SPECTRAL FLOW IN MORSE THEORY.

Notes by Maksim Maydanskiy.

1 Introduction

Spectral flow is a general formula or computing the Fredholm index of an operator $\frac{d}{ds} + A(s) : L^{1,2}(\mathbb{R}, \mathcal{H}) \to L^2(\mathbb{R}, \mathcal{H})$ for a family A_s of (almost) self-adjoint operators on a Hilbert space $\mathcal{H}$. The simplest version of such an operator is one where $\mathcal{H} = \mathbb{R}^n$. This appears in Morse theory when studying the spaces of gradient trajectories. Thus Morse theory is a natural setting to consider such operators, with the view towards other geometric applications in various flavours of Floer theory.

These notes are based on the two talks on spectral flow that I gave at the Stanford student symplectic geometry seminar. The content is for the most part not original, but rather a compilation of various sources with a few modifications. There is a number of references for the basic setup of Morse theory, of which the most detailed and concise are the notes of Alex Ritter [8], which I generally recommend. The approach to spectral flow is based on the book of Kronheimer and Mrowka [4]. The goal, to paraphrase a famous saying, was to give the simplest detailed exposition to spectral flow possible, but not simpler.

We note that the two parts – Morse theory setup and Spectral flow – are independent of each other.

2 Setup.

2.1 Basic setup – negative gradient trajectories.

We start with a manifold M of dimension n with f a morse function on M and g a Riemannian metric on it.

We denote by f'_g the gradient of f with respect to g that is $g(f'_g, \cdot) = df(\cdot)$. This is a vector field on M.

Let p and q be critical points of f. We consider the space of smooth paths $\mathcal{P} = \{\gamma : \mathbb{R} \to M | \lim_{s\to+\infty} = p, \lim_{s\to-\infty} = q\}$. Such a path is a *negative gradient trajectory* if

$$\frac{d\gamma(s)}{ds} + f'_g(\gamma(s)) = 0.$$

For any $\gamma \in \mathcal{P}$, even one that doesn't solve this equation, the left hand side of it defines a map $S_g(\gamma) : \mathbb{R} \to \gamma^*TM$. Putting the spaces γ^*TM for all $\gamma \in \mathcal{P}$

together gives a bundle $\mathcal{B} \to \mathcal{P}$ and we view S_g as a section of this bundle. The negative gradient trajectories are then the intersection of S_g with zero.

The "modern" or "function-analytic" or "moduli space" approach to Morse homology studies this space of trajectories and uses it and related moduli spaces to define algebraic invariants of M. The advantage of this approach (as opposed to "classical" or "geometric" one) is that it generalizes to other settings (see a nice table in [8, Lecture 10]).

To understand how this zero set of the section F_g behaves it helps to consider zero-sets of sections of finite dimensional bundles first. When we later consider the infinite-dimensional version some of the words will get translated. We mention these translation in parenthesis.

With this in mind, let $B^{n+m} \to P^n$ be a fiber bundle and s its section. We identify P with the zero section (and hence any point $p \in P$ with it's image in B). In that case, if s and the zero section are transverse, their intersection is a submanifold of P of dimension $n - m$ by implicit function theorem (it is empty if $m > n$). Transversality at a point p in $\operatorname{Im}(s) \cup P$ is the statement that $T_p \operatorname{Im}(s) \oplus T_p P = T_p B = T_p P \oplus T_p F_p$, where F_p is the fiber of B over p. Of course $T_p \operatorname{Im}(s)$ is the $\operatorname{Im}(Ds)$, where Ds is the derivative (or "linearizarion") of s, $Ds : T_p P \to T_p B$. Composing Ds with projection to $T_p F_p$ (taking "vertical part of the linearization") gives the linear map (operator) $Ds^{vert} : T_p P \to T_p F_p$. The section is transverse to the zero section at p if and only if Ds^{vert} is surjective. If it is, the kernel of it is the tangent space to the "moduli space" $\operatorname{Im}(s) \cap P$. Note that the dimension of that moduli space is the "index" of Ds^{vert} – dimension of kernel minus dimension of cokernel. Of course if Ds^{vert} is surjective cokernel is 0, but the index makes sense even if the map is not surjective and is constant – equal to $n - m$ – essentially by rank-nulity theorem.

