Generalised Lagrangian Translations

Alex Epelde

September 2021

Contents

0 Introduction

In this write-up we will be looking at Lagrangian torus fibrations with singularities, in dimensions 2 and 3. The main idea is to, given Lagrangian sections L, L' of such a fibration, construct a fiber-preserving symplectomorphism of the total space from the difference $L - L'$. A concept of equivalence for Lagrangian sections helps to narrow down interesting symplectomorphisms. We also look at 'boundary conditions' for Lagrangian sections and the compactly supported symplectomorphisms this produces.

1 Symplectic Geometry Background

We go over some of the basic definitions in symplectic geometry which we will be using throughout. Refer to [3] and [7] for details.

Definition 1.1. A *differential k-form* ω on a (smooth) manifold M is the choice of, at each $x \in M$, a multilinear alternating form

$$\omega|_x : \underbrace{T_xM \times \cdots \times T_xM}_{k \text{ times}} \to \mathbb{R}$$

depending smoothly on x.

Remark. One way to make this question of smooth dependence precise is to demand that given smooth vector fields $V^{(1)}, \ldots, V^{(k)}$, the function

$$M \to \mathbb{R}, \qquad x \mapsto \omega|_x \left(V^{(1)}\big|_x, \ldots, V^{(k)}\big|_x \right)$$

be smooth.

Definition 1.2. A *symplectic manifold* is a pair (M, σ) where M is a manifold and the *symplectic form* $\sigma \in \Omega^2(M)$ is a differential 2-form which is

- closed, i.e. $d\sigma = 0$, and
- non-degenerate, i.e. at each $x \in M$, $\sigma|_x : T_xM \to T_x^*M$ is an isomorphism.

Example 1.3. Let $M = \mathbb{R}^{2n}$ with coordinates $q_1, \ldots q_n, p_1, \ldots p_n$ and set

$$\sigma_0 = \sum_{i=1}^{n} dq_i \wedge dp_i.$$

This is the *standard symplectic form on* $\mathbb{R}^{2n}$ and satisfies

$$\sigma_0 \left(\frac{\partial}{\partial q_i}, \frac{\partial}{\partial q_j} \right) = \sigma_0 \left(\frac{\partial}{\partial p_i}, \frac{\partial}{\partial p_j} \right) = 0,$$

$$\sigma_0 \left(\frac{\partial}{\partial q_i}, \frac{\partial}{\partial p_j} \right) = -\sigma_0 \left(\frac{\partial}{\partial p_j}, \frac{\partial}{\partial q_i} \right) = \delta_{ij}.$$

Example 1.4. Let $M = \mathbb{C}^n$ with coordinates $z_1, \ldots, z_n$ and set

$$\sigma_0 = \frac{i}{2} \sum_{j=1}^{n} dz_j \wedge d\bar{z}_j.$$

This is the *standard symplectic form on* $\mathbb{C}^n$.

Example 1.5. Let X be an n-dimensional manifold and $M = T^*X$ its cotangent bundle. Then there is a canonical choice of symplectic form σ on M, defined as follows.

Given a coordinate chart on X, say $q = (q_1, \ldots, q_n) : U \to \mathbb{R}^n$, we can write, for $x \in U$ and $\xi \in T_x^* M$,

$$\xi = \sum_{i=1}^{n} p_i dq_i|_x,$$

where the components p_i are uniquely determined by ξ.

This gives a coordinate chart on T^*X,

$$(q, p) = (q_1, \ldots q_n, p_1, \ldots p_n) : T^*U \to \mathbb{R}^{2n},$$

and we define σ on T^*U as

$$\sigma = \sum_{i=1}^{n} dq_i \wedge dp_i.$$

Different choices of chart (q, U) on X give the same σ at each point of M, so patch up to define σ globally.

Definition 1.6. Let (M, σ) be a symplectic manifold, say with $\dim M = 2n$. Then a *Lagrangian submanifold* $X \subset M$ is an n-dimensional submanifold of M such that for all $x \in X$ and $u, v \in T_x X \subset T_x M$, we have $\sigma(u, v) = 0$.

Remark. By linear algebra, for a symplectic vector space (V, σ), $\dim V$ is even and $\frac{1}{2} \dim V$ is the greatest possible dimension of a subspace W with $\sigma|_W = 0$.

Example 1.7. We continue with example 1.5. Given a 1-form α on X, we may view it as a section $\alpha : X \to T^*X$. Then the graph $\Gamma_\alpha \subset T^*X$ of α is an n-dimensional embedded submanifold of T^*X, and we can ask whether it is Lagrangian with respect to the canonical form σ.

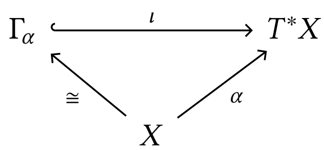

The map α induces a diffeomorphism between X and Γ_α. Thus Γ_α is Lagrangian, i.e. $\iota^*\sigma = 0$, if and only if $\alpha^*\sigma = 0$.

If (q, U) is a chart on X, and we write $\alpha = \sum_{i=1}^n \alpha_i dq_i$, then

$$\begin{aligned}
(\alpha^*\sigma)\left(\frac{\partial}{\partial q_i}, \frac{\partial}{\partial q_j}\right) &= \sigma\left(\alpha_*\left(\frac{\partial}{\partial q_i}\right),\ \alpha_*\left(\frac{\partial}{\partial q_j}\right)\right) \\
&= \sigma\left(\frac{\partial}{\partial q_i} + \sum_{k=1}^{n} \frac{\partial \alpha_k}{\partial q_i}\frac{\partial}{\partial p_k},\ \frac{\partial}{\partial q_j} + \sum_{l=1}^{n} \frac{\partial \alpha_l}{\partial q_j}\frac{\partial}{\partial p_l}\right) \\
&= \frac{\partial \alpha_i}{\partial q_j} - \frac{\partial \alpha_j}{\partial q_i}.
\end{aligned}$$

We see that Γ_α is Lagrangian precisely when $d\alpha = 0$, i.e. α is a closed 1-form.

Note that the converse also holds: if a section of T^*X is Lagrangian then it is the graph of a closed 1-form α.

Example 1.8. We continue still with the cotangent bundle T^*X. Let $\alpha : X \to T^*X$ be a 1-form, and consider the fiber-preserving diffeomorphism

$$\phi : T^*X \to T^*X, \qquad (x, \xi) \mapsto (x, \xi + \alpha|_x),$$

which is 'translation by α'.

When is ϕ a symplectomorphism, i.e. when does ϕ preserve σ?

As usual, let (q, U) be a chart on X and (q, p) the corresponding coordinates on T^*X. A straightforward calculation gives

$$(\phi^*\sigma)\left(\frac{\partial}{\partial q_i}, \frac{\partial}{\partial q_j}\right) = \frac{\partial \alpha_i}{\partial q_j} - \frac{\partial \alpha_j}{\partial q_i},$$

$$(\phi^*\sigma)\left(\frac{\partial}{\partial q_i}, \frac{\partial}{\partial p_j}\right) = \delta_{ij},$$

$$(\phi^*\sigma)\left(\frac{\partial}{\partial p_i}, \frac{\partial}{\partial p_j}\right) = 0.$$

So when α is closed, not only is Γ_α Lagrangian: translation by α preserves the whole symplectic structure of T^*X.

Definition 1.9. Let (M, σ) be a symplectic manifold and $H : M \to \mathbb{R}$ a smooth function. Non-degeneracy of σ implies the existence of a unique vector field X_H such that

$$\iota_{X_H}\sigma = \sigma(X_H, \cdot\,) = dH.$$

This X_H is the *Hamiltonian vector field* associated to the *Hamiltonian function H*.

Example 1.10. Let $M = \mathbb{R}^{2n}$ with the standard symplectic form. Then

$$X_{q_i} = -\frac{\partial}{\partial p_i}, \qquad X_{p_i} = \frac{\partial}{\partial q_i}.$$

Given the vector field X_H, integrating it gives the corresponding *Hamiltonian flow ρ_t^H*. An important result is

Proposition 1.11. *Hamiltonian flows preserve the symplectic form.*

Proof. Use Cartan's magic formula:

$$\begin{aligned} \mathcal{L}_{X_H}\sigma &= d\iota_{X_H}\sigma + \iota_{X_H}d\sigma \\ &= d(dH) + \iota_{X_H}0 \\ &= 0. \end{aligned}$$

□

Definition 1.12. Let (M, σ) be a symplectic manifold and $f, g : M \to \mathbb{R}$ smooth functions. The *Poisson bracket* of f and g is another function $M \to \mathbb{R}$, defined as

$$\{f, g\} = \mathcal{L}_{X_g} f = \sigma\left(X_f, X_g\right).$$

Proposition 1.13. *The Poisson bracket is bilinear and satisfies*

- $\{f, g\} = -\{g, f\}$, (Antisymmetry)
- $\{f, gh\} = \{f, g\}h + \{f, h\}g$, (Leibniz rule)
- $\{f, \{g, h\}\} + \{g, \{h, f\}\} + \{h, \{f, g\}\} = 0.$ (Jacobi identity)

Moreover $X_{\{f,g\}} = [X_f, X_g]$, *where* $[\cdot,\cdot]$ *is the usual Lie bracket of vector fields.*

Example 1.14. In $M = \mathbb{R}^{2n}$ with the standard symplectic form, we have the *canonical commutation relations*

$$\{q_i, q_j\} = \{p_i, p_j\} = 0, \qquad \{q_i, p_j\} = \delta_{ij}.$$

We conclude the section with

Theorem 1.15 (Darboux)**.** *Let* (M,σ) *be a symplectic manifold and* $x \in M$*. Then there is an open set* $U \ni x$ *and a coordinate chart* $(q,p) : U \longrightarrow \mathbb{R}^{2n}$ *such that* $\sigma = \sum_{i=1}^n dq_i \wedge dp_i$ *on* U*. We call* (q,p) *a* Darboux chart.

Darboux's theorem means that any statement of local character and invariant under symplectomorphisms which holds in $(\mathbb{R}^{2n}, \sigma_0)$ holds in any symplectic manifold (M,σ).

For example, given $x \in M$ and a Lagrangian subspace W of T_xM, we can locally find a Lagrangian submanifold X with $T_xX = W$.

A useful variant of Darboux's theorem is the following. Suppose we are given smooth functions $q_1,\dots,q_n$ in a neighbourhood of $x \in M$ and which Poisson commute. Then we can extend the q to a Darboux chart (q,p), in a possibly smaller neighbourhood of x.

