\phi \in \text{Sp}(V_g)$. 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 $\dim((L_1^\omega \oplus L_2^\omega) \cap (L_1 \cap L_2)) = \dim((\ker(L) \oplus \text{halo}(L)) \cap V_g) = 0$. \square

Despite the continuity transversality requirement (namely $\ker(L) \cap \text{halo}(L) = \emptyset$) being equivalent to $r = 0$, we see that here a larger r seems to add room for potential domains analogous to our V_g to fit inside $L_1 \cap L_2$ and survive the quotienting process, thereby 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 ideal 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 required to ensure continuity of composition that $r = 0$, particularly considering there's often little concern with a non-zero κ in such applications. It suggests that these two situations are disjoint in the sense that, even if they each might work for a large number of the possible invariants (r, κ, k) individually, they're likely to only work simultaneously in the nicest of situations (e.g. our assumption that $\kappa = r = 0$).

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 $\text{dom}(L) = \text{ran}(L)$ 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

$$\phi = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

with respect to the 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

$$(k_L, r_L, \kappa_L) = (\dim(\ker(L)), \dim(\ker(L) \cap \text{ran}(L)), \dim(\text{dom}(L) \cap \text{ran}(L))),$$

associated to $(\ker(L), \text{halo}(L)) \in \mathcal{I}_{\kappa_L}$ as shown in [62]. Letting $\hat{H} := \{L \in \Lambda_{2n} \mid 0 < r_L < \kappa_L\}$, then $\hat{\rho}$ 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 $H = \{L \in \Lambda_{2n} \mid 0 < r_L\}$ so that $\hat{H} \subset H$ and $\mathcal{L}_{2n} \subset \hat{\mathcal{L}}_{2n}$. It is unlikely $\text{codim}(\hat{H}) > 2$ although the added elements in $\hat{\mathcal{L}}_{2n}$ may have another benefit in enlarging \mathcal{L}_{2n} to satisfy certain groupoid axioms like those found in [65].

We conclude the thesis with the very essence of anti-climax; a mundane 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,

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

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 writing $\lim_{k \rightarrow \infty} A_k = A$ and $\lim_{k \rightarrow \infty} B_k = B$ we see as $k \rightarrow \infty$ we may write the following with respect to the above Darboux coordinates,

$$\begin{aligned} \text{Gr}(A_k) &= \left\langle \left(e_1, \frac{e_1}{k} \right), \dots, \left(e_n, \frac{e_n}{k} \right), (f_1, kf_1), \dots, (f_n, kf_n) \right\rangle \\ &\rightarrow \langle (e_1, 0), \dots, (e_n, 0), (0, f_1), \dots, (0, f_n) \rangle, \\ \text{Gr}(B_k) &= \left\langle \left(e_1, \frac{-e_1}{k} \right), \dots, \left(e_n, \frac{e_n}{k} \right), (f_1, -kf_1), \dots, (f_n, kf_n) \right\rangle \\ &\rightarrow \langle (e_1, 0), \dots, (e_n, 0), (0, -f_1), \dots, (0, f_n) \rangle. \end{aligned}$$