We want to repeat this argument in our – infinite dimensional* – situation. To be able to do so we need to overcome several technical difficulties. We now mention some of them.

1) To be able to apply implicit function theorem in infinite-dimensional setting we need $\mathcal{B}$ to be a Banach bundle over Banach manifold $\mathcal{P}$. However, the space of smooth paths in $\mathcal{P}$ is not Banach (ultimately, because $C^\infty(\mathbb{R})$ is not Banach), so we need to work with appropriate completions.

Our first hunch is to take $L^{1,2}(\mathbb{R}, M)$ (we need once differentiable to make sense of $\frac{d(\gamma(s))}{ds}$ and $L^{1,2}$ is a Hilbert space, which is handy). However, because $\mathbb{R}$ is non-compact this space does not really make sense. Indeed, usually to define

*We are still doing finite-dimensional Morse theory. It's the abstract "opeartor as a section" setup that is infinite-dimensional.

$L^{k,2}(N, M)$ one either embeds M into some $\mathbb{R}^N$ or uses charts on M which works fine for compact *domain* N (see, for example, [8, Lecture 11]). However, for a non-compact domain, in both of these approaches a constant map $\gamma : \mathbb{R} \to p$ may be assigned infinite norm and excluded from $L^2(\mathbb{R}.M)$. This is a problem in general.

For us, the solution is to notice that there is no issue defining $L^{1,2}_{loc}(\mathbb{R}, M)$ or $L^{1,2}(\mathbb{R}, \mathbb{R}^n)$. Moreover by Sobolev embedding, paths in $L^{1,2}_{loc}(\mathbb{R}, M)$ are also in $C(\mathbb{R}, M)$ (see Lemma 3.5 for a proof). So it makes sense to talk about such a path converging to a point. Now for a path $\gamma(s)$ converging to a fixed p for large s we land close enough to p that we can write $\gamma(s) = \exp v(s)$ for a unique $v \in T_pM$ (here $\exp v$ is the endpoint of the geodesic of the metric g starting from p and tangent to v). Identifying $T_pM = R^n$ we require $v(s)$ to be in $L^{1,2}(\mathbb{R}, \mathbb{R}^n)$ (this is independent of the $T_pM = \mathbb{R}^n$ identification). That is the space of paths $\mathcal{P}$ that we actually work with. It is locally modelled on $L^{1,2}(\mathbb{R}, \mathbb{R}^n)$ and so is Banach.

2) Just as in the finite dimensional case, to get the moduli spaces of trajectories to be a manifold we need the section F_g to be transverse to the zero-section. To achieve this we need to perturb the metric g. To that end, we cross $\mathcal{B} \to \mathcal{P}$ with $\mathcal{G}$ – the space of all C^k smooth metrics on M and consider the "universal" section F given by the same formula as F_g but now for varying g. One proves that this universal section is transverse to the zero section, and uses infinite dimensional version of Sard's theorem to conclude that F_g is transverse for a "generic" (co meagre and hence second category) set of gs. However, since g is not smooth, F is also not smooth, but rather only C^k ([8, Lecture 13]). Even the finite-dimensional Sard theorem only holds when the smoothness k is large enough (see [2, p. 69] for an amazing (counter)example of what happens when this smoothness bound does not hold). The same issue appears in the infinite dimensional version and imposes the requirement $k > \max(0, \operatorname{index} F_q)$. In the Morse theory case it is enough to take $k > n$. However, one would still want to have generic smooth metrics to work with. An additional lemma ([8, Lecture 5], see also [6, p. 52]) has to be proved to get that. Alternatively, one can avoid going through C^k metrics altogether by using a Floer space – a Banach space made out of smooth functions (see [1, p.807] also mentioned in [6, Remark 3.2.7, p.53]).