2 Lagrangian Torus Fibrations

2.1 The Arnol'd-Liouville Theorem

Let $\pi : M \longrightarrow B$ be a submersion of the symplectic $2n$-manifold M onto the n-dimensional B, such that the fibers $F_b = \pi^{-1}(b)$ are compact, connected, Lagrangian submanifolds of M.

Remark. The preimage theorem forces the F_b to be n-dimensional submanifolds of M, so the Lagrangian condition is reasonable.

There is an action of T_b^*B on F_b via Hamiltonian flows as follows. Given $\xi \in T_b^*B$, let f be defined in a neighbourhood of b such that $df|_b = \xi$. Consider the Hamiltonian flow on M corresponding to $f \circ \pi$. For $x \in F_b$, the differential $d(f\circ\pi) = \pi^*(df)$ kills T_xF_b. Hence by the Lagrangian assumption, $X_\xi := X_{f\circ\pi}$ lies in T_xF_b.

So the Hamiltonian flow preserves F_b as desired.

Remark. Note that on F_b, $X_{f\circ\pi}$ depends on f only through $df|_b = \xi$.

The tangent space T_xF_b being Lagrangian also means the $f\circ\pi$ Poisson commute. Hence that the vector fields $X_{f\circ\pi}$ commute. Let ρ_t^ξ denote the time t flow of the vector field X_ξ corresponding to $\xi \in T_b^*B$. Then given $\xi_1, \xi_2 \in T_b^*B$, using that X_{ξ_1}, X_{ξ_2} commute, we can show that $\rho_t^{\xi_1+\xi_2} = \rho_t^{\xi_1}\rho_t^{\xi_2}$ for $t \in \mathbb{R}$.

So $\xi \mapsto \rho_1^\xi$ defines a Lie group action of T_b^*B on the fiber F_b.

Next note that for $x \in F_b$, $\xi \mapsto (X_\xi)\big|_x$ surjects onto T_xF_b. So by connectedness, the action on F_b is transitive, and the kernel of the action

$$P_b := \left\{\xi \in T_b^*B : \rho_1^\xi = \mathrm{id}\right\}$$

is a discrete subgroup of T_b^*B.

Finally, compactness of F_b implies that P_b is a lattice of maximal rank, and hence that F_b is diffeomorphic to an n-torus. Hence the name *Lagrangian torus fibration.* This is one half of

Theorem 2.1 (Arnol'd-Liouville)**.** *Let* $\pi : M \longrightarrow B$ *be as above. Then*

- *the fibers* F_b *are* n*-tori, and*
- *the period lattices* P_b *join up to form a Lagrangian submanifold of* T^*B*.*

Proof. See [1], theorem 1.1. □

The second part of the theorem means that given $T_0 \in P_b$ and a sufficiently small neighbourhood U of b, the piece of $P := \coprod_{b \in B} P_b$ through T_0 is the graph of a closed 1-form over U.

From here it is a small step to the more usual statement of the Arnol'd-Liouville theorem in terms of action-angle coordinates. Letting T_0 run through a basis of P_b, we obtain smooth functions $\chi_1, \ldots, \chi_n$ in a neighbourhood U of b such that $d\chi_1, \ldots, d\chi_n$ form a basis of P over each point of U.

Then $a_i := \chi_i \circ \pi$ are our action coordinates, which we may complete to a set of Darboux coordinates (a, α). The periodicity of the Hamiltonian flows for the a_i means the α_i are indeed angle variables, defined modulo 1.

Note that the set $\alpha = 0$ is a generic local Lagrangian section of $\pi : M \to B$. Other Lagrangian sections are obtained by translating by the graph of a closed 1-form over B.

2.2 Monodromy

In practice we will work with $M = T^*B/P$ where $P \subset T^*B$ is Lagrangian, $P \to B$ is a smooth covering, and for $b \in B$, $P|_b$ is a rank n lattice. In the notation of section 2.1, this corresponds to $M \to B$ carrying a global Lagrangian section, here the image of the zero section under the quotient $T^*B \to T^*B/P$.

As in example 1.8, P being Lagrangian means translation by local sections of P preserves the symplectic form on T^*B. So the symplectic form descends to the quotient T^*B/P.

We will be interested in the monodromy of the bundle $M \to B$, a further restriction on its being trivial (other than the existence of a global section).

Example 2.2. Let $B = \mathbb{R}^2 \setminus \{0\}$ with polar coordinates (r, θ), and $P \subset T^*B$ be spanned by

$$d(r\theta) = r d\theta + \theta dr, \qquad 2\pi dr.$$

Let $\gamma : I \to B$ be a loop in B based at $b = (1, 0)$, as in figure 1. Then there is a unique lift $\lambda : P_b \times I \to P$ making the following diagram commute.

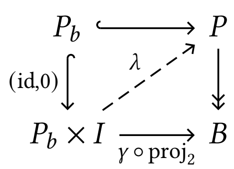

This sets up a group homomorphism $\phi_b : \pi_1(B; b) \to \mathrm{Aut}(P_b)$, where $[\gamma]$ gets sent to the induced map $P_b \to P_b$ at $t = 1$:

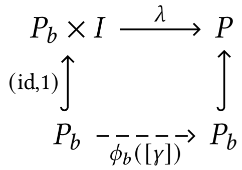

In this example, $\pi_1(B; b) \cong \mathbb{Z}$, and choosing suitable bases, a generator of π_1, e.g. the class of $\gamma(t) = (\cos t, \sin t)$, gets sent to

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \in \mathrm{GL}(2, \mathbb{Z}) = \mathrm{Aut}(P_b).$$

Indeed, as we circle the origin, $d(r\theta)$ jumps by $2\pi dr$.

There is a similar monodromy representation

$$\pi_1(B;b) \longrightarrow \mathrm{MCG}(F_b) = \mathrm{GL}(2,\mathbb{Z}),$$

obtained by considering instead the following diagram.

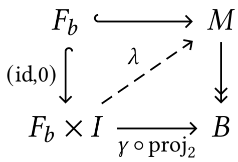

Note however that since $M \longrightarrow B$ is not a covering map, but rather a fiber bundle with fiber T^2, the lift λ is not uniquely determined. Still, once we pass to the mapping class group of F_b, the resulting $\phi_b([\gamma])$ *is* determined by $[\gamma]$.

This non-triviality of the monodromy prevents not only the bundle $M \longrightarrow B$ from being trivial, but also its extending over the origin in $\mathbb{R}^2$ while remaining an honest Lagrangian torus fibration.

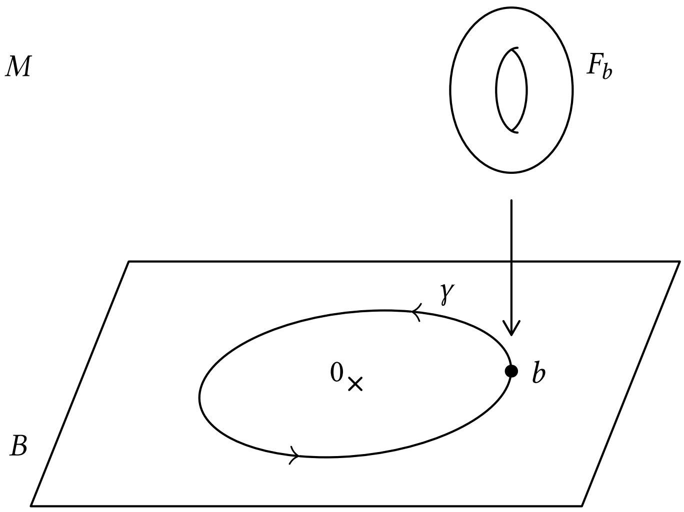

Figure 1: Example 2.2

3 Singularities for $n = 2$

It turns out that fibrations such as that in example 2.2 *can* be extended so long as we allow the presence of singularities.

We will look at *nodal singularities*, where $\pi : M \longrightarrow B$ fails to be a submersion at isolated points, but around which we can find Darboux coordinates (x, y) putting π in the form

$$(\pi_1, \pi_2)(x, y) = (x_1 y_1 + x_2 y_2, x_1 y_2 - x_2 y_1).$$

Introducing complex coordinates $x = x_1 + i x_2,\ y = y_1 + i y_2$, in which the symplectic form becomes $\sigma = \mathrm{Re}\,(d\bar{x} \wedge dy)$, this corresponds to

$$\pi(x, y) = \bar{x} y.$$

We demand that the fiber through such a singular point remains a Lagrangian submanifold away from it. Near the singular point, the fiber looks like the planes $x = 0,\ y = 0$ intersecting transversely in $\mathbb{C}^2$.

Example 3.1 (Neighbourhood of a nodal fiber)**.** We consider a particular disk bundle over an S^2, covered by the charts

$$U = \left\{(x, y) \in \mathbb{C}^2 : |xy| < \epsilon, |x| < 2, |y| < \tfrac{1}{4}\right\},$$
$$V = \left\{(u, v) \in \mathbb{C}^2 : |uv| < \epsilon, |u| < 2, |v| < \tfrac{1}{4}\right\},$$

glued together by identifying

$$\{(x, y) \in U : \tfrac{1}{2} < |x| < 2\} \quad \text{and} \quad \{(u, v) \in V : \tfrac{1}{2} < |u| < 2\}$$

using the symplectomorphism $\varphi_1(x, y) = (1/x, \bar{x}^2 y)$.

Here U, V carry the symplectic form $\sigma = \text{Re}\,(d\bar{x} \wedge dy) = -\text{Re}\,(d\bar{u} \wedge dv)$, and $0 < \epsilon < \frac{1}{8}$ is an otherwise arbitrary constant.

We further identify neighbourhoods of the origin in U, V, gluing

$$\left\{(x, y) \in U : |x|, |y| < \tfrac{1}{4}\right\} \quad \text{and} \quad \left\{(u, v) \in V : |u|, |v| < \tfrac{1}{4}\right\}$$

along the map $\varphi_2(x, y) = (\bar{v}, \bar{u})$, again preserving σ.

The resulting space is our M, and we let $\pi : M \to \mathbb{C}$ send

$$(x, y) \mapsto \bar{x}y, \qquad (u, v) \mapsto \bar{u}v.$$

The singular fiber $F_0 = \pi^{-1}(0)$ is an immersed S^2 intersecting itself transversely at $(x, y) = (u, v) = 0$. Topologically it is a T^2 with a $S^1 \times \{\text{pt}\}$ collapsed to a point. We call the $z \in H_1(F_b)$ which gets collapsed to a point in the singular fiber the *vanishing cycle.*

The cotangent space T_0^*B still acts on the smooth part of F_0, but the period lattice $P|_0$ has only rank 1. Indeed a quick calculation gives the Hamiltonian flows of π_1, π_2 as

$$\begin{aligned}
\rho_t^{\pi_1} : (x, y) &\mapsto (e^t x, e^{-t} y),\\
(u, v) &\mapsto (e^{-t} u, e^t v),\\
\rho_t^{\pi_2} : (x, y) &\mapsto (e^{it} x, e^{-it} y),\\
(u, v) &\mapsto (e^{-it} u, e^{it} v).
\end{aligned}$$

Letting (w, z) be coordinates on B, we have that $P|_0$ is spanned by $2\pi dz$.

