

THE LAGRANGIAN FLUX-MONODROMY MORPHISM

JOÉ BRENDEL, JEAN-PHILIPPE CHASSÉ, AND RÉMI LECLERCQ

Happy birthday, Octav!

ABSTRACT. We define a new morphism on the fundamental group of the space of Lagrangian submanifolds, which takes into account the Lagrangian flux and monodromy. We study the discreteness of its image and relate this discreteness to the (C^∞) topology of the Hamiltonian orbit of the Lagrangian in question. We completely describe this new morphism for Lagrangian tori in $\mathbb{R}^4$ and $\mathbb{CP}^2$ and give explicit constructions. Similar constructions allow us to study the shape invariant of the Clifford torus in $\mathbb{CP}^n$. Finally, we upgrade our results on Lagrangian tori in $\mathbb{R}^4$ from the C^∞ topology to the Hausdorff one.

1. INTRODUCTION

The *flux morphism* of a symplectic manifold (X, ω) can be interpreted as measuring how much a symplectic isotopy fails to be Hamiltonian. It is defined as follows. Given a symplectic isotopy Ψ starting at the identity, one can associate to any loop γ in X the symplectic area of the cylinder swept out by γ under Ψ . Because ω is closed, this quantity only depends on the homotopy class of γ and the homotopy class of Ψ with fixed endpoints. Thus, this induces a map from the universal cover of the identity component of the symplectomorphism group of (X, ω) to the first de Rham cohomology group of M :

$$\text{Flux} : \widetilde{\text{Symp}}_0(X, \omega) \rightarrow H^1(X; \mathbb{R}).$$

This map is a morphism, since the two ends of a cylinder are homologous cycles in X . The following facts make the flux morphism particularly meaningful:

- (F1) Banyaga [Ban97] proved that its kernel is the universal cover of the Hamiltonian group of (X, ω) . Equivalently, a symplectic isotopy has vanishing flux if and only if it is homotopic relative to its endpoints to a Hamiltonian isotopy.
- (F2) Ono [Ono06] proved the deep fact that *the flux group* $\Gamma_X := \text{Flux}(\pi_1(\text{Symp}_0(X, \omega)))$ is discrete for any symplectic manifold.

The latter is known as the C^1 *flux conjecture* because of its important topological reformulation. Indeed, it is well known that Γ_X being discrete is equivalent to $\text{Ham}(X, \omega)$ being closed in $\text{Symp}_0(X, \omega)$ in the C^1 -topology. It is also equivalent to $\text{Ham}(X, \omega)$ being locally path connected in that topology.

There is a natural variant of the flux for Lagrangian submanifolds: In $\mathbb{R}^{2n}$, it was already considered by Gromov [Gro85], Viterbo [Vit90], and Sikorav [Sik91], among other people. The Lagrangian version of the C^1 flux conjecture was proved in special cases by Ono [Ono08]. The Lagrangian flux has garnered some attention — see, for example, [Sol13, EGM18, STV19, ACLS25] and §2.1 below for more details — but many of its fundamental properties are still somewhat ill-understood, especially when compared to how much is known about the symplectic flux. One of the striking differences is that *the Lagrangian flux is not a morphism*. The starting point of the present work is a natural upgrade of the Lagrangian flux map to an actual morphism and a study of the resulting morphism from various points of view. This study also illustrates that, as is often the case, the Lagrangian setting can be more subtle than the absolute case.

1.1. The flux-monodromy morphism. The Lagrangian variant of the flux is similar to the symplectic flux morphism. Fix a closed connected Lagrangian L in (X, ω) , and let $\Pi = \{L_t\}_{t \in [0,1]}$ be a Lagrangian isotopy starting at $L_0 = L$. To such a path, one can associate an element $\text{Flux} \Pi \in H^1(L; \mathbb{R})$, called the *Lagrangian flux* of Π , by associating to any 1-cycle γ of L the symplectic area swept out by γ under the isotopy.