2.2 Linearizing the vertical part of F_g.

We saw above that the operator in question is the vertical part of DF_g at a negative gradient trajectory γ. Thus we want to compute it in local coordinates with the goal of then proving it Fredholm and computing its index. This is done in homeworks 2 and 12 and lectures 13 and 14 in [8].

The chart around γ is given by the geodesic flow with respect to the metric g,

so for a vector field v along γ we compute

$$DF_g^{vert}(v) = \nabla_t|_{t=0}(\partial_s w + f_g'(w)) = (\nabla_t \partial_s w + \nabla_t f_g'(w))_{t=0} = \nabla_s v + (\nabla_v f_g')_v.$$

Here $w(s,t) = \exp_{v(s)} tv(s)$. Using parallel frame e_i write $v = \sum v^i e_i$ and compute $\nabla_s v = \sum \partial_s v^i e_i$, $\nabla_v f' g = \sum v^j \nabla_{e_j} f' g(\gamma) = \sum (A_s)^i_j v^j e_i$ (this defines the matrices A_s). Thus in these coordinates

$$\begin{aligned} L^{1,2}(\mathbb{R}, \mathbb{R}^n) &\to L^2(\mathbb{R}, \mathbb{R}^n) \\ DF_g^{vert}(v) &= (\partial_s + A_s)v. \end{aligned}$$

We note that $\lim_{s\to+\infty} = \operatorname{Hess} p$ and $\lim_{s\to-\infty} = \operatorname{Hess} q$, non-degenerate symmetric matrices. We will prove the "index=spectral flow" formula for such operators.

3 The spectral flow.

We have now reduced our problem to studying the Fredholmness and index properties of

$$F_{A_s} = \partial_s + A_s : L^{1,2}(\mathbb{R}, \mathbb{R}^n) \to L^2(\mathbb{R}, \mathbb{R}^n).$$

A discussion of Fredholm operators can be found, among other places in Tom Mrowka's notes [5].

We assume that $\lim_{s\to+\infty} A_s = A_+$ and $\lim_{s\to-\infty} A_s = A_-$ with $A_\pm$ invertible. Recall that for a symmetric matrix $\operatorname{Index} A$ is the numbers of negative eigenvalues.

Our goal is to prove the spectral flow theorem: $\operatorname{Index} F_A = \operatorname{Index} A_- - \operatorname{Index} A_+$. Our exposition is an adaptation of the approach of Kronheimer and Mrowka [4, Chapter 14].

3.1 The constant case – Invertibility.

Consider the case when $A_s = A$ is independent of t. In fact, we will start even simpler, with $F_\lambda = \partial_s + \lambda : L^{1,2}(\mathbb{R}, \mathbb{R}) \to L^2(\mathbb{R}, \mathbb{R})$.

Firstly, if $\lambda = 0$ the operator is *not* Fredholm! †.

As an illustration, consider $f(t) = \frac{1}{t}$ for $t > 1$ smoothly cut of to zero on $[0, 1]$ and zero on $(-\infty, 0]$. Then f is in $L^2(\mathbb{R}, \mathbb{R})$, but is not in the image of any $L^{1,2}$ function (as it's integral is $\ln(t) + C$ on $t > 1$ and is not square integrable. The

† If you are used to compact domains this may be surprising, as it is elliptic. In particular $\partial_s : L^{1,2}(S^1, \mathbb{R}) \to L^2(S^1, \mathbb{R})$ is Fredholm. Things are different on the non-compact domain $\mathbb{R}$.

same thing happens for all functions decaying as t^α for $\alpha \in (-3/2, -1/2)$ and so the cokernel of ∂_s is not finite dimensional.