What about $P|_{\delta e^{i\theta}}$ for δ small, positive? Starting at $(x, \delta e^{i\theta}/\bar{x}) \in U$ and flowing with π_1, we end up at

$$(e^{-t}/x, \delta e^t e^{i\theta} \bar{x}) \in V,$$

and eventually back at

$$(\delta e^t e^{-i\theta} x, e^{-t}/\bar{x}) \in U.$$

So we see that $P|_{\delta e^{i\theta}}$ is spanned by $(-\log\delta)dw + \theta dz$ and $2\pi dz$.

Notice how the lattice P has the same monodromy as in example 2.2.

Example 3.2 (A fiber with several singular points)**.** It is entirely possible for a single fiber to contain more than one singular point.

A model for the neighbourhood of such a fiber can be constructed much like in the previous example. Start with k copies of U, V each. Then glue U_j to V_j along a φ_1 and V_j to U_{j+1} along a φ_2, cycling the indices modulo k.

Defining $\pi : M \to \mathbb{C}$ in the obvious way, the fiber $F_0 = \pi^{-1}(0)$ ends up with k nodal singularities. It consists of a cycle of k Lagrangian spheres each intersecting the next transversely. See figure 2 for an attempt at representing such a fiber.

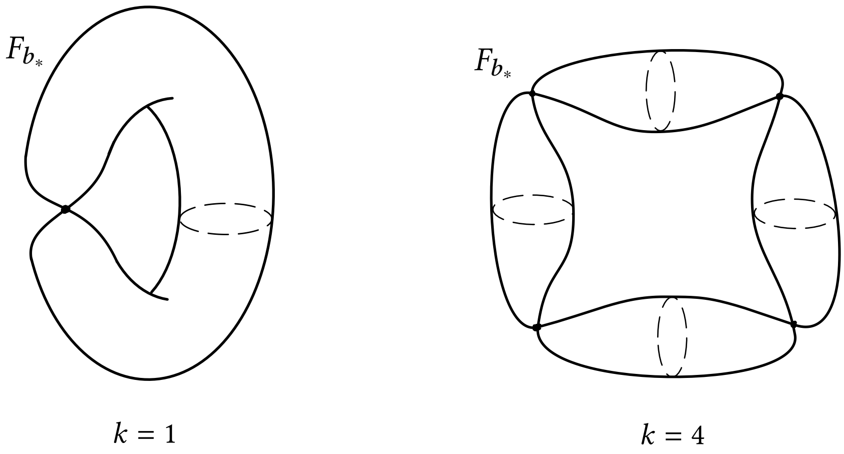

Figure 2: Fibers with k nodal singularities.

We note that the monodromy in P around such a fiber is represented by the matrix

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

the same as in going around k simple singular fibers with common vanishing cycle.

4 Algebraic Topology Background

Before moving on let us introduce some concepts from algebraic topology. The reader may check [4] for details and proofs.

4.1 Relative Homology

Recall that given a topological space X and $n \geq 0$, we define its group of *singular n-chains* $C_n(X)$, free abelian with basis the continuous maps from a standard n-simplex Δ^n into X.

One then defines *boundary homomorphisms*

$$\partial_n : C_n(X) \to C_{n-1}(X)$$

satisfying $\partial^2 = 0$ to obtain a chain complex of abelian groups. The homology groups of this complex are the *singular homology groups $H_n(X)$* of the space X.

Definition 4.1. Let X be a topological space and $A \subseteq X$. Then $C_n(A)$ is naturally a subgroup of $C_n(X)$. The quotient group

$$C_n(X, A) := C_n(X)/C_n(A)$$

is the *group of relative n-chains*. The boundary maps ∂_n send $C_n(A)$ to $C_{n-1}(A)$, and hence induce *relative boundary maps*

$$\partial_n : C_n(X, A) \to C_{n-1}(X, A).$$

The homology groups of the resulting chain complex are then known as the *relative homology groups $H_n(X, A)$*.

Proposition 4.2. *Let $A \subseteq X$ as above. Then there is a long exact sequence of abelian groups*

$$\cdots \longrightarrow H_n(A) \longrightarrow H_n(X) \longrightarrow H_n(X, A) \xrightarrow{\partial^*} H_{n-1}(A) \longrightarrow \cdots .$$

Here ∂^* is the usual boundary map: a relative cycle $z \in Z_n(X, A)$ is represented by an element of $C_n(X)$ whose boundary lies in $Z_{n-1}(A)$.

Example 4.3. Let A be the annulus $\{z \in \mathbb{C} : 1 \leq |z| \leq 2\}$. Then

$$H_n(A) = \begin{cases} \mathbb{Z} & \text{if } n = 0, 1, \\ 0 & \text{otherwise,} \end{cases} \qquad H_n(\partial A) = \begin{cases} \mathbb{Z}^2 & \text{if } n = 0, 1, \\ 0 & \text{otherwise,} \end{cases}$$

from which one may compute

$$H_n(A, \partial A) = \begin{cases} \mathbb{Z} & \text{if } n = 1, 2, \\ 0 & \text{otherwise.} \end{cases}$$

Suitable generators are shown in figure 3.

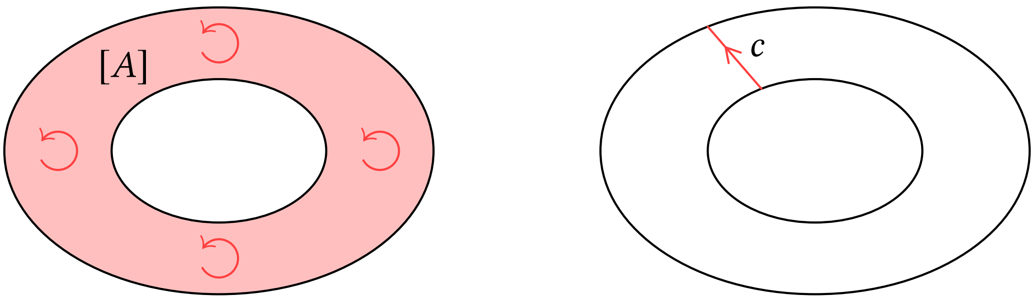

generator of $H_2(A, \partial A)$ generator of $H_1(A, \partial A)$

Figure 3: Relative homology of the annulus.

4.2 Cohomology and Poincaré Duality

In addition to the homology groups $H_n(X)$, associated to a space X are its *singular cohomology groups* $H^n(X)$. These are obtained by dualising at the chain level.

Definition 4.4. Let X be a topological space. Its group of *n-cochains* is

$$C^n(X) = \mathrm{Hom}(C_n(X), \mathbb{Z}).$$

Upon dualising, the boundary maps ∂_n change direction to give *coboundary maps*

$$\delta_n : C^n(X) \to C^{n+1}(X).$$

As usual $\delta^2 = 0$ and we define the n^{th} *singular cohomology group* as

$$H^n(X) := \ker \delta_n / \mathrm{im}\, \delta_{n-1}.$$

Note that in general $H^n(X)$ is *not* simply $\mathrm{Hom}(H_n(X), \mathbb{Z})$. There is however a natural surjection $H^n(X) \to \mathrm{Hom}(H_n(X), \mathbb{Z})$ whose kernel is given in terms of the Ext functor. In the case where $H_{n-1}(X)$ is free, this kernel vanishes.

Theorem 4.5 (Poincaré-Lefschetz). *Let M be a compact, orientable n-manifold, with fundamental class $[M] \in H_n(M, \partial M)$. Then for all k, the map*

$$D : H^k(M) \to H_{n-k}(M, \partial M), \qquad \alpha \mapsto [M] \frown \alpha$$

is an isomorphism.

The fundamental class $[M]$ is a generator of $H_n(M, \partial M) \cong \mathbb{Z}$, the choice of which amounts to a choice of orientation for M. The Poincaré isomorphism is in terms of cap product $\frown$.

Remark. The case where $\partial M = \emptyset$ is the usual Poincaré duality.

Example 4.6. Take the annulus A of example 4.3. Since all the $H_n(A)$ are free,

$$H^n(A) = \mathrm{Hom}(H_n(A), \mathbb{Z}).$$

The Poincaré isomorphism D sends a generator α of $H^1(A)$ to one of $H_1(A, \partial A)$, say the c of figure 3. We can then interpret the action of α on $H_1(A)$ as counting intersection with c.

4.3 Cohomology with twisted coefficients

Let X be a topological space and $x_0 \in X$ be a basepoint. If $\tilde{X}$ is the universal cover of X, and $\pi_1 = \pi_1(X; x_0)$, recall that there is a right action of π_1 on $\tilde{X}$ by deck transformations. This action turns the chain group $C_n(\tilde{X})$ into a right $\mathbb{Z}[\pi_1]$-module.

Definition 4.7. Let X be as above, and let M be a right $\mathbb{Z}[\pi_1]$-module. Then the groups

$$\mathrm{Hom}_{\mathbb{Z}[\pi_1]}\left(C_\bullet(\tilde{X}), M\right)$$

fit into a cochain complex whose cohomology groups $H^n(X; M)$ are the *cohomology groups of X with coefficients in M.*

In practice, X is often a CW complex and we can avoid working directly with the singular chain groups $C_\bullet(\tilde{X})$. Instead, we pull back the CW structure to $\tilde{X}$, where π_1 can then act on the cellular chain groups $C_\bullet^{\mathrm{cell}}(\tilde{X})$. The cohomology groups $H^n(X, M)$ may then be computed by considering the chain complex $\mathrm{Hom}_{\mathbb{Z}[\pi_1]}\left(C_\bullet^{\mathrm{cell}}(\tilde{X}), M\right)$.

Example 4.8. Let $F \to M \to B$ be a fiber bundle, and $b \in B$ some basepoint. Recall the monodromy representation $\pi_1(B; b) \to \mathrm{MCG}(F_b)$, c.f. example 2.2. In this way we obtain a right action of π_1 on $H_t(F)$, making it a right $\mathbb{Z}[\pi_1]$-module.