The Lagrangian flux is invariant under endpoints-preserving homotopies of isotopies. Restricting our attention to loops, i.e. to the case $L_0 = L_1$, we obtain the *Lagrangian flux map*

$$\text{Flux}: \pi_1(\mathcal{L}, L) \rightarrow H^1(L; \mathbb{R}), \quad (1)$$

where $\mathcal{L}$ is the space of Lagrangians in X that are *Lagrangian isotopic* to L . This space is equipped with the C^1 -topology — see §2.2 for details. As for the symplectic flux, the Lagrangian flux of an isotopy can be interpreted as measuring how much it fails to be Hamiltonian: Π is induced by an ambient Hamiltonian isotopy of X if and only if the Lagrangian flux $\text{Flux}\{L_t\}_{t \in [0, \tau]}$ vanishes for all $\tau \in [0, 1]$.

Remark 1.1. In contrast with the case of symplectic flux (see Fact (F1) above), it is not true that a Lagrangian isotopy with vanishing *total* flux can be deformed into a Hamiltonian isotopy. Indeed, there is a Lagrangian isotopy from the Clifford torus in $(\mathbb{R}^4, \omega_0)$ to the Chekanov torus with vanishing total flux. However, these tori are not Hamiltonian isotopic [Che96]. On the other hand, we prove the analogue of Fact (F1) for Lagrangian *immersions* in Appendix B. •

At the outset of this paper, let us highlight another substantial difference between the Lagrangian and symplectic versions of the flux:

Observation 1.2. *The Lagrangian flux map (1) is not a morphism.*

This comes from the fact that, *seen as 1-cycles in L* , the end loops of a cylinder in M need not be homologous, as illustrated by the following example.

Example 1.3. Let X be the unit 2-sphere equipped with the standard area form. Denote by $S_h^1 = \{z = h\}$ the circle of constant height $h \in (-1, 1)$, and fix a Lagrangian $L = S_a^1$ for some $a \in (0, 1)$. Let $\Lambda \in \pi_1(\mathcal{L})$ be the loop obtained by concatenation of the Lagrangian deformation $t \mapsto S_{(1-2t)a}^1$ for $t \in [0, 1]$, which deforms S_a^1 to S_{-a}^1 , with a rotation of the sphere mapping S_{-a}^1 back to S_a^1 . For ξ a generator of $H_1(L)$, we obtain $\text{Flux } \Lambda \cdot \xi = 4\pi a$, the area of the cylinder bounded by S_a^1 and S_{-a}^1 . This is because the rotation is realized by a Hamiltonian isotopy and thus has vanishing flux. However, Λ maps ξ to $-\xi$, and hence

$$\text{Flux}(\Lambda * \Lambda) \cdot \xi = \text{Flux } \Lambda \cdot \xi + \text{Flux } \Lambda \cdot (-\xi) = 4\pi a - 4\pi a = 0. \quad \bullet$$




FIGURE 1. The Lagrangian loop Λ from Example 1.3.

This example suggests that the defect of the flux map (1) can be fixed by incorporating information about the action of Λ on L . To that end, we define the *monodromy* $\mathcal{M}(\Lambda) \in \text{MCG}(L) = \pi_0(\text{Diff}(L))$ of $\Lambda \in \pi_1(\mathcal{L})$ as the class of the self-map of L induced by Λ .

Definition 1.4. Let $L \subset X$ be a closed connected Lagrangian and $\mathcal{L}$ its connected component in the space of Lagrangians. We define the Lagrangian flux-monodromy morphism by

$$\begin{aligned} \text{FM} = \text{FM}_L: \pi_1(\mathcal{L}) &\rightarrow H^1(L; \mathbb{R}) \rtimes \text{MCG}(L) \\ \Lambda &\mapsto (\text{Flux } \Lambda, \mathcal{M}(\Lambda)), \end{aligned}$$

where $\text{MCG}(L)$ acts on $H^1(L; \mathbb{R})$ by pullback.

1.2. Discreteness of the flux-monodromy group. In contrast to the symplectic case, not only is the image of the Lagrangian flux map (1)

$$\text{F}(L) = \text{Flux}(\pi_1(\mathcal{L}, L)) \subset H^1(L; \mathbb{R})$$

not a subgroup of $H^1(L; \mathbb{R})$, but it also need not be discrete in general. Indeed, $\text{F}(L)$ is known to be non-discrete for many product tori in $\mathbb{R}^{2n}$ when $n \geq 3$ by [Che96]. Based on Chekanov's work, [ACLS25, Theorem 1] provides Lagrangian tori whose associated flux map accumulates at $0 \in H^1(L; \mathbb{R})$ in any symplectic manifold of dimension $2n \geq 6$. Examples in certain 4-manifolds can be constructed using, for example, [CS16, Theorem 1.5] or symmetric probes, as in [Bre25, BK25].