However if $\lambda \neq 0$ the operator is in fact invertible.

Proposition 3.1 ([4]14.1.2)**.** *$F_\lambda = \partial_s + \lambda : L^{1,2}(\mathbb{R}, \mathbb{R}) \to L^2(\mathbb{R}, \mathbb{R})$ is invertible.*

Proof. The proof is based on Fourier transform. Recall that $FT(u) = \hat{u}$ is the function given by $\hat{u}(\xi) = \int_{-\infty}^{+\infty} u(s) e^{it\xi}\, ds$. Plancherel theorem says that FT is an isomorphism of $L^2(\mathbb{R}, \mathbb{C})$ to itself. Its more general version says that FT of $L^{1,2}(\mathbb{R}, \mathbb{C})$ is the completion of $C_c^\infty(\mathbb{R}, \mathbb{C})$ in the norm $||\hat{u}(\xi)||_{FT(L^{1,2})} = \int_{-\infty}^{+\infty} (1 + \xi^2)|u(\xi)|^2\, d\xi$.

Of course the reason we like FT is that it takes $\partial_s + \lambda$ to D – the multiplication by $i\xi + \lambda$. Thus in the transformed form it has inverse D^{-1} which is multiplication by $(i\xi + \lambda)^{-1}$. We compute:

$$\begin{aligned} ||D^{-1}\hat{u}||_{FT(L^{1,2}} = \int_{-\infty}^{+\infty} (1+\xi^2)|(i\xi + \lambda)^{-1} u(\xi)|^2\, d\xi \quad &= \quad \int_{-\infty}^{+\infty} \frac{1+\xi^2}{\lambda^2 + \xi^2} |u(\xi)|^2\, d\xi \\ &\leqslant \quad \max\left(1, \frac{1}{\lambda^2}\right) ||\hat{u}(\xi)||_{L^2}. \end{aligned}$$

So the inverse of F_λ (given by Fourier transforming, dividing and Fourier transforming back) lands in $L^2(\mathbb{R}, \mathbb{R})$ so F_λ is invertible as a map $L^{1,2}(\mathbb{R}, \mathbb{C}) \to L^2(\mathbb{R}, \mathbb{C}$. Finally note that the space of real functions is Fourier transformed to the space of those functions that have $\overline{u(\xi)} = u(-\xi)$ and this condition is unaffected by multiplying or dividing by $(i\xi + \lambda)$, so both F_λ and its inverse restrict to that subspace, proving our proposition.

□

Note: There is a more explicit way to do this given in [7, p.6]. Namely, the fact that the equation $F_A(u) = \eta$ for $\eta \in L^2(\mathbb{R}, \mathbb{R})$ has at most one solution in $L^{1,2}(\mathbb{R}, \mathbb{R})$ is easy (the difference solves the homogeneous equation and can not be in $L^2(\mathbb{R}, \mathbb{R})$). This solution can be found explicitly: in the case $\lambda > 0$ set $u(s) = \int_{-\infty}^{s} e^{-\lambda(s-t)} \eta(t)\, dt = \Phi * \eta$ where Φ is 0 for $s < 0$ and $e^{-\lambda t}$ for $t > 0$ (this is coming from Laplace transforms, ∂_s transforming to ζ, Φ transforming to $\frac{1}{\zeta + \lambda}$ and convolution transforming to multiplication), and Young inequality gives $||u||_{L^2} \leqslant ||\Phi||_{L^1} ||\eta||_{L^2}$. Since u solves the equation $\dot{u} = \eta - Au$, the derivative $\dot{u}$ is also in L^2. In the case $\lambda > 0$ the solution for $\eta(s)$ is $-u(-s)$ where $u(s)$ is the solution for $-\lambda$ and $\eta(-s)$.