Thus both $\text{Gr}(A_k) \rightarrow L \leftarrow \text{Gr}(B_k)$ 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, Dokazatelstvo teoremy A.N. Kolmogorova o sohranении uslovnopriodičeskikh dviženij pri malom izmenenii funkicii Gamiltona [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 harakterističeskom klasse, vhodjašem v uslovija kvantovanija [On the Characteristic Class coming into Quantization Conditions], *Funkc. analiz i ego pril.* **1** (1967), no. 1, 1–14, (Russian). *Funct. Anal. Appl.*, **1**, (1967), no. 1, 1–13, (English).
- [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] D. Ayala, J. Francis, H. Tanaka, Local Structures on Stratified Spaces, *Adv. in Math.* **307** (2017), 903–1028.
- [5] J. Barge, E. Ghys, Cocycles d'Euler et de Maslov [Euler and Maslov cocycles], *Mathematische Annalen.* **294** (1992), 235–265.
- [6] G. Ben Simon, D. Salamon, Homogeneous Quasimorphisms on the Symplectic Linear Group, *Isr. Jour. of Math.* **175** (2010), 221–224.
- [7] S. Benenti, W. Tulczyjew, Symplectic Linear Relations, *Mem. Accad. Sci. Torino Cl. Sci. Fis. Mat. Natur.* **5** (1981), no. 4, 71–140.
- [8] 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.
- [9] R. Bott, On the Iteration of Closed Geodesics and the Sturm Intersection Theory, *Comm. on Pure and Appl. Math.* **9** (1956), 171–206.
- [10] 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.
- [11] E. Çineli, Conley Conjecture and Local Floer Homology, *Arch. Math.* **111** (2018), no. 6, 647–656.
- [12] E. Çineli, V. Ginzburg, B. Gürel, Pseudo-rotations and Holomorphic Curves, Preprint *arXiv:1905.07567v2 [math.SG]* (2019).
- [13] 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.
- [14] 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.

- [15] J. Eells, A Setting for Global Analysis, *Bull. Amer. Math. Soc.* **72** (1966), 751–807.
- [16] I. Ekeland, H. Hofer, Convex Hamiltonian Energy Surfaces and Their Periodic Trajectories, *Comm. Math. Phys.* **113** (1987), no. 3, 419–469.
- [17] Ekeland, I.: Une Théorie de Morse pour les Systèmes Hamiltoniens Convexes [A Morse Theory for Convex Hamiltonian Systems], *Ann. Inst. H. Poincaré, Anal. non linéaire* **1** (1984), no. 1, 19–78, (French edition).
- [18] 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.
- [19] 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).
- [20] 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.
- [21] A. Floer, A Relative Morse Index for the Symplectic Action, *Comm. Pure Appl. Math.* **41** (1988), no. 4, 393–407.
- [22] A. Floer, The Unregularized Gradient Flow of the Symplectic Action, *Comm. Pure Appl. Math.* **41** (1988), no. 6, 775–813.
- [23] A. Floer, Morse Theory for Lagrangian Intersections, *J. Differential Geom.* **28** (1988), no. 3, 513–547.
- [24] A. Floer, Witten’s Complex and Infinite-Dimensional Morse Theory, *J. Differential Geom.* **30** (1989), no. 1, 207–221.
- [25] A. Floer, Symplectic Fixed Points and Holomorphic Spheres, *Comm. Math. Phys.* **120** (1989), no. 4, 575–611.
- [26] I. M. Gelfand, V. B. Lidskii, O Strukture Oblastej Ustojčivosti Linejnyh Kanoničeskij Sistem Differencialnyh Uravnenij s Periodičeskimi Koëfficientami [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).
- [27] V. Ginzburg, B. Gürel, Action and Index Spectra and Periodic Orbits in Hamiltonian Dynamics, *Geom. Topol.* **13** (2009), no. 5, 2745–2805.
- [28] V. Ginzburg, B. Gürel, Local Floer Homology and the Action Gap, *Jour. of Symplectic Geo.* **8** (2010), no. 3, 323–357.
- [29] V. Ginzburg, B. Gürel, The Conley Conjecture and Beyond, *Arnold Math. Jour.* **1** (2015), no. 3, 299–337.
- [30] V. Ginzburg, B. Gürel, Pseudo-rotations versus Rotations, Preprint *arXiv:1812.05782v1 [math.SG]* (2018).
- [31] V. Ginzburg, B. Gürel, Conley Conjecture Revisited, *Int. Math. Res. Not. IMRN* (2019), no. 3, 761–798.
- [32] V. Ginzburg, B. Gürel, Lusternik-Schnirelmann Theory and Closed Reeb Orbits, *Math. Zeit.* **295** (2020), no. 1-2, 515–582.