Definition 1.5. *The flux-monodromy group of L is defined as*

$$\Gamma(L) = \text{FM}_L(\pi_1(\mathcal{L})) \subset H^1(L; \mathbb{R}) \rtimes \text{MCG}(L). \quad (2)$$

As follows from Theorem 1.7 below, this group turns out to be discrete in all examples of the above paragraph. It is therefore natural to ask whether this is always true.

Question 1.6. *Is $\Gamma(L) \subset H^1(L; \mathbb{R}) \rtimes \text{MCG}(L)$ discrete?*

The answer to this question depends on the group of periods $\omega(H_2(X; \mathbb{Z}))$ and, more particularly, on two of its subgroups:

$$\begin{aligned} \mathcal{P}(X) &= \{ \langle \omega, S \rangle \in \mathbb{R} \mid S \in \ker(c_1 : H_2(X; \mathbb{Z}) \rightarrow \mathbb{Z}) \}, \text{ and} \\ \mathcal{P}_S(X) &= \{ \langle \omega, S \rangle \in \mathbb{R} \mid S \in \ker(c_1 \circ h_2 : \pi_2(X) \rightarrow \mathbb{Z}) \} \end{aligned} \quad (3)$$

with $h_2 : \pi_2(X) \rightarrow H_2(X; \mathbb{Z})$ being the second Hurewicz map. As usual, $c_1 \in H^2(X; \mathbb{Z})$ denotes the first Chern class of (TX, J) for any ω -compatible almost complex structure J .

Theorem 1.7. *If the flux-monodromy $\Gamma(L)$ is discrete for all Lagrangians $L \subset X$, then the group $\mathcal{P}_S(X)$ is discrete. Conversely, if $\mathcal{P}(X)$ is discrete, then $\Gamma(L)$ is discrete for all Lagrangians $L \subset X$.*

We suspect that the discreteness of $\Gamma(L)$ is actually equivalent to that of $\mathcal{P}(X)$, but we were unable to prove it. Note that the first claim in Theorem 1.7 relies on work by Chekanov–Schlenk [CS16].

Combining information on both the flux set $F(L)$ and the flux-monodromy group $\Gamma(L)$ of a given Lagrangian L yields interesting observations. Assume, for example, that 0 is not isolated in $F(L)$ but that $\Gamma(L)$ is discrete. This means that the smaller the flux of a loop in $\pi_1(\mathcal{L}, L)$ is, the more complicated is their monodromy. This happens, for example, in the case of product tori in $\mathbb{R}^{2n}$ for $n \geq 3$, as proved by Chekanov [Che96]. Here is a 4-dimensional example where this occurs.

Example 1.8. Let X_a be $S^2 \times S^2$ equipped with $(a\omega_{S^2}) \oplus \omega_{S^2}$, with $a \in [1, +\infty)$. Note that if $a \notin \mathbb{Q}$, the group of periods $\omega(H_2(X_a; \mathbb{Z}))$ of X_a is not discrete while $\mathcal{P}(X_a)$ is, since $\ker c_1 \subset H_2(X_a)$ has rank one. Therefore, Theorem 1.7 ensures discreteness of $\Gamma(L)$ for any Lagrangian L in X_a . When $a \neq 1$, it was shown in [BK25] that there are product tori $L \subset X_a$ for which there is a sequence of loops $\{\Lambda_i\} \subset \pi_1(\mathcal{L}, L)$ with $\text{Flux } \Lambda_i \rightarrow 0$. Since $\Gamma(L)$ is discrete, the corresponding sequence of monodromies $\{\mathcal{M}(\Lambda_i)\}$ is unbounded. •

1.3. Topological consequences. The relevant analogue of the Hamiltonian group in the Lagrangian setting is the orbit $\mathcal{L}\text{Ham}(L)$ of a Lagrangian L under the action of $\text{Ham}(X)$ on the space of Lagrangians $\mathcal{L}$, that is

$$\mathcal{L}\text{Ham}(L) := \{ \phi(L) \mid \phi \in \text{Ham}(X) \}.$$

We also consider the orbit $\mathcal{L}\text{Ham}(\iota)$ of a Lagrangian *embedding* $\iota : L \hookrightarrow X$ in the space $\tilde{\mathcal{L}}$ of Lagrangian embeddings $L \hookrightarrow X$ that are homotopic through Lagrangian embeddings to ι .