We can also make $H^t(F)$ into a right $\mathbb{Z}[\pi_1]$-module, but this needs some tweaking because of the contravariance of cohomology. We change direction by letting $[\gamma] \in \pi_1$ act on F_b as the inverse $[\overline{\gamma}]$.

Example 4.9. Let B be the annulus $\{z \in \mathbb{C} : 1 \leq |z| \leq 2\}$, and take a lattice P in T^*B with the same monodromy as in example 2.2, i.e. given by $\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$.

Let $M = T^*B/P$, and $\pi : M \to B$ be the natural projection to obtain a T^2 bundle.

Choose a suitable basis μ, λ for $H_1(F_b) \cong P|_b \cong \mathbb{Z}^2$ with respect to which the action of a generator α of $\pi_1(B)$ is by the matrix $\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$.

The base B deformation retracts onto the unit circle $\{|z| = 1\}$, which can be given a CW structure consisting of one cell in dimensions 0 and 1. Pulling back to its universal cover $\mathbb{R}$ gives a 0-cell at n and a 1-cell at $[n, n+1]$ for each $n \in \mathbb{Z}$.

An element

$$\varphi \in \mathrm{Hom}_{\mathbb{Z}[\pi_1]}\left(C_0^{\mathrm{cell}}(\tilde{B}), H_1(F)\right)$$

is determined by $\varphi(0) \in H_1(F)$ since we may send 0 to any other 0-cell using deck transformations. Similarly

$$\psi \in \mathrm{Hom}_{\mathbb{Z}[\pi_1]}\left(C_1^{\mathrm{cell}}(\tilde{B}), H_1(F)\right)$$

is determined by $\psi([0,1]) \in H_1(F)$.

As for the coboundary map δ, we compute

$$\begin{aligned}\delta\varphi([0,1]) &= \varphi(1) - \varphi(0) \\ &= (\alpha_* - 1)\varphi(0).\end{aligned}$$

So we arrive at

$$\begin{aligned}H^0(B, H_1(F)) &\cong \ker(\alpha_* - 1) \cong \mathbb{Z}, \\ H^1(B, H_1(F)) &\cong \operatorname{coker}(\alpha_* - 1) \cong \mathbb{Z}.\end{aligned}$$

5 Lagrangian Sections I

We come to the main idea of the project: to study the Lagrangian sections $L : B \to M$ of a Lagrangian torus fibration $\pi : M \to B$, possibly with singular fibers.

Given two Lagrangian sections L, L', and $b \in B$ such that F_b has no singularities, for $U \ni b$ a small enough neighbourhood we find that

$$L|_U = L'|_U + \xi,$$

where ξ is a closed 1-form over U, well-defined up to adding a section of P.

Translation by ξ gives a symplectomorphism taking $L|_U$ to $L'|_U$. This local translation by the difference $L' - L$ is defined globally, even if ξ isn't, and gives a fiber-preserving symplectomorphism $M \to M$ (at least away from singular fibers).

We will be interested in the case where there is no global 1-form ξ such that $L = L' + \xi$. Indeed, otherwise the resulting symplectomorphism, simply translation by ξ, is homotopic to the identity (through symplectomorphisms even): at time $t \in [0,1]$ translate by $t\xi$.

This motivates the classification of Lagrangian sections up to translation by global 1-forms.

5.1 The Regular Case

Suppose B is compact, connected and given an orientation, possibly with non-empty boundary. Then a section L gives an element of the relative homology group $H_n(M, \partial M)$.

As a connected and oriented (a symplectic manifold carries a natural orientation) manifold with boundary, M gets an intersection pairing via cap product

$$\frown : H_n(M, \partial M) \times H_n(M) \to \mathbb{Z}.$$

If $\gamma \in H_n(M)$ is the class of a fiber F_b, then L being a section (as opposed to a multisection) corresponds to $[L] \frown \gamma = 1$. Any other section L' satisfies

$$[L' - L] \in \ker(\cdot \frown \gamma : H_n(M, \partial M) \to \mathbb{Z}).$$

And if L', L differ by a global 1-form, then certainly $[L] = [L']$ in homology.

If $n = 2$ and there are no singularities this turns out to be the whole story.

Theorem 5.1 (Gross). *Let $L_0 : B \to M$ be a Lagrangian section of M, where $\pi : M \to B$ is without singularities and* $\dim B = 2$. *Then*

$$L \mapsto [L - L_0]$$

is a bijection between the Lagrangian sections of M, up to translation by global 1-forms, and

$$\ker(\cdot \frown \gamma : H_2(M, \partial M) \to \mathbb{Z}).$$

Proof. See proposition 6.69 of [6] for a proof, albeit presented in different terms. □

Example 5.2. Let the base B be the annulus $\{z \in \mathbb{C} : 1 \leq |z| \leq 2\}$, and take a lattice P in T^*B with the same monodromy as in example 2.2, i.e. given by $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$.

Let $M = T^*B/P$ and let $\pi : M \longrightarrow B$ be the natural projection.

Let λ, μ be a basis for $H_1(F_b)$ for which the monodromy is

$$\lambda \mapsto \lambda + \mu, \quad \mu \mapsto \mu.$$

One may then compute that $H_2(M, \partial M) \cong \mathbb{Z}^2$, generated by

- the class of the zero section, and
- the class of the cylinder $\lambda \times c$, where c is an interval in the base generating $H_1(B, \partial B)$.

Considering what capping against γ does to each of these, theorem 5.1 gives us a $\mathbb{Z}$ worth of Lagrangian sections over B. What are they?

We may understand them by considering

$$(\mathbb{R}^2/\mathbb{Z}^2) \times [0, 1]$$

with $(\mathbb{R}^2/\mathbb{Z}^2) \times \{1\}$ glued back to $(\mathbb{R}^2/\mathbb{Z}^2) \times \{0\}$ along $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$.

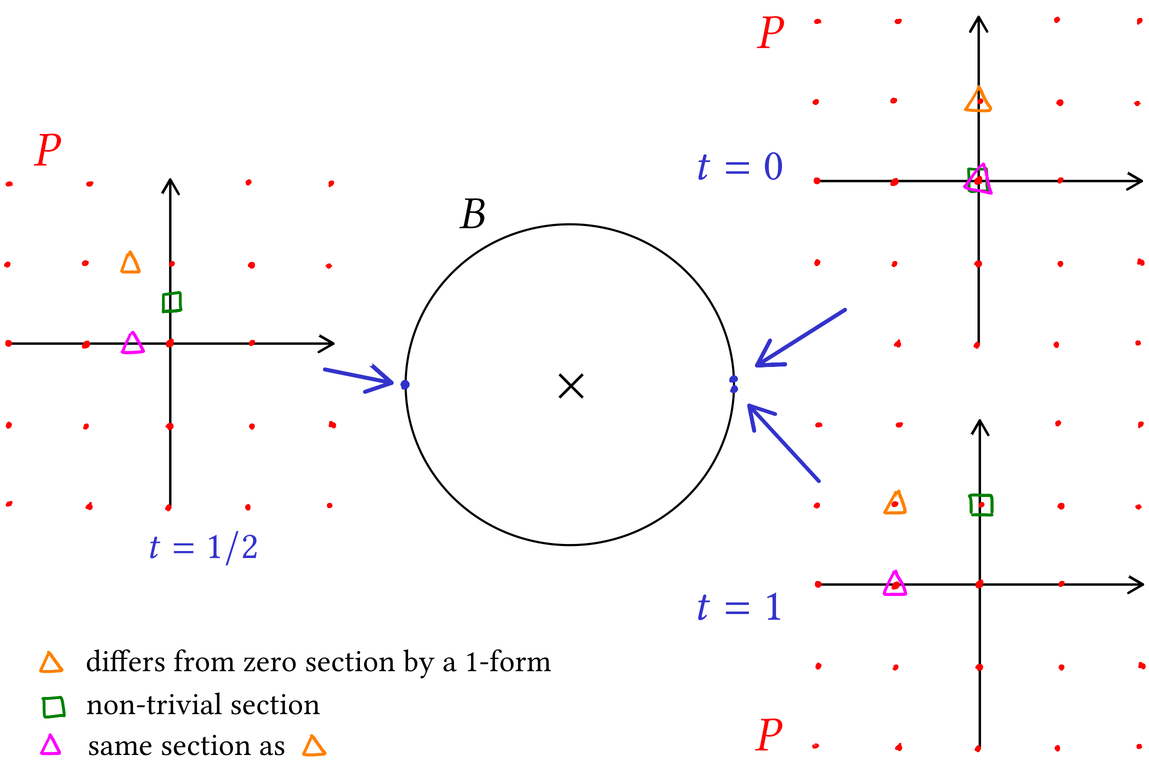

Figure 4: Some Lagrangian sections.

Figure 4 shows what some sections end up looking like. Note in particular how the section marked in orange is obtained from the zero section L_0 by adding a global 1-form. Marked in pink is the same section, translated by an element of P. Finally, in green we have a non-trivial section, L_1. The other Lagrangian sections may be obtained translating further by the difference $L_1 - L_0$.

Example 5.3. Consider the same B as in the previous example, but with monodromy given by $\begin{pmatrix} 1 & k \\ 0 & 1 \end{pmatrix}$, corresponding to a fiber with k singularities over $0 \in B$.

Then this time $H_2(M, \partial M) \cong \mathbb{Z}^2 \oplus \mathbb{Z}/k\mathbb{Z}$, with the $\mathbb{Z}^2$ factor generated by the zero section and $\lambda \times c$ as before, and the $\mathbb{Z}/k\mathbb{Z}$ factor generated by $\mu \times c$. Theorem 5.1 then gives a $\mathbb{Z} \oplus \mathbb{Z}/k\mathbb{Z}$ worth of sections over B.

5.2 Extending over singularities

How does the situation change if we allow $M \to B$ to have singularities in the $n = 2$ case? If, say, the fiber F_{b_*} has one or more nodal singularities, there are still Lagrangian sections over b_*, but we demand they go through the smooth part of F_{b_*}.

If F_{b_*} has just one singularity, and $U \ni b_*$ is a small enough neighbourhood, we may take any Lagrangian section over U to any other by adding the graph of a 1-form ξ. In particular if L, L' are global Lagrangian sections then the symplectomorphism given by the difference $L - L'$ extends without issue over b_*.

If instead F_{b_*} has $k > 1$ singularities, there are k different Lagrangian sections over b_*, depending on the piece of F_{b_*} they go through. If Lagrangian sections L, L' go through different pieces of F_{b_*} then it is not as easy to extend the symplectomorphism for $L - L'$ over b_*. However the extension, if it exists, is unique by continuity.