Now if we consider $F_A : L^{1,2}(\mathbb{R}, \mathbb{R}^n) \to L^2(\mathbb{R}, \mathbb{R}^n)$ with constant A then by changing basis on the target $\mathbb{R}^n$ we get $A = \operatorname{diag}(\lambda_1, \ldots, \lambda_n)$. So if A is not invertible then F_A is not Fredholm, and if A is invertible then all $F_i = \partial_s + \lambda_i : L^{1,2}(\mathbb{R}, \mathbb{R}) \to L^2(\mathbb{R}, \mathbb{R})$ are invertible and so is F_A. We have proved the following.

Theorem 3.2. *If $A_s = A$ is constant and A invertible, then F_A is invertible.*

3.2 The general case – Fredholmness.

If A_s is not constant we can not expect F_A to be invertible (after all its index should be the spectral flow, not 0). But if it is Fredholm it should be invertible "up to compact operators". That is F is Fredholm if and only if there us a $Q : L^2(\mathbb{R}, \mathbb{R}^n) \to L^{1,2}(\mathbb{R}, \mathbb{R}^n)$ such that $FP - I$ and $PF - I$ are compact. We want to "glue the inverses of F_{A_+} and F_{A_-}" to get the "almost-inverse" P.

To start with, we would like to perturb the "constant" operators F_{A_-} and F_{A_+} to agree with F_{A_s} at the ends while keeping them invertible.

Lemma 3.3. *Given a family of matrices A_s with $\lim_{s\to-\infty} A_s = A_-$, for any $\varepsilon > 0$ exists $S_1 > 0$ and family of matrices $A_1(s)$ and such that $A_1(s) = A_s$ for $s < -S_1$ and $||F_{A_1} - F_{A_-}|| < \varepsilon$ (using the operator norm).*

Proof. Take $S_1 > 1$ such that $|A_s - A_-| < \varepsilon$ for all $s < -S_1 + 1$ (using the operator norm on $n \times n$ matrices). Choose a smooth cut off function $\beta(s)$ equal to 0 for $s < -S$ and 1 for $s > -S_1 + 1$ and define $A_1(s) = A_s + \beta(s)(A_- - A_s)$, so that it agrees with A_s or $s < -S_1$ as wanted and agrees with A_- for $s > -S_1 + 1$.

Then for any u in $L^{1,2}(\mathbb{R}, \mathbb{R}^n)$ with $||u||_{L^{1,2}} = 1$ we have

$$\begin{aligned}
\left(||(F_{A_1} - F_{A_-})(u)||_{L^2}\right)^2 &= \int_{-\infty}^{+\infty} |(A_1 - A_-)u(s)|^2\, ds = \int_{-\infty}^{+\infty} |(1-\beta)(A_s - A_-)u(s)|^2\, ds \\
&= \int_{-\infty}^{-S_1+1} |(1-\beta)(A_s - A_-)u(s)|^2\, ds + 0 \\
&< \int_{-\infty}^{-S_1+1} \varepsilon^2 |u(s)|^2\, ds \leqslant \varepsilon^2 \int_{-\infty}^{+\infty} |u(s)|^2\, ds \leqslant \varepsilon^2
\end{aligned}$$

□

Now, since invertibility is open in the operator norm, we can find ε small enough that F_{A_1} constructed this way is invertible. Similarly, we can find a family of matrices $A_2(s)$ such that F_{A_2} is invertible and $A_2(s) = A_s$ for all $s > S_2$.

Moving on to the patching argument, let G_1 be the inverse of F_{A_1} and G_2 be the inverse on F_{A_2}. We take a partition of unity $\eta_1 + \eta_2 = 1$ subordinate to the cover $U_1 = (-\infty, 1), U_2 = (-1, +\infty)$ of $\mathbb{R}$ (recall this means that η_i is zero outside U_i). Take also smooth μ_1 to be 1 on U_1 and zero for $s > 2$ and μ_1 to be 1 on U_2 and zero for $s < -2$. Define $P_1 : L^2(\mathbb{R}, \mathbb{R}^n) \to L^{1,2}(\mathbb{R}, \mathbb{R}^n)$ by

$$P_1 v = \mu_1 G_1 \eta_1 v$$

and similarly for P_2. Finally let $P = P_1 + P_2$.