- [33] 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).
- [34] M. Gromov, Pseudo-holomorphic curves in Symplectic Manifolds, *Invent. Math.* **82** (1985), no. 2, 307–347.
- [35] 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.
- [36] J. Gutt, The Conley-Zehnder Index for a Path of Symplectic Matrices, Preprint *arXiv:1201.3728v2 [math.DG]*, (2012).
- [37] J. Gutt, Generalized Conley-Zehnder index, *Ann. Fac. Sci. Toulouse Math. (6)* **23** (2014), no. 4, 907–932.
- [38] J. Gutt, Normal Forms For Symplectic Matrices, Preprint *arXiv:1307.2403v2 [math.SG]*, (2014).
- [39] W. Hamilton, On a General Method in Dynamics, *Phil. Trans. R. Soc.* **124** (1834), 247–308.
- [40] A. Hatcher, Algebraic Topology, *Cambridge Univ. Press* (2002), 56–78.
- [41] H. He, Compactification of Classical Groups, *Comm. Anal. Geom.* **10** (2002), no. 4, 709–740.
- [42] G. A. Hedlund, Poincaré’s Rotation Number and Morse’s Type Number, *Trans. Amer. Math. Soc.* **34** (1932), no. 1, 75–97.
- [43] H. Hofer, On Strongly Indefinite Functionals with Applications, *Trans. of the Amer. Math. Soc.* **275** (1983), no. 1, 185–214.
- [44] 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.
- [45] L. Hörmander, Fourier Integral Operators I, *Acta Mathematica* **127** (1971), no. 1-2, 79-183.
- [46] M. Hutchings, Quantitative Embedded Contact Homology, *J. Differential Geom.* **88** (2011), no. 2, 231–266.
- [47] M. Hutchings, Beyond ECH Capacities, *Geom. Topol.* **20** (2016), no. 2, 1085–1126.
- [48] S. Janeczko, Lagrangian Submanifolds in Product Symplectic Spaces, *Journal of Math. Phys.* **41** (2000), no. 8, 5642–5655.
- [49] T. Kato, Perturbation Theory for Linear Operators, *Springer Berlin Heidelberg, Classics in Mathematics* **132** (1995).
- [50] M. Klein, Mathematical Thought from Ancient to Modern Times, Vol. I and II, *New York: Oxford University Press.* (1972), 167–168.
- [51] 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.
- [52] I. Kapovich, A. Myasnikov, Stallings Foldings and the Subgroup Structure of Free Groups, *J. Algebra* **248** (2002), no 2, pp. 608-668.
- [53] D. Li-Bland, A. Weinstein, Selective Categories and Linear Canonical Relations, *Symmetry Integrability Geom. Methods Appl.* **10** (2014).

- [54] C. Liu and Y. Long, An Optimal Increasing Estimate for Maslov-type Indices of Iterations, *Chinese Sci. Bull.* **43** (1997), 2275–2278.
- [55] Y. Long, Multiple Solutions of Perturbed Superquadratic Second Order Hamiltonian Systems, *Trans. Amer. Math. Soc.* **311** (1989), no. 2, 749–780.
- [56] Y. Long, Forced Oscillations for the Triple Pendulum, *Analysis, et cetera* (1990), 177–208, *Academic Press, Boston, MA*.
- [57] 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).
- [58] Y. Long, A Maslov-type Index Theory for Symplectic Paths, *Top. Meth. Nonl. Anal.* **10** (1997), 47–78.
- [59] 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).
- [60] Y. Long, Index Theory for Symplectic Paths with Applications, *Progress in Mathematics* **207**, Birkhäuser Verlag, Basel, (2002).
- [61] 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.
- [62] J. Lorand, Classifying Linear Canonical Relations (Master’s Thesis), *arXiv:1508.0–4568v1[math.SG]*, (2014).
- [63] J. Lorand, Alan Weinstein, Coisotropic Pairs, Preprint *arXiv:1408.5620v1 [math-.SG]*, (2014).
- [64] J. Lurie, Higher Algebra, <https://www.math.ias.edu/~lurie/papers/HA.pdf>.
- [65] K. Mackenzie, On Extensions of Principal Bundles, *Ann. of Global Anal. and Geo.* **6** (1988), 141–163.
- [66] V. P. Maslov, Perturbation Theory and Asymptotic Methods, *Moscow University Press, Moscow*, (1965), (Russian edition).
- [67] V. P. Maslov, V. I. Arnol’d, V. S. Buslaev, Théorie des Perturbations Et Méthodes Asymptotiques [Theory of Perturbations and Asymptotic Methods], *Etudes math.*, Dunod, (1972).
- [68] D. McDuff, F. Schlenk, The Embedding Capacity of 4-dimensional Symplectic Ellipsoids, *Ann. of Math. Second Series* **175** (2012), No. 3, 1191–1282.
- [69] J. Milnor, On the Existence of a Connection with Curvature Zero, *Comment. Math. Helv.* **32** (1958), 215–223.
- [70] M. Morse, The Calculus of Variations in the Large, (Reprint of the 1932 original), *Amer. Math. Soc. Col. Pub.* **18** (1996).
- [71] J. Moser, On Invariant Curves of Area-Preserving Mappings of an Annulus, *Nachr. Akad. Wiss. Gttingen Math.-Phys.* **II** (1962), 1–20.
- [72] J. Nelson, Reflections on Cylindrical Contact Homology, *Instit. For Adv. Study*, (14 May 2020), <https://youtu.be/t5NiA4Xdgk4>, (lecture).