On the one hand, the flux set of L is connected to topological properties of $\mathcal{L}\text{Ham}(L)$ in $\mathcal{L}$. On the other hand, the discreteness of the flux-monodromy group gives information on the parametrized objects, as shown by the following two propositions.

Proposition 1.9. *If the flux-monodromy group $\Gamma(L)$ of L is discrete, then the Hamiltonian orbit $\mathcal{L}\text{Ham}(\iota)$ of ι is C^1 -locally path connected and C^1 -closed in $\tilde{\mathcal{L}}$.*

Conversely, suppose that the map $\iota_ : H_1(L; \mathbb{R}) \rightarrow H_1(X; \mathbb{R})$ induced by ι is injective. If $\mathcal{L}\text{Ham}(\iota)$ is either C^1 -locally path connected or C^1 -closed in $\tilde{\mathcal{L}}$, then $\Gamma(L)$ is discrete.*

Proposition 1.10. *If 0 is isolated in the flux set $F(L)$, then $\mathcal{L}\text{Ham}(L)$ is C^1 -locally path connected and C^1 -closed in $\mathcal{L}$.*

Conversely, suppose either that $\iota_ : H_1(L; \mathbb{R}) \rightarrow H_1(X; \mathbb{R})$ is injective or that L is a product torus in $\mathbb{C}^n$. If $\mathcal{L}\text{Ham}(L)$ is either C^1 -locally path connected or C^1 -closed in $\mathcal{L}$, then $F(L)$ is discrete.*

Remark 1.11. (1) The condition that $H_1(L; \mathbb{R}) \rightarrow H_1(X; \mathbb{R})$ be injective ensures that there cannot be any monodromy at the homological level. Therefore, in that case, the flux map is actually a morphism. In particular, 0 is isolated in $F(L)$ if and only if $F(L)$ is discrete.

(2) There is no hope to prove the converse statement of Proposition 1.10 without hypotheses on ι_* : some Lagrangians have dense flux set $F(L)$, but C^1 -locally path-connected Hamiltonian orbit; see Remark 5.15. •

We note that we can actually prove discreteness of the flux set in a few important cases, under the same condition on ι as above:

Proposition 1.12. *Let L be a Lagrangian such that the inclusion induces an injection $H_1(L; \mathbb{R}) \hookrightarrow H_1(X; \mathbb{R})$. Moreover, suppose that one of the following holds:*

- (a) L is weakly exact;
- (b) L is a non-displaceable torus;
- (c) X is H -monotone, i.e. monotone over $H_2(X; \mathbb{Z})$.

Then, $F(L)$ is discrete, and $\mathcal{L}\text{Ham}(L)$ is both C^1 -locally path connected and C^1 -closed in $\mathcal{L}$. Moreover, in the last case, $F(L) = 0$.

In particular, this means that, when X is H -monotone, its symplectic flux group Γ_X is equal to the image of the Lagrangian flux group of the diagonal $\Delta \subseteq X \times X$ under the isomorphism $\iota^* : H^1(\Delta; \mathbb{R}) \xrightarrow{\sim} H^1(X; \mathbb{R})$.

From Proposition 1.10, discreteness of $F(L)$ ensures that $\mathcal{L}\text{Ham}(L)$ is locally path connected and closed in $\mathcal{L}$ in the C^1 topology, for any Lagrangian torus L in $\mathbb{C}^2$. Using techniques inspired by [ACLS25], we can prove the much stronger following statement.

Theorem 1.13. *Let L be a Lagrangian torus in $\mathbb{C}^2$. Its Hamiltonian orbit $\mathcal{L}\text{Ham}(L)$ is locally path connected and closed in $\mathcal{L}$ in the Hausdorff topology.*

In fact, we prove a stronger form of local connectivity: *The path ensuring local path-connectedness can be realized by a Hamiltonian isotopy*; see Corollary 4.23 for a precise statement. This extends [ACLS25, Theorem 2] to not-necessarily H -rational tori in $\mathbb{C}^2$.