Theorem 3.4. *The operators $PF_{A_s} - I$ and $F_{A_s}P - I$ are compact, and hence F_{A_s} is Fredholm (this is, for example, [5, Theorem 16.26]).*

Proof. We will have several occasions to use the product rule $\partial_s(f(s)v) = f(s)(\partial_s v) + \dot{f}(s)v$, or written another way $f(s)(\partial_s v) = \partial_s(f(s)v) - \dot{f}(s)v$

We compute:

$$\begin{aligned} F_{A_s}P_1 v &= (\partial_s + A_s)\mu_1 G_1 \eta_1 v \\ &= \dot{\mu}_1 G_1 \eta_1 v + \mu_1 \partial_s(G_1\eta_1 v) + A_s \mu_1 G_1 \eta_1 v \\ &= \dot{\mu}_1 G_1 \eta_1 v + \mu_1(\partial_s + A_1)(G_1\eta_1 v) - \mu_1 A_1 G_1 \eta_1 v + \mu_1 A_s G_1 \eta_1 v \\ &= \mu_1 \eta_1 v + \dot{\mu}_1 G_1 \eta_1 v + \mu_1(A_s - A_1) G_1 \eta_1 v \\ &= \eta_1 v + \dot{\mu}_1 G_1 \eta_1 v + \mu_1(A_s - A_1) G_1 \eta_1 v \end{aligned}$$

Since $\dot{\mu}_1$ is supported on $[-2,-1]$, μ_1 is supported on $s \leqslant 1$ and $A_s - A_1$ is supported on $s \geqslant -S_1$. The difference $F_{A_s}P_1 - \eta_1$ is supported on $[-S_1 - 2, 1]$ and since $L^{1,2}(\mathbb{R}) \to L^2([-S_1 - 2, 1])$ is compact by the lemma below, this difference is compact as well.

Lemma 3.5. *[9, Lemma 1.1] The restriction map $L^{1,2}(\mathbb{R}, \mathbb{R}^n) \to L^2([A,B], \mathbb{R}^n)$ is compact (for any pair of real numbers $A < B$).*

Proof. Of course it is enough to do $n = 1$. We show $L^{1,2}(\mathbb{R}, \mathbb{R}) \subset C(\mathbb{R})$ (actually the proof shows $L^{1,2}(\mathbb{R}, \mathbb{R}) \subset C(\mathbb{R})$, as promised in the first section of these notes) and the restriction $L^{1,2}(\mathbb{R}, \mathbb{R}) \to C([A,B], \mathbb{R})$ is compact, which is enough since convergence in C^0 norm implies convergence in L^2.

By convolving with mollifiers one sees that $C^{1,2} = C \cap L^{1,2}$ is dense in $L^{1,2}$ (mollifiers actually give denseness of $C^\infty \cap L^{1,2}$ but we will only need this weaker version). For any f in $C^{1,2}$ we have

$$|f(x) - f(y)| \leqslant \int_x^y |f_k'(t)|\, dt \leqslant ||f'(x)||_{L^2}\sqrt{|x-y|} \leqslant ||f(x)||_{L^{1,2}}\sqrt{|s-t|} \tag{3.5}$$

Let f_n be a sequence in $C^{1,2}$ converging to f in $L^{1,2}$. We want to apply Arzela-Ascoli (on $[-S,S]$ for arbitrary S) to it to conclude that it has a subsequence that converges to a continuous $g \in C[A,B]$ uniformly on $[A,B]$. Then f_{n_k} converges to g in L^2 and hence $g = f$, so f is continuous.