- [73] B. Nica, Group Quasimorphisms, Vanderbilt University (2006, rev. 2008), <http://www.math.mcgill.ca/bnica/quasimorphisms.pdf>, (lecture).
- [74] R. Palais, Lusternik-Schnirelmann Theory for Banach manifolds, *Topology* **5** (1966), 115–132.
- [75] H. Poincaré, Sur les Courbes Définies par les Équations Différentielles [On Curves Defined by Differential Equations], *Journal de Mathématiques* **4** (1885), vol. 1, 167–244.
- [76] P. H. Rabinowitz, Periodic Solutions of Hamiltonian Systems, *Comm. Pure Appl. Math.* **31** (1978), no. 2, 157–184.
- [77] J. Rawnsley, On the Universal Covering Group of the Real Symplectic Group, *Jour. of Geo. and Phys.* **62** (2012), no. 10, 2044–2058.
- [78] J. Robbin, D. Salamon, The Maslov Index for Paths, *Jour. of Topo.* **32** (1993), no. 4, 827–844.
- [79] P. Robinson, J. Rawnsley, The Metaplectic representation, Mp^c Structures and Geometric Quantization, *Mem. of the Amer. Math. Soc.* **81** (1989), no. 410.
- [80] D. Salamon, Morse Theory, the Conley Index and Floer Homology, *Bulletin of the London Mathematical Society* **22** (1990), no. 2, 113–140.
- [81] 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).
- [82] 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.
- [83] K. Siegel, Higher Symplectic Capacities, Preprint *arXiv:1902.01490v3 [math.SG]*, (2020).
- [84] J. Towber, Linear Relations, *Journal of Algebra* **19** (1971), 1–20.
- [85] A. Usher, A History of Mechanical Invention (Revised edition), *Cambridge, MA: Harvard Univ. Press*, (1954).
- [86] C. Viterbo, A New Obstruction to Embedding Lagrangian Tori, *Invent. Math.* **100** (1990), no. 2, 301–320.
- [87] K. Wehrheim, C. Woodward, Functoriality for Lagrangian Correspondences in Floer theory, *Quantum Topol.* **1** (2010), no. 2, 129–170.
- [88] A. Weinstein, Symplectic Categories, *Proc. of Geom. Summer School: Lisbon*, *arXiv:0911.4133v1 [math.SG]*, (2009), (lecture).
- [89] S. Weinberger, The Topological Classification of Stratified Spaces, *The Univ. of Chicago Press, Ltd., London*, (1994).
- [90] Wolfram Research Inc., Mathematica, Version 12.1, Champaign, IL, (2013).
- [91] V. Yakubovich, V. Starzhinskii, Linejnye differentsial'nyye uravneniya s periodičeskimi koëffitsiyentami i ikh prilozheniya [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).