1.4. From monodromy to flux-monodromy. To determine the image $\Gamma(L)$ of the flux-monodromy morphism for a given Lagrangian $L \subset X$, it is often useful to understand the group of all possible monodromies first. For this discussion, it is convenient to consider the *homological Lagrangian monodromy* group

$$\mathcal{M}^{\mathbb{Z}}(L) = \{(\phi^\Lambda)_* \in \text{Aut}(H_1(L; \mathbb{Z})) \mid \Lambda \in \pi_1(\mathcal{L}, L)\},$$

and the *extended homological Lagrangian monodromy* group

$$\widehat{\mathcal{M}}^{\mathbb{Z}}(L) = \{\Phi^\Lambda \in \text{Aut}(H_2(X, L; \mathbb{Z})) \mid \Lambda \in \pi_1(\mathcal{L}, L)\},$$

where the automorphisms $(\phi^\Lambda)_*$ and Φ^Λ are the ones induced by the Lagrangian loop Λ — see §2.2 for details.

Question 1.14. *Can any $\Psi \in \text{Aut}(H_2(X, L; \mathbb{Z}))$ that*

- (1) *preserves the Maslov class, $\Psi^* m_L = m_L$, and*
- (2) *acts by the identity on the subspace $\text{im}(j: H_2(X) \rightarrow H_2(X, L))$ of ambient elements be realized as the extended monodromy $\Phi^\Lambda = \Psi$ of some $\Lambda \in \pi_1(\mathcal{L}, L)$?*

Those two conditions are obviously necessary. Similarly to [Che96, Proposition 5.1 (a)], which answers positively for product tori in $\mathbb{R}^{2n}$, we prove the following.

Theorem 1.15. *The answer to Question 1.14 is positive for the Clifford torus in $\mathbb{CP}^n$, and thus, for every Lagrangian torus that is Lagrangian isotopic to it.*

Remark 1.16. (1) If $\partial: H_2(X, L) \rightarrow H_1(L)$ is surjective, then $\widehat{\mathcal{M}}^{\mathbb{Z}}(L)$ determines the group of homological monodromies $\mathcal{M}^{\mathbb{Z}}(L)$.

(2) Furthermore, if X is (homologically) monotone, then it turns out that $\mathcal{M}^{\mathbb{Z}}(L)$ determines the flux-monodromy group $\Gamma(L)$; see Proposition 3.3.

(3) Every Lagrangian torus in $\mathbb{C}^2$ and $\mathbb{CP}^2$ is Lagrangian isotopic to the monotone product torus. This deep fact from [DRGI16] allows us to get the general description from the monotone case through a basepoint-change formula; see Proposition 2.4. •

When $n = 2$, we get a completely explicit description of the flux-monodromy group $\Gamma(L)$ of any Lagrangian torus L . In order to state these results, we fix an appropriate basis of $H_1(L)$ and its dual basis in cohomology.

Theorem 1.17. *For every Lagrangian torus L in $\mathbb{C}^2$, there exists $x \in \mathbb{R}_{\geq 0}$ such that the flux-monodromy group associated to L is given by*

$$\Gamma(L) = \left\{ \left(x \begin{pmatrix} (-1)^\delta - 1 & \\ & k \end{pmatrix}, \begin{pmatrix} (-1)^\delta & k \\ 0 & 1 \end{pmatrix} \right) \mid k \in \mathbb{Z}, \delta \in \{0, 1\} \right\}. \quad (4)$$

In particular, both the associated flux set $F(L)$ and flux-monodromy group $\Gamma(L)$ are discrete.

Moreover, there is an isomorphism $\pi_1(\mathcal{L}) \cong (\mathbb{Z} \rtimes \mathbb{Z}_2) \ltimes N$ for some normal subgroup $N \triangleleft \ker \text{FM}$.

By rescaling in $\mathbb{C}^2$, we can ensure that a Lagrangian loop is homotopic to one that is supported in any ball that contains the base Lagrangian L . Therefore, if L is any Lagrangian torus in a symplectic 4-manifold that is contained in a Darboux ball, then $\pi_1(\mathcal{L}, L)$ contains a subgroup isomorphic to $\mathbb{Z} \rtimes \mathbb{Z}_2$.