To verify that Arzela-Ascoli applies we need to see that f_n is equicontinuous and uniformly bounded on $[-S,S]$. Since $||f_n||_{L^{1,2}}$ converge to $||f||_{L^{1,2}}$, the above equation 3.5 shows that f_n are equicontinuous (and actually uniformly equicontinuous).

We prove that f_n are uniformly bounded by contradiction. If $f_{n_k}(x_k)$ goes to infinity, then by Bolzano-Weirstrass we pass to x_{k_i} converging to x and by equicontinuity at x the sequence $f_{n_{k_i}}(x)$ diverges and so, by equicontinuity at x again, the norms $||f_{n_{k_i}}||_{L^2}$ diverge. This proves that f_n are uniformly bounded, allowing us to apply Arzela-Ascoli and finishing the proof that f is continuous on $[-S, S]$ and as S is arbitrary on all of $\mathbb{R}$ as well.

Now to prove compactness of the restriction, consider a sequence of f_n in $L^{1,2}$ with bounded $L^{1,2}$ norm. We want to show that their restriction to $[A, B]$ has uniformly convergent subsequence. Since f_n are continuous by what we just proved, we can apply the exact same argument (with convergence of f_n replaced by boundedness of $||f_n||_{L^{1,2}}$) to see that by Arzela-Ascoli applies and gives us what we want.

$\square$

Similarly $F_{A_s}P_2 - \eta_2$ is compact, so that $(F_{A_s}P_1 - \eta_1) + (F_{A_s}P_2 - \eta_2) = F_A P - I$ is compact, as wanted.

In the other direction we get

$$\begin{aligned} P_1 F_{A_s} w &= \mu_1 G_1 \eta_1 (\partial_s + A_s) w \\ &= \mu_1 G_1 (\partial_s(\eta_1 w) - \dot{\eta_1} v + \eta_1 A_s w) \\ &= \mu_1 G_1 (\partial_s + A_1)(\eta_1 w) + \eta_1 (A_s - A_1) w - \dot{\eta_1} w \\ &= \eta_1 w + \mu_1 G_1 (\eta_1 (A_s - A_1) w - \dot{\eta_1} w) \end{aligned}$$

As $\dot{\eta_1}$ and $\eta_1(A_s - A_1)$ are both supported on $[-S_1 - 1, 1]$ and the difference $P_1 F_{A_s} - \eta_1$ is compact. Similarly $P_2 F_{A_s} - \eta_2$ and hence $PF_{A_s} - I$ are compact.

$\square$

3.3 The diagonal case – Index.

We now turn to computing the index of F_{A_s}. To start, consider the case $n = 1$ and $A_s = \lambda(s)$. Then $F_{A_s} v = 0$ is the ODE $\partial_s v = -\lambda v$ which we can solve explicitly $\partial_s \ln(v) = -\lambda(s)$, $v(s) = A_s e^{\int_0^s -\lambda(x)\,dx}$. This is in $L^2(\mathbb{R}, \mathbb{R})$ if and only if $\lambda_+ > 0$ and $\lambda_- < 0$.

Since we know that F_{A_s} is Fredholm, the cokernel of F_{A_s} is the kernel of the adjoint of F_{A_s} which (by integration by parts) is just $-F_{-A_s}$, and that is one dimensional if $\lambda_+ < 0$ and $\lambda_- > 0$ and empty otherwise.

Putting this two together proves that index of F_{A_s} is given by the spectral flow $\operatorname{index} A_- - \operatorname{index} A_+$ in this case.

Now in higher dimensions, but still keeping $A_s = \operatorname{diag}(\lambda_1(s), \ldots, \lambda_n(s))$ we see that $F_{A_s}(v) = 0$ splits into n coordinate ODE's $\partial_s v_i = -\lambda_i v_i$, and so the

dimension of the kernel is the number of λ_i switching from negative to positive. Just as before, since F_{A_s} is Fredholm its cokernel is the kernel of $-F_{-A_s}$ and so has dimension equal to the number of λ_i switching from positive to negative. This proves the "index=spectral flow" for the case $A_s = \text{diag}(\lambda_1(s), \ldots, \lambda_n(s))$.