Suppose now that B contains finitely many points $b_*^{(1)}, \dots, b_*^{(m)}$ over which π has singularities. Deleting a small open disk D_i about each $b_*^{(i)}$ gives a base B^0 over which π is regular. Given a Lagrangian section L over B^0 we may then ask if L can be extended to the whole of B, perhaps after adding to it the graph of a 1-form over B^0.

Let $D_i' \supset D_i$ be a slightly larger open disk around $b_*^{(i)}$, and set $A_i = \overline{D_i'} \setminus D_i$. Then the Lagrangian sections over the annulus A_i form a copy of $\mathbb{Z} \oplus \mathbb{Z}/k\mathbb{Z}$, where k is the number of singularities over $b_*^{(i)}$ c.f. example 5.3. Those coming from sections over $\overline{D_i'}$ correspond to the $\mathbb{Z}/k\mathbb{Z}$ subgroup. So a necessary condition for L to extend over $b_*^{(i)}$ is the vanishing of $L|_{A_i}$ upon projection to the $\mathbb{Z}$ factor in $\mathbb{Z} \oplus \mathbb{Z}/k\mathbb{Z}$.

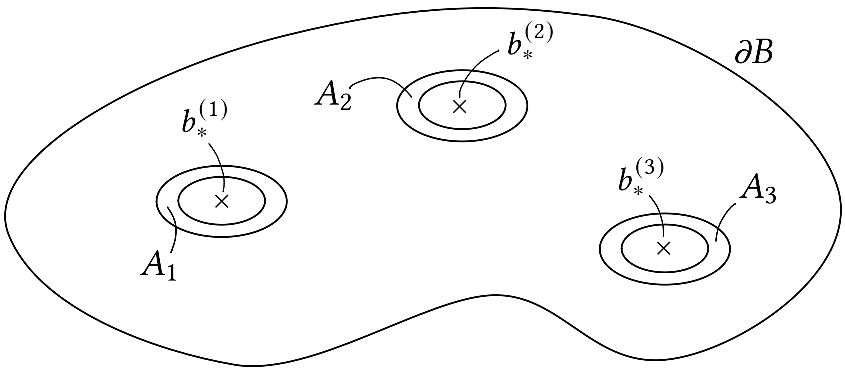

Figure 5: Extending over finitely many singular fibers.

Now suppose each $L|_{A_i}$ is suitable for extension over $b_*^{(i)}$. Then we can find sections L_i^{ref} over $\overline{D_i'}$ for $i = 1, \dots, m$ such that $L_i^{\mathrm{ref}}\big|_{A_i}$ and $L|_{A_i}$ differ by a (closed) 1-form ζ_i over A_i. For a suitable base B, e.g. a domain in $\mathbb{R}^2$, cohomology considerations mean we can find a closed 1-form ζ over B^0 whose restriction to each A_i is ζ_i. Adding ζ to L we can patch it up with the L_i^{ref} to obtain a section over the whole B.

We can do even better, and extend theorem 5.1 to this setting, at least when the fibers contain at most one singularity.

Theorem 5.4. *Let $M \to B$ be as above, i.e. with finitely many singular fibers and $n = 2$. Assume the fibers contains at most one singularity, and that B is a bounded domain in $\mathbb{R}^2$. If L_0 is a fixed Lagrangian section over B, then*

$$[L - L_0] \in \ker(\,\cdot \frown \gamma : H_2(M, \partial M) \to \mathbb{Z})$$

classifies Lagrangian sections L over B up to translation by global 1-forms.

Proof. We prove injectivity and surjectivity of $L \mapsto [L - L_0]$ separately, reducing in each instance to the regular case, where theorem 5.1 holds.

Let L, L' be Lagrangian sections over B such that $[L' - L] = 0$ in $H_2(M, \partial M)$. Let B^0 be as above, obtained from B by removing a small neighbourhood D_i of each singular point, and let $M^0 = M|_{B^0}$. Then $[L'|_{B^0} - L|_{B^0}] = 0$ in $H_2\big(M^0, \partial M^0\big)$ as well.

Note that by Poincaré-Lefschetz $H_2(M, \partial M)$ is naturally isomorphic to $H^2(M)$, and similarly for M^0, from which we get a natural restriction map

$$H_2(M, \partial M) \to H_2\big(M^0, \partial M^0\big)\,.$$

By theorem 5.1, there is a 1-form ξ over B^0 taking $L|_{B^0}$ to $L'|_{B^0}$.

Also, since we are assuming the singular fibers to be simple, over each D_i' we have a 1-form ζ_i taking $L|_{D_i'}$ to $L'|_{D_i'}$. Then over the annuli A_i we have two different 1-forms taking L to L', namely ζ_i and ξ. Then ζ_i and ξ differ by a section of the period lattice P over A_i. But any section of P over A_i extends to the whole of D_i, c.f. example 3.1.

So after adding a section of P to each ζ_i, we can patch them up with ξ to obtain a 1-form over the whole of B taking L to L', as desired.

Next let $z \in \ker(\,\cdot \frown \gamma : H_2(M, \partial M) \to \mathbb{Z})$. We want a Lagrangian section L over B such that $[L - L_0] = z$. We can restrict as above to B^0 and appeal to theorem 5.1, obtaining L over B^0 such that $\big[L - L_0|_{B^0}\big] = z|_{B^0}$.

We now note that the homology of M over the disk D_i' is such that

$$\ker\big(\,\cdot \frown \gamma : H_2\big(M|_{D_i'}, M|_{\partial D_i'}\big) \to \mathbb{Z}\big) = 0.$$

So $z|_{D_i'} = 0$, whence $z|_{A_i} = 0$ as well. In particular, again by theorem 5.1, $L|_{A_i}$ and $L_0|_{A_i}$ differ by a 1-form ζ_i over A_i. From the assumption on B, we can find a 1-form ζ over B^0 extending the ζ_i. Adding ζ to L allows us to patch up with L_0 and obtain a section over B.

It remains to confirm that $[L - L_0] = z$. Consider the Mayer-Vietoris sequence for

$$M = M^0 \cup \Big(\bigsqcup M|_{D_i'}\Big),$$

and in particular

$$\begin{aligned} \cdots \longrightarrow H^1(M^0) \oplus \Big(\bigoplus_i H^1\big(M|_{D_i'}\big)\Big) \longrightarrow \bigoplus_i H^1\big(M|_{A_i}\big) \longrightarrow \\ \longrightarrow H^2(M) \longrightarrow H^2(M^0) \oplus \Big(\bigoplus_i H^2\big(M|_{D_i'}\big)\Big) \longrightarrow \cdots\,. \end{aligned}$$

By construction $[L - L_0] = z$ when restricted both to M^0, and to each $M|_{D_i'}$. We'll be done if we can show that the connecting homomorphism

$$\bigoplus_i H^1\big(M|_{A_i}\big) \longrightarrow H^2(M)$$

is the zero map.

One may compute (e.g. using the Serre Spectral Sequence) that $H^1(M|_{A_i}) \cong \mathbb{Z}^2$, with one generator coming from $H^1(M|_{D_i'})$ and the other coming from $H^1(M^0)$. By exactness of the Mayer-Vietoris sequence we are done. □

5.3 Boundary Conditions for Lagrangian Sections

Suppose we are given a Lagrangian section L_0 over B, where $M \longrightarrow B$ is allowed singularities, and still $n = 2$. If L is another section agreeing with L_0 over a neighbourhood of ∂B, i.e. outside a compact subset of int B, then the symplectomorphism corresponding to $L - L_0$ is notable in having compact support.

It turns out such L can be classified fairly neatly in terms of their homology.

Theorem 5.5. *Let B be a closed disk in $\mathbb{R}^2$ and let L_0 be as above. Suppose further that the symplectic form σ on M is exact. Then*

$$L \mapsto [L - L_0]$$

sets up a bijection between the Lagrangian sections agreeing with L_0 outside of some compact set, up to translation by global 1-forms, and

$$\ker(\partial : H_2(M, \partial M) \longrightarrow H_1(\partial M)).$$

Proof. One direction is clear: if indeed L, L_0 agree over a neighbourhood of ∂B then certainly $\partial([L]) = \partial([L_0])$ in $H_1(\partial M)$.

Conversely suppose L is such that

$$[L - L_0] \in \ker(\partial : H_2(M, \partial M) \longrightarrow H_1(\partial M)).$$

The claim is that we can add a global 1-form to L so as to make it agree with L_0 over a neighbourhood of ∂B.

Let $A \subset B$ be a neighbourhood of ∂B which deformation retracts onto ∂B. Then

$$H_2(M|_A) \longrightarrow H_2(M|_A, M|_{\partial A})$$

is the zero map, whence

$$H_2(M|_A, M|_{\partial A}) \longrightarrow H_1(M|_{\partial A})$$

is injective. Note that $H_1(M|_{\partial A}) \cong H_1(M|_{\partial B})^2$ and that the image of this last map is contained in the diagonal copy of $H_1(M|_{\partial B})$.

All of this means that $\left[L|_A - L_0|_A\right] = 0$ in $H_2(M|_A, M|_{\partial A})$. So by theorem 5.1 we can find a 1-form ξ over A taking L_0 to L.

Now we use the assumption on the form of B and exactness of $\sigma = d\alpha$ to extend ξ to the whole of B. Consider the closed 2-chain in M given by L, L_0 and C, the graph of ξ over ∂B taking L_0 to L (see figure 6). Exactness of σ implies

$$\int_{L_0 + C - L} \sigma = 0\,.$$

Since L, L_0 are Lagrangian, the integral of σ over them vanishes and we are left with

$$\int_C \sigma = \int_{\partial B} \xi = 0\,.$$

This last equality follows noting that locally $\sigma = d\tau$ for τ the *tautological 1-form* on T^*B. The main property of τ is that if $\omega : B \longrightarrow T^*B$ is a 1-form over B, then $\omega^*\tau = \omega$.

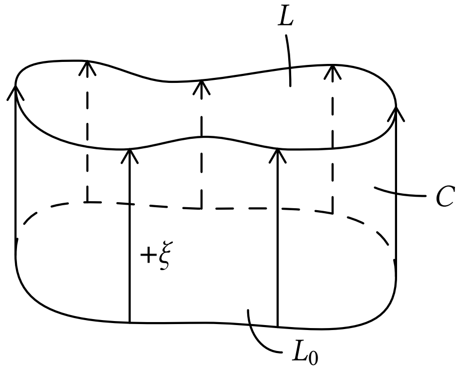

Figure 6: The 2-chain in the proof of theorem 5.5

Now, the integral of ξ over ∂B vanishing means we can extend ξ to the whole of B, whence $L - \xi$ can be defined over B and agrees with L_0 near ∂B as desired. □

6 Singularities for $n = 3$

Moving up in dimension allows for more involved singularities in $M \to B$. The main feature is that the discriminant locus $\Delta \subset B$ remains codimension 2.