Remark 1.18. Equivalent descriptions of the monodromy group of Lagrangian tori in $\mathbb{C}^2$ already appeared in the works of Chekanov [Che96, Proposition 5.1 (a)] and of Yau [Yau12]. We give a more explicit description in terms of almost toric fibrations. This allows us to construct a section $s : \Gamma(L) \rightarrow \pi_1(\mathcal{L})$ to the flux-monodromy morphism from which the splitting of $\pi_1(\mathcal{L})$ follows, which is entirely new; see Theorem 4.8. •

In the case of *product tori*, that is,

$$T(a, b) = S^1(a) \times S^1(b) = \{\pi|z_1|^2 = a, \pi|z_2|^2 = b\},$$

for $a, b \in \mathbb{R}_{>0}$, the number $x \in \mathbb{R}$ appearing in (4) is simply $x = |a - b|$. Thus, loops based at the monotone product torus $L = T(a, a)$ have vanishing flux.¹ Moreover, the copy of $\mathbb{Z}_2$ in $\Gamma(L)$ is then generated by an ambient *Hamiltonian* isotopy whose time one map coincides with the coordinate swap $(z_1, z_2) \mapsto (z_2, z_1)$ on a compact subset. On the other hand, we show that no element of the copy of $\mathbb{Z}$ can ever be realized by an ambient Hamiltonian isotopy.

In order to state our result for $\mathbb{CP}^2$, we denote by $\rho : \text{GL}(2; \mathbb{Z}) \rightarrow \text{GL}(2; \mathbb{Z}/3\mathbb{Z})$ the reduction mod 3 and by $\mathbb{1}$ the vector $(1, 1)$.

Theorem 1.19. *The monodromy and flux-monodromy groups associated with a Lagrangian torus L in $\mathbb{CP}^2$ are isomorphic to the preimage of the stabilizer of $\mathbb{1}$ under ρ :*

$$H' := \rho^{-1}(\text{Stab}(\mathbb{1})) = \{h \in \text{GL}(2; \mathbb{Z}) \mid \mathbb{1}h = \mathbb{1} \pmod{3}\}. \quad (5)$$

It is a congruence subgroup of index 8 in $\text{GL}(2; \mathbb{Z})$. Moreover, it is generated by

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \begin{pmatrix} 2 & 1 \\ -1 & 0 \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}.$$

Notice that the modulo 3 condition comes from the fact that the Maslov class takes values in $6\mathbb{Z}$ when evaluated on $H_2(\mathbb{CP}^2; \mathbb{Z})$. Therefore, it yields a well-defined class $2 \cdot \mathbb{1}$ in the mod-6 reduction of $H^1(L)$.

As for the three generating matrices, the first two correspond to the obvious generators of $\mathbb{Z}_2 \rtimes \mathbb{Z}$, the flux-monodromy group of a Lagrangian torus in $\mathbb{C}^2$. The last matrix comes from a rotation of the moment polytope of $\mathbb{CP}^2$.

¹In Corollary 2.6 below, we prove that this is a general phenomenon for monotone Lagrangians when $H_1(X; \mathbb{R}) = 0$.

1.5. From flux-monodromy to flux. Knowing the homological flux-monodromy group $\Gamma^{\mathbb{Z}}(L) \subset H^1(L; \mathbb{R}) \rtimes \text{Aut } H_1(L; \mathbb{Z})$, we can understand the image $F(L) \subset H^1(L; \mathbb{R})$ of the flux map on $\pi_1(\mathcal{L}, L)$ by projecting to the H^1 -component. By the discussion in §1.4, we can thus understand $F(L)$ for any L which is Lagrangian isotopic to a product torus in $\mathbb{C}^n$ or to the Clifford torus in $\mathbb{CP}^n$ — though, again, these sets are difficult to write down explicitly. In $\mathbb{CP}^2$, we prove that the flux set $F(L)$ is dense in $H^1(L; \mathbb{R}^2) \cong \mathbb{R}^2$ whenever there is a Lagrangian isotopy from the Clifford torus T_{Cl} to L which has flux $(x, y) \in \mathbb{R}^2$ with $\frac{x}{y} \notin \mathbb{Q}$; see Proposition 5.14. This contrasts with the discrete flux set we got in Theorem 1.17 for all tori in $\mathbb{C}^2$.