3.4 The general case – Index.

The general case now follows readily by homotopying it to the diagonal case. Namely we need two facts

1) Fredholm index is invariant under homotopy through Fredholm operators.
2) Any path $A_s$of symmetric matrices with A_- and A_+ invertible is homotopic through other such paths to the path $\tilde{A}_s = A_- + \lambda\beta(s)A_-$ for some $\lambda \in \mathbb{R}$ and $\beta(s)$ a function with $\beta = 0$ for $s < -1$ and $\beta(s) = 1$ for $s > 1$. In particular $\tilde{A}_s$ is diagonal in some coordinate system on the target $\mathbb{R}^n$ and so "index= spectral flow" holds for it.

The first one is a standard fact about Fredholm operators, see for example [5, Lemma 16.18]. For the second one we first observe that there is a $\lambda \in \mathbb{R}$ such that A_+ has the same index (as a symmetric matrix) as $A_- + \lambda$, since we just need to bump the right number of eigenvalues of A_- across 0.

Lemma 3.6. *Any two invertible symmetric matrices A and B of the same index are homotopic through invertible symmetric matrices.*

Proof. An invertible symmetric matrix A has an orthonormal basis in which it is diagonal (which can be chosen to give the canonical orientation on $\mathbb{R}^n$) and conversely a choice of such a basis together with n non-zero real numbers (to serve as eigenvalues in that basis) specifies A. Given A and B first rotate a basis of A to a basis of B keeping the λ's constant and in such a way that the negative eigenvalue eigenvectors of A are taken to negative eigenvalue eigenvectors of B (this is possible since indexes of A and B agree, and the group $SO(n)$ is path connected). Then simply rescale λs of A to λs of B linearly. Since the corresponding λs are of the same sign, they never become zero during this process. Considering the corresponding symmetric matrix for each "labeled basis" gives a path of invertible symmetric matrices from A to B.

□

Now we "pre-homotope" A_s to be constant A_- for $s < -S$ and A_+ for $s > S$ (for some S) and then just use the homotopy from $A_- + \lambda$ to A_+ given by the lemma above to homotope A_s to a path between A_- and $A_- + \lambda$. Finally ,we homotope that path to $\tilde{A}$ linearly.

The index stays fixed throughout and hence is equal to spectral flow from A_- to $A_- + \lambda$ which is the same as that from A_- to A_+, which completes the proof.

References

[1] A. Floer. The unregularised gradient flow of the symplectic action. Comm. Pure Appi. Math. 41 (1988) 775-813.

[2] M. W. Hirsch. Differential topology. Vol. 33. New York: Springer, 1976.

[3] E. Kerman, from Morse Homology to Floer Homolgy, informal lecture notes, http://www.math.uiuc.edu/ ekerman/mf-notes.ps.

[4] P. B. Kronheimer and T.S. Mrowka. Monopoles and three-manifolds. Vol. 10. Cambridge University Press, 2007.

[5] T. Mrowka, Lecture notes for Geometry of manifolds course, http://ocw.mit.edu/courses/mathematics/18-965-geometry-of-manifolds-fall-2004/lecture-notes/.

[6] D. McDuff, Dusa and D. Salamon. J-holomorphic curves and symplectic topology. Vol. 52. American Mathematical Society, 2011.

[7] J. Robbin and D. Salamon. The spectral flow and the Maslov index. Bulletin of the London Mathematical Society 27.1 (1995): 1-33.

[8] A. Ritter, Course notes on Morse homology, http://morsehomology.wikispaces.com/.

[9] E. P. van den Ban, Fredholm index and spectral flow, www.staff.science.uu.nl/ ban00101/lecnotes/flow.pdf.