Example 6.1 (Model for an edge)**.** Let $\pi_{(2)} : M_{(2)} \to \mathbb{R}^2$ be a Lagrangian 2-torus fibration with a simple singularity over the origin, c.f. example 3.1.

Then consider

$$\begin{aligned}\pi : M = M_{(2)} \times \mathbb{C}^* &\longrightarrow \mathbb{R}^3 \\ (p, z) &\longmapsto \big(\pi_{(2)}(p), \log|z|\big).\end{aligned}$$

This gives a Lagrangian 3-torus fibration with singular fibers over the x_3-axis, which is our discriminant locus Δ. The fiber over $(0, 0, b_3)$ is $\pi_{(2)}^{-1}(0)$, the degenerate 2-torus of example 3.1, times a circle in $\mathbb{C}^*$.

As for the period lattice P, it similarly splits as a product of that for $\pi_{(2)}$ with a copy of $\mathbb{Z}$ corresponding to the last component. Do note the flow of π_3 by itself is not globally periodic.

The monodromy in P, or equivalently in $H_1(F)$, in going around Δ, is given by

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

The combinatorial data attached to the edge includes $P|_\Delta$, corresponding to the subspace of P invariant under monodromy, and the span of the vanishing cycle $c \in H_1(F)$ by which there is a shift, here the first basis element.

In a general singular 3-torus fibration, $\Delta \subset B$ takes the form of a trivalent graph, where the structure near each edge is as above. We have two types of vertex according to the how the edges meeting at it are related.

Example 6.2 (Model for a positive vertex)**.** This extended example is 1.2 in [2].

Let $M = \mathbb{C}^3 \setminus \{1 + z_1 z_2 z_3 = 0\}$, with the canonical symplectic form.

Consider the map $\pi : M \to \mathbb{R}^3$ given by

$$\pi_1 = |z_1|^2 - |z_2|^2,$$
$$\pi_2 = |z_1|^2 - |z_3|^2,$$
$$\pi_3 = \log|1 + z_1 z_2 z_3|.$$

Then one may check that π is a submersion away from

$$Q_{12} = \{z_1 = z_2 = 0\}, \quad Q_{23} = \{z_2 = z_3 = 0\}, \quad Q_{13} = \{z_3 = z_1 = 0\},$$

and that the Poisson brackets $\{\pi_j, \pi_k\}$ all vanish.

To understand the structure of π, note that the flows of π_1, π_2 generate the T^2 action

$$(z_1, z_2, z_3) \mapsto \left(e^{i\theta_1} z_1, e^{i\theta_2} z_2, e^{i\theta_3} z_3\right)$$

with $\sum_i \theta_i = 0$, which is faithful away from $Q_{12} \cup Q_{23} \cup Q_{13}$. Moreover this action gives the part of the fiber of π through a given point and with a fixed value of $z_1 z_2 z_3$. The possible values of $z_1 z_2 z_3$ on a fiber are determined by π_3, sitting on a circle of radius e^{π_3} and centre -1.

Away from the discriminant locus

$$\Delta = \pi(Q_{12} \cup Q_{23} \cup Q_{13}) \subset B,$$

which consists of three edges meeting at the origin (see figure 7), π is a regular Lagrangian torus fibration.

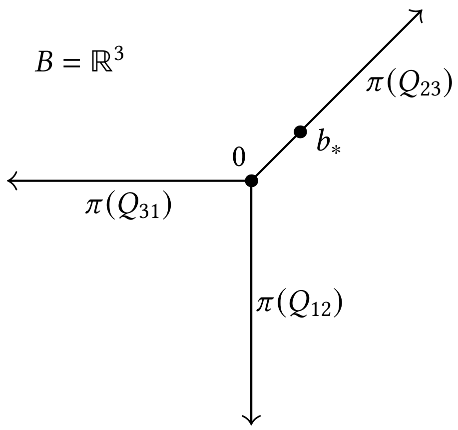

Figure 7: The discriminant locus Δ

The structure near each leg of Δ is as in the previous example, splitting as a product of a singularity one dimension down with an S^1 bundle. The fiber over the vertex 0 is even more singular, being a $T^2 \times S^1$ with $T^2 \times \{\text{pt}\}$ collapsed to a point.

The global T^2 action generated by π_1, π_2 helps not only understand the fibers of π, but also gives a rank 2 globally invariant sublattice of P, which must then correspond to the restriction of P to each leg of Δ.

As for the monodromy, the fundamental group of $B \backslash \Delta$ is generated by loops $\alpha_1, \alpha_2, \alpha_3$ about each of the legs of Δ, subject only to the relation $\alpha_1 \alpha_2 \alpha_3 = 1$. Fixing a basepoint $b \in B \setminus \Delta$ and choosing a suitable basis for $H_1(F_b) \cong P|_b$, the monodromy corresponding to each is

$$T_1 = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad T_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \quad T_3 = \begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \\ 0 & 0 & 1 \end{pmatrix}.$$

We have then that the behaviour near one of these *positive vertices* is that of a common invariant sublattice for each of the edges, with changing vanishing cycles.

Example 6.3 (Model for a negative vertex)**.** Our next example, construction 6.90 of [6], is only topological, as it turns out to be quite tricky to produce $\pi : M \to B$ with the desired monodromy while remaining smooth over the vertex.

Let $\Delta \subset B$ be as in the previous example, with legs l_1, l_2, l_3, and let $Y = B \times (\mathbb{R}^2/\mathbb{Z}^2)$. We construct a particular surface $S \subset Y$ sitting over Δ. Let C_1 be the circle $(1,0)\mathbb{R}/(1,0)\mathbb{Z}$ in $\mathbb{R}^2/\mathbb{Z}^2$, and let C_2 be the circle $(0,1)\mathbb{R}/(0,1)\mathbb{Z}$.

Over l_1, l_2, we take the cylinders $S_i = l_i \times C_i$. Note that this gives a figure eight over the vertex, the union of C_1, C_2 in $\mathbb{R}^2/\mathbb{Z}^2$. Next choose a suitably nice homotopy H between $C_1 - C_2$ and the circle $C_3 = (1,-1)\mathbb{R}/(1,-1)\mathbb{Z}$. Use this along with a parametrisation $\alpha : [0,1) \to B$ of l_3 with $\alpha(0) = 0$ to obtain

$$S_3 = \left\{(\alpha(t), H(s,t)) : t \in [0,1), s \in S^1\right\}.$$

Then $S = S_1 \cup S_2 \cup S_3$ is a pair-of-pants fibering over Δ, as shown in figure 8. It turns out there is a 'non-trivial' S^1 bundle $M' \to Y' = Y \setminus S$, compactifying to $M \to Y$ such that the fibers over S collapse to single points. Our singular 3-torus fibration is the composition with projection $M \to Y \to B$.

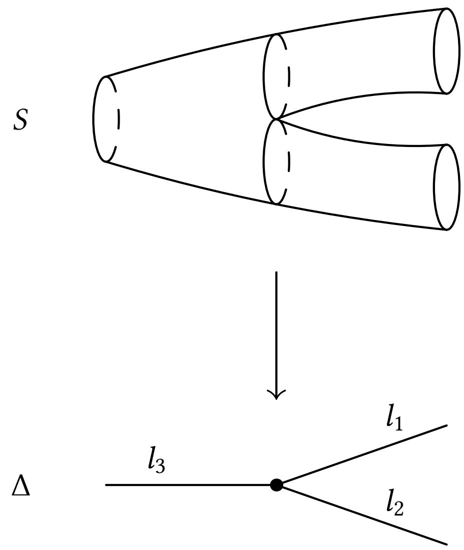

Figure 8: The surface $S \subset Y$ over which M is singular

The upshot of this construction is that the vanishing cycle over each edge of Δ is the same S^1, the one from the bundle $M' \to Y'$. The monodromy in $H_1(F)$ around the loops α_i of the previous example is this time given by

$$T_1 = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad T_2 = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad T_3 = \begin{pmatrix} 1 & -1 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Remark. There is a sense in which positive and negative vertices are 'dual' to each other, related to ideas in mirror symmetry. Suffice it to note that the monodromy representations for each are indeed dual to each other. By the way, the terms 'positive' and 'negative' refer to the Euler characteristic of the fiber over the vertex.

7 Lagrangian Sections II

7.1 The Sheaf Cohomology approach

In the $n > 2$ case, theorem 5.1 does not hold as stated, and we resort to sheaf cohomology in order to understand the Lagrangian sections of $M = T^*B/P$ (this is how theorem 5.1 is proven in the first place). If $\Lambda(T^*B/P)$ denotes the sheaf of Lagrangian sections $B \to T^*B/P$, then we have the short exact sequence of sheaves

$$0 \to P \to \Lambda(T^*B) \to \Lambda(T^*B/P) \to 0\,.$$

Taking cohomology gives the long exact sequence

$$0 \to H^0(B,P) \to H^0(B,\Lambda(T^*B)) \to H^0(B,\Lambda(T^*B/P)) \to H^1(B,P) \to$$
$$\to H^1(B,\Lambda(T^*B)) \to H^1(B,\Lambda(T^*B/P)) \to H^2(B,P) \to \cdots,$$

where H^0 is the space of global sections.

We are after the Lagrangian sections of T^*B/P, i.e. $H^0(B,\Lambda(T^*B/P))$, modulo those of T^*B, i.e. $H^0(B,\Lambda(T^*B))$. But this is precisely the image in $H^1(B,P)$ above. So a given $[g] \in H^1(B,P)$ determines an equivalence class of Lagrangian sections of M if and only if its image in

$$H^1(B,\Lambda(T^*B))$$

is zero.

Note that the short exact sequence

$$0 \to \mathbb{R} \to C^\infty(B) \xrightarrow{d} \Lambda(T^*B) \to 0$$

gives an isomorphism

$$H^1(B,\Lambda(T^*B)) \cong H^2(B,\mathbb{R}).$$

This last group is easier to work with, and we will often have B for which it is zero. In this case the Lagrangian sections of M up to translation by global 1-forms end up classified by $H^1(B,P)$, which we can compute using twisted coefficients.