This begs the question of what happens for *Lagrangian paths*, meaning Lagrangian isotopies which do not necessarily close up. The *shape* of L is defined as

$$\text{Sh}(L) = \{\text{Flux } \Pi \in H^1(L; \mathbb{R}) \mid \Pi_0 = L\}.$$

See §2.1 for a review of known facts about this invariant.

Our understanding of Lagrangian monodromy in $\mathbb{CP}^n$ allows us an almost complete description of the set $\text{Sh}(T_{\text{Cl}})$ for the Clifford torus. We concentrate here on the case $n = 2$; see §5.3 for more details on the case $n \geq 3$.




FIGURE 2. Illustration of the statement of Theorem 1.20. The simplex is the moment polytope $\Delta_{\mathbb{CP}^2} \subset \mathbb{R}^2 = H^1(T_{\text{Cl}}; \mathbb{R})$. The colours of the solid dots indicate their orbit type in the decomposition (6). The purple and pink dots and half-lines emanating from them are obstructed. The dashed green lines are the cases that remain unknown. The area shaded in grey can be realized as the flux of a Lagrangian isotopy.

We show that the action of the subgroup H' , defined by Equation (5) above, on the set $\mathcal{P} \subset H_1(T_{\text{Cl}}; \mathbb{Z}) = \mathbb{Z}^2$ of primitive elements has three orbits:

$$\mathcal{P} = ((1, 1) \cdot H') \sqcup ((-1, -1) \cdot H') \sqcup ((0, 1) \cdot H'). \quad (6)$$

Our result is the following.

Theorem 1.20. *Let $\eta \in H^1(T_{\text{Cl}}; \mathbb{R}) = \mathbb{R}^2$. If $\eta \notin \mathbb{R} \cdot \mathbb{Z}^2$, then $\eta \in \text{Sh}(T_{\text{Cl}})$. If $\eta \in \mathbb{R} \cdot \mathbb{Z}^2$, then write $\eta = r\lambda$ for $r \geq 0$ and $\lambda \in \mathcal{P}$. We have the following three cases:*

- (1) *if $\lambda \in (1, 1) \cdot H'$, then $\eta = r\lambda \in \text{Sh}(T_{\text{Cl}})$ if and only if $r < \frac{1}{2}$;*
- (2) *if $\lambda \in (-1, -1) \cdot H'$, then $\eta = r\lambda \in \text{Sh}(T_{\text{Cl}})$ if and only if $r < 1$;*
- (3) *if $\lambda \in (0, 1) \cdot H'$, then $\eta = r\lambda \in \text{Sh}(T_{\text{Cl}})$ if $r < 1$.*

The only case remaining open is thus $\eta = r\lambda$ with $\lambda \in (0, 1) \cdot H'$ and $r \geq 1$; see Question 6.6. The construction part of the argument, showing that $\mathbb{R}^n \setminus (\mathbb{R} \cdot \mathbb{Z}^n) \subset \text{Sh}(T_{\text{Cl}})$, is true for all n . The obstruction part is based on [CM18, Theorem 1.1], meaning that it works for general n , too. However, it is much trickier to understand the set of directions which is obstructed, mainly because we do not understand the orbit structure of the analogue of H' in higher dimensions.

ACKNOWLEDGEMENTS

This project was started during a visit of RL at the Forschungsinstitut für Mathematik at ETH Zürich and continued during visits of JB and JPC at the Laboratoire de Mathématiques d'Orsay. We thank both these institutions for the fantastic research environment they provide. We have benefited from discussions with Yves Benoist and Manfred Einsiedler about Theorems 1.19 and 1.20. We are especially grateful to Simon Machado for his patience in explaining the methods used in §5.2 to us. Furthermore, we thank Paul Biran, Emmy Murphy, and Felix Schlenk for various enlightening discussions about the content of this paper.

JB acknowledges the support of the Swiss National Science Foundation Ambizione Grant PZ00P2-223460. JPC acknowledges the partial support from the Swiss National Science Foundation (grant 200021_204107) and the Fondation Courtois.

CONTENT AND ORGANIZATION OF THE PAPER

Acknowledgements	9
2. Background	11
We review related literature in §2.1 and present basic facts about the flux-monodromy morphism in §2.2.	

3.	Discreteness and topological properties	15
	We first prove Theorem 1.7, which gives necessary and sufficient conditions to ensure discreteness of the flux-monodromy group in § 3.1. We then turn to the topological consequences of said discreteness, namely by proving Propositions 1.9 and 1.10. These are divided into two parts: local path connectedness in § 3.2 and closedness in § 3.3. Proposition 1.12, which ensures discreteness of the flux set in various settings, is proved at the end of § 3.2.	
4.	In $\mathbb{C}^2$: A case study	23
	We review preliminary material about Arnol'd–Liouville tori in § 4.1 and almost toric fibrations in § 4.2. We compute the flux-monodromy group of Lagrangians tori in $\mathbb{C}^2$ in § 4.3, and prove the existence of a section to the flux-monodromy morphism in § 4.4. Theorem 1.17 follows from these results. We also establish a complete description of the fundamental group of the space of <i>immersed</i> Lagrangians in § 4.5. In § 4.6, we study the Hamiltonian action on $\mathcal{L}(L)$, when L is a Lagrangian torus in several natural 4-manifolds. In § 4.7, we prove Proposition 1.12, which states a strong form of local path connectedness of the Hamiltonian orbit of any Lagrangian torus in $\mathbb{C}^2$ in the Hausdorff metric.	
5.	Monodromy and flux in $\mathbb{CP}^n$	38
	We compute the monodromy group of the Clifford torus in complex projective spaces. The case $n = 2$ is done in § 5.1, where we prove Theorem 1.19, while the case $n \geq 3$ is proven in § 5.2. Theorem 1.15 then follows directly from these results. The shape invariant is computed in § 5.3, when $n \geq 3$, and in § 5.4, when $n = 2$. This latter section is also where we prove Theorem 1.20.	
6.	Questions and outlook	53
	We close the paper with a series of open questions related to Lagrangians and their flux-monodromy morphism.	
	Appendix A. Paths of Hamiltonian isotopic Lagrangians	56
	We prove some basic results characterizing when Lagrangian paths are generated by Hamiltonian isotopies. These results are known to experts, but are scattered throughout the litterature or folkloric — we simply gather them into one easy-to-cite place.	
	Appendix B. Lagrangian homotopies with vanishing flux	57
	We study when Lagrangian isotopies with vanishing flux are homotopic to exact <i>homotopies</i> , an immersed variant of Hamiltonian isotopies. We give a complete answer for so-called H -rational immersions by developing an h -principle for such immersions.	
	References	61

2. BACKGROUND

2.1. Review of literature and relation to previous work.

Monodromy of Lagrangian submanifolds. The first result about monodromy of Lagrangian submanifolds that we are aware of appears in Chekanov's work [Che96]. More precisely, he answers Question 1.14 positively for product tori in $\mathbb{R}^{2n}$. See also the work by Yau [Yau09, Yau12] on the subject. We are not aware of any results about the possible monodromies of Lagrangian loops beyond these.

However, when we replace loops of Lagrangian isotopies by *Hamiltonian diffeomorphisms* mapping the Lagrangian to itself, there has been much more work. Chekanov [Che96, Theorem 4.5] gave a full description for product tori in $\mathbb{R}^{2n}$. Hu–Lalonde–Leclercq [HLL11] showed that Hamiltonian loops of closed weakly exact Lagrangians have trivial homological monodromy; this was recently extended to generalized homology theories by Porcelli [Por24, Por25]. Another recent work by Augustynowicz–Smith–Wornbard [ASW22] gave a full description of the (homological) Hamiltonian monodromy of many monotone Lagrangian tori. In [Bre25] and [BK25], the use of symmetric probes and new obstructions allowed for a full picture of this group for many (not-necessarily monotone) toric fibres in the 4-dimensional setting. For a full literature review, we recommend the excellent introduction of [ASW22].

Values of the Lagrangian flux. The question discussed in §1.5 of which possible values in $H^1(L; \mathbb{R})$ can be realized as the flux of a Lagrangian isotopy has a long history. For a Lagrangian isotopy ι_t in $(\mathbb{R}^{2n}, \omega_0 = d\lambda)$, the flux is equal to $[\iota_1^* \lambda] - [\