Remark. A slightly stronger notion of equivalence of Lagrangian sections is to consider L, L' to be equivalent if they differ by the graph of an *exact* 1-form $d\varphi$, rather than just by the graph of a closed 1-form. In this context $[g] \in H_1(B,P)$ may determine several non-equivalent Lagrangian sections of M, which may be understood in terms of lifts to $H^1(B,\mathrm{Aff}_{\mathbb{Z}}(B,\mathbb{R}))$. See proposition 6.69 of [6] for details.

7.2 Sections near an edge

Suppose $M \to B$ is as in example 6.1, with singular locus Δ only one edge, and the monodromy in P given by

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

For compactness take B to be the closed unit ball in $\mathbb{R}^3$.

Like in the $n = 2$ case, any two Lagrangian sections L, L' defined over the whole of B can be taken to each other using a global 1-form ξ. Indeed, first do so over Δ, then extend ξ to all of B. Note ξ is defined up to a global section of P, and that these correspond to the $\mathbb{Z}^2$ sublattice of P invariant under monodromy.

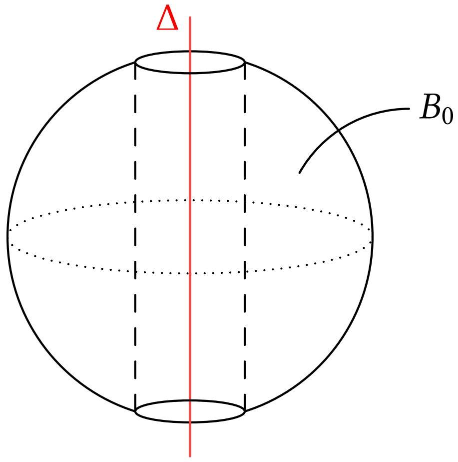

Figure 9: The base B^0 near an edge

Remove an open cylinder around Δ from B to obtain B^0, so that $M^0 = M|_{B^0}$ is a regular Lagrangian 3-torus bundle; see figure 9. What does sheaf cohomology imply about the sections of M^0?

Since $H^2(B^0, \mathbb{R}) = 0$, sections are classified by $H^1(B^0, P)$. We can compute this cohomology group as in example 4.9. Indeed, B^0 deformation retracts to a circle, and all that matters is the action of a generator α of $\pi_1(B)$. We get

$$H^1(B^0, P) \cong \operatorname{coker}(\alpha_* - 1) \cong \mathbb{Z}^2,$$

corresponding to a $\mathbb{Z}^2$ worth of different Lagrangian sections over B^0. It is worth comparing with example 5.2. Like in that situation, if the monodromy were trivial, we would have a full $\mathbb{Z}^3$ of sections. But the monodromy ends up killing the $\mathbb{Z}$ subgroup associated to the vanishing cycle.

Also similarly to in the $n = 2$ case, if L_0 is a fixed section over B, and L is a section over B^0, then L extends over Δ (possibly after translating by a 1-form over B^0) if and only if $[L - L_0|_{B^0}]$ is zero in this $\mathbb{Z}^2$.

7.3 Sections near a positive vertex

Next suppose $M \to B$ is as in example 6.2, with singular locus Δ three edges meeting at the origin, and the monodromy associated to a positive vertex. Again take B to be the closed unit ball.

It is still the case that any two sections L, L' defined over all of B may be taken to each other via a global 1-form ξ. Start by defining ξ over the vertex, then extend to the whole of Δ, then to the whole of B. As usual ξ is unique up to a global section of P. Or equivalently, up to an element of P over the vertex.

Next, as above, let us remove a tubular neighbourhood of Δ to obtain a base B^0 over which $M^0 = M|_{B^0}$ is a regular Lagrangian 3-torus fibration; see figure 10. When can Lagrangian sections over B^0 be extended to all of B? As was the case over a single edge, this depends on how many different Lagrangian sections there are over B^0.

Since $H^2(B^0, \mathbb{R}) = 0$, we need just compute $H^1(B^0, P)$. Note B^0 deformation retracts to the graph A, which in turn is homotopy equivalent to a wedge of two circles X.

If the loops β, γ generating $\pi_1(X)$ are as shown, the monodromy in P about each of them

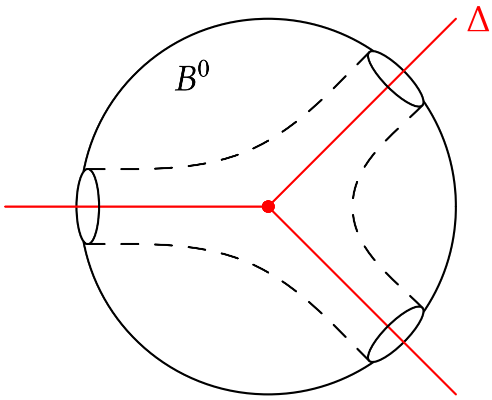

Figure 10: The base B^0 near a vertex

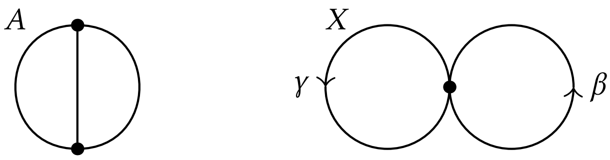

Figure 11: The base B^0 is homotopy equivalent to $X = S^1 \vee S^1$.

can be put in the form

$$\beta_* = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \gamma_* = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}.$$

If $\tilde{X}$ is the universal cover of X, then an element

$$\varphi \in \mathrm{Hom}_{\mathbb{Z}[\pi_1]}\left(C_0^{\mathrm{cell}}(\tilde{X}), P\right)$$

is determined by its value at a fixed vertex $\tilde{x}$, while

$$\psi \in \mathrm{Hom}_{\mathbb{Z}[\pi_1]}\left(C_1^{\mathrm{cell}}(\tilde{X}), P\right)$$

is determined by its values at two fixed 1-cells $\tilde{\beta}, \tilde{\gamma}$, projecting down to β, γ respectively.

The coboundary map satisfies

$$\delta\varphi(\tilde{\beta}) = (\beta_* - 1)\,\varphi(\tilde{x}),$$
$$\delta\varphi(\tilde{\gamma}) = (\gamma_* - 1)\,\varphi(\tilde{x}),$$

so that $H^1(B^0, P)$ is given by

$$\mathrm{coker}\left[(\beta_* - 1) \oplus (\gamma_* - 1)\right] \cong \mathbb{Z}^5.$$

One can try to understand what these sections look like, as we did in example 5.2. The situation is not that different; of the $\mathbb{Z}^6$ sections we would have if the monodromy were trivial, there is a $\mathbb{Z}$ subgroup which gets killed off.

7.4 Sections near a negative vertex

Finally, let us look at the situation near a negative vertex, with the base B and B^0 just as in the previous section.

To start off there is the problem that we have no model of $M \to B$ over the vertex which is smooth. We can somewhat get around this, given that we are mostly just interested in constructing symplectomorphisms $M \to M$.

The solution is to pretend that we do have a (smooth) Lagrangian section L_0 over the whole of B, and that we agree that any other section L agrees with L_0 in a neighbourhood of the vertex. This way given Lagrangian sections L, L', we can set the symplectomorphism corresponding to $L - L'$ as the identity over the vertex.

This setup also hard-codes the fact, familiar from above, that only one of the classes of sections over B^0 extends over Δ. So, what is $H^1(B^0, P)$ in this case?

Proceeding exactly as in the case of a positive vertex, this time the monodromy is

$$\beta_* = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \gamma_* = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

and we obtain

$$H^1(B^0, P) \cong \operatorname{coker}\left[(\beta_* - 1) \oplus (\gamma_* - 1)\right] \cong \mathbb{Z}^4.$$

7.5 Sections over a component of Δ

So far we have been looking at sections very locally. As a first step towards globalising, let us consider the case where B deformation retracts to Δ, and is connected. In general this will correspond to the sections over a tubular neighbourhood of a component of Δ.

Let L, L' be Lagrangian sections over such a B. For each vertex v of Δ, we can choose a 1-form ζ_v over a neighbourhood of v taking L to L'. Each ζ_v extends uniquely to each of the edges meeting at v. If we want a global 1-form taking L to L', the difficulty is in making sure that for each edge e, say joining the two vertices v_1 and v_2, the extensions ζ_{v_1}, ζ_{v_2} agree over e. Note that ζ_{v_1}, ζ_{v_2} at worst differ by an element of $P|_e$.

The freedom we have in choosing each ζ_v depends on $P|_v$. For positive vertices, this is a copy of $\mathbb{Z}^2$, and is canonically isomorphic to the P over each of the edges meeting at the vertex. On the other hand, for negative vertices, $P|_v$ has only rank 1, instead embedding into the P over the edges meeting at the vertex. Disregard any edges meeting only a single vertex. Choose an orientation for each of the other edges of Δ, and consider the map

$$\Phi : \bigoplus_{\text{vertices } v} P|_v \to \bigoplus_{\text{edges } e} P|_e,$$

where the term over e is the difference of those over the vertices at either end of e. Then $\ker \Phi$ corresponds to the global sections of P, while $\operatorname{im} \Phi$ tells us when we can modify our choices of ζ_v to obtain a global 1-form ζ taking L to L'. We can go further, and classify all Lagrangian sections over such a B.

Proposition 7.1. *Let B be as above, and let L_0 be a given Lagrangian section over B. Then there is a natural bijection between the equivalence classes of Lagrangian sections over B up to translation by 1-forms, and* $\operatorname{coker} \Phi$.

Proof. The definition of the map from Lagrangian sections to $\operatorname{coker} \Phi$ follows as above. Given L, for each vertex v choose ξ_v taking L_0 to L. Then on each edge e take the difference between the ξ_v coming from either vertex.

Injectivity is clear, so let us prove surjectivity. Work over a single edge e_*. Outside the very interior of e_*, we let L agree with L_0. Then given a section p of P over e_*, interpolate between the zero section at one vertex and p at the other, and add this 1-form to L_0. Note we can ensure that the 1-form by which we translate is closed, so that L is Lagrangian.

The result is that the image of L in $\bigoplus_e P|_e$ is zero except for a p in $P|_{e_*}$. This proves that we reach the whole of coker Φ and we are done. □

Example 7.2 (Interplay between negative vertices)**.** Consider the case where Δ is as in figure 12, with two negative vertices v_1, v_2 and one edge e between them.

Suppose that under an isomorphism $P|_e \cong \mathbb{Z}^2$, both $P|_{v_i}$ embed as $(1,0)\mathbb{Z}$. Then

$$\operatorname{im}\Phi = (1,0)\mathbb{Z} \subset P|_e\,,$$

so coker $\Phi \cong \mathbb{Z}$, giving a $\mathbb{Z}$ worth of non-equivalent Lagrangian sections over B.

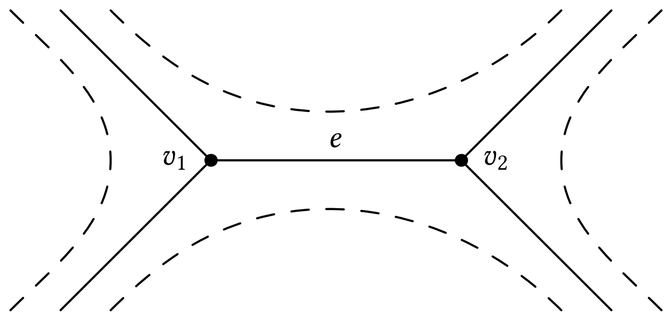

Figure 12: A simple example of what Δ might look like

Suppose instead that $P|_{v_1}$ embeds as $(1,0)\mathbb{Z}$, while P_{v_2} embeds as $(1,k)\mathbb{Z}$. Then in this case

$$\operatorname{im}\Phi = (1,0)\mathbb{Z} \oplus (0,k)\mathbb{Z} \subset P|_e\,,$$

so coker $\Phi \cong \mathbb{Z}/k\mathbb{Z}$ and we have a $\mathbb{Z}/k\mathbb{Z}$ worth of Lagrangian sections.

Remark. We see that negative vertices allow for variety in Lagrangian sections, thanks to the invariant sublattice of P in their neighbourhood having only rank 1.

The situation with positive vertices is less exciting; for example if above the vertices v_1, v_2 had been positive then all Lagrangian sections over B would have been equivalent. However, one may still hope to find several Lagrangian sections if, for example, Δ is not simply connected.

The case $\Delta = S^1$

The reader may have noticed there is a particular Δ which the above fails to address. This is where Δ consists of a single edge which loops to give an S^1.

Let $B \cong S^1 \times D^2$ be a solid torus around Δ , and let $B^0 \simeq T^2$ be B with a smaller tubular neighbourhood of Δ removed. As usual, the restriction $M^0 = M|_{B^0}$ is a regular Lagrangian 3-torus fibration.

Fix a basepoint $b \in B^0$ and generators α, β of $\pi_1(B^0; b) \cong \mathbb{Z}^2$ as shown in figure 13. Then we can find a basis of $P|_b$ with respect to which α acts by

$$\alpha_* = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Since α, β commute in $\pi_1(B^0; b)$, the monodromy of β preserves the rank 2 sublattice of $P|_b$ invariant under α. Indeed, this sublattice is just $P|_\Delta$, and β_* gives us the monodromy in $P|_\Delta$ as we move along Δ.

Now suppose L, L' are Lagrangian sections over B. If b_* is a given point of Δ, then we can find a 1-form ζ in a neighbourhood of b_* taking L to L', well-defined up to adding an element of $P|_{b_*}$. As we move along Δ then there is a unique way of extending ζ such that it still takes L to L'. Eventually we come back to b_*, and the question is whether ζ comes back to itself to give a global 1-form taking L to L'.

The situation is fairly similar to the above, in that we come back to a 1-form which differs from ζ by at worst an element of $P|_{b_*}$. The effect adding an element of $P|_{b_*}$ to ζ at the start has will depend on the monodromy β_*. We omit the details, but the end result is

Proposition 7.3. *There is a bijection between the Lagrangian sections over B up to translation by 1-forms, and* $\operatorname{coker}(\beta_* - 1)$.

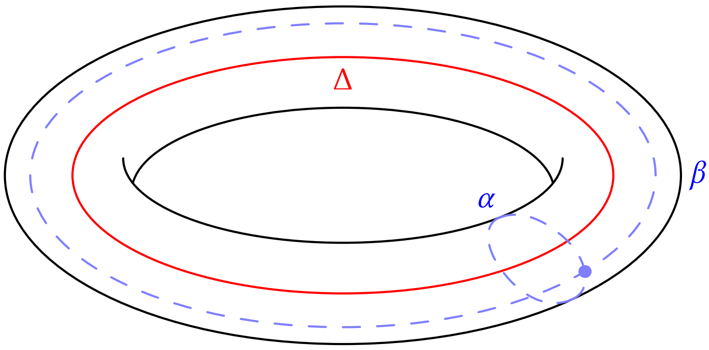

Figure 13: The loops α, β in B^0

7.6 Sections at large

We can now discuss the Lagrangian sections over a general base B and with arbitrary singular locus Δ. First split Δ into its connected components

$$\Delta = \bigsqcup_{i=1}^{m} \Delta_i \,,$$

and let U_i be a tubular neighbourhood of Δ_i c.f. the previous section. Let $V_i \subset U_i$ be a smaller tubular neighbourhood of Δ_i, and set

$$A_i = \overline{U_i} \setminus V_i \,,$$
$$B^0 = B \setminus \left(\bigcup V_i \right).$$

Remark. Note the similarity in our setup with what we had in section 5.2.

Our first result concerns the equivalence of Lagrangian sections. In a sense, a section L is determined by its restriction to B^0 and to each U_i .

Proposition 7.4. *Let L, L' be Lagrangian sections over B. Then L, L' differ by a global 1-form if and only if*

- *$L|_{B^0}, L'|_{B^0}$ differ by a 1-form over B^0, and*
- *for all i, $L|_{U_i}, L'|_{U_i}$ differ by a 1-form over U_i .*

Proof. One direction is immediate.

Now suppose ξ is the 1-form over B^0 taking L to L', and ζ_i is the one over U_i.

Then over A_i both ξ, ζ_i take L to L', and hence differ by a section of P over A_i. But sections of P over A_i correspond to elements of P invariant under monodromy, which are precisely those which extend over Δ.

So after adding a section of P over U_i to ζ_i, we make it agree with ξ and are able to extend ξ to a 1-form over all of B taking L to L'. □

The natural next step is an analogous existence result. Given Lagrangian sections L' over B^0 and $L^{(i)}$ over the U_i, can we find a section L over B whose restrictions to B^0, U_i differ from $L', L^{(i)}$ by a 1-form?

Proposition 7.5. *Let $L', L^{(i)}$ be as above. Then a necessary condition for the desired L to exist is that for each i, $L'|_{A_i}$ and $L^{(i)}|_{A_i}$ differ by a 1-form ζ_i over A_i. Moreover this is sufficient if we can choose the ζ_i such that the map δ in the Mayer-Vietoris sequence*

$$\cdots \longrightarrow H^1_{dR}\Big(\bigsqcup U_i\Big) \oplus H^1_{dR}(B^0) \longrightarrow H^1_{dR}\Big(\bigsqcup A_i\Big) \xrightarrow{\delta} H^2_{dR}(B) \longrightarrow \cdots$$

kills the 1-form ζ obtained by putting the ζ_i together.

Proof. That the first condition is necessary is clear. Indeed, $L'|_{A_i}, L^{(i)}|_{A_i}$ will both differ by a 1-form from $L|_{A_i}$.

Do note that ζ_i will not be unique if there are non-zero sections of P over A_i.

Now, if our choice of ζ_i is such that $\delta(\zeta) = 0$, then by exactness we can find 1-forms ξ over B^0 and ρ over $\bigsqcup U_i$ such that $\rho|_{A_i} - \xi|_{A_i} = \zeta_i$. So

$$L' + \Gamma_\xi \quad \text{and} \quad \bigsqcup L^{(i)} + \Gamma_\rho$$

agree over $\bigsqcup A_i$ and glue to give the desired Lagrangian section L. □

We can now more or less completely understand the different Lagrangian sections over B. The results of 7.5 on sections over a neighbourhood of Δ give us 'boundary conditions', which modulo cohomology restrictions c.f. the above proposition tell us when a Lagrangian section over B^0 extends over Δ.

7.7 Boundary conditions at infinity

Let L_0 be a Lagrangian section of $M \to B$ as above. Let us finish by looking at the Lagrangian sections L over B agreeing with L_0 outside a compact set. As discussed in section 5.3, the resulting Lagrangian translation by $L - L_0$ has compact support.

There is a sense in which the situation is simpler, in that we don't need to assume the symplectic form σ to be exact in proving the analogue of theorem 5.5, thanks to $H^1(S^2)$ being trivial.

Proposition 7.6. *Suppose $B \subset \mathbb{R}^3$ is a closed ball, and that L_0, L are Lagrangian sections over B. Then we can find a 1-form ζ such that $L - \zeta$ and L_0 agree in a neighbourhood of ∂B if and only if $L - L_0$ gets killed by the composition*

$$H^0(B, \Lambda(T^*B/P)) \to H^1(B, P) \to H^1(\partial B, P).$$

Proof. Considering a neighbourhood A of ∂B which deformation retracts onto ∂B, an argument similar to that in the proof of theorem 5.5 shows that L, L_0 differ by a closed 1-form ξ over A.

Since $H^1(\partial B) = 0$, ξ is exact. We can then find a 1-form ζ over B agreeing with ξ in a neighbourhood of ∂B.

So $L - \zeta$ and L_0 agree in a neighbourhood of ∂B and we are done. □

References

[1] J. J. Duistermaat. "On global action-angle coordinates". In: *Communications on Pure and Applied Mathematics* 33.6 (1980), pp. 687–706. ISSN: 0010-3640.

[2] Mark Gross. "Special Lagrangian fibrations. I. Topology". In: *Integrable systems and algebraic geometry (Kobe/Kyoto, 1997)*. World Sci. Publ., River Edge, NJ, 1998, pp. 156–193.

[3] Ana Cannas da Silva. *Lectures on symplectic geometry*. Vol. 1764. Lecture Notes in Mathematics. Springer-Verlag, Berlin, 2001. ISBN: 3-540-42195-5.

[4] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002. ISBN: 0-521-79160-X.

[5] Margaret Symington. "Four dimensions from two in symplectic topology". In: *Topology and geometry of manifolds (Athens, GA, 2001)*. Vol. 71. Proc. Sympos. Pure Math. Amer. Math. Soc., Providence, RI, 2003, pp. 153–208.

[6] Paul S. Aspinwall et al. *Dirichlet branes and mirror symmetry*. Vol. 4. Clay Mathematics Monographs. American Mathematical Society, Providence, RI; Clay Mathematics Institute, Cambridge, MA, 2009. ISBN: 978-0-8218-3848-8.

[7] John M. Lee. *Introduction to smooth manifolds*. Second. Vol. 218. Graduate Texts in Mathematics. Springer, New York, 2013. ISBN: 978-1-4419-9981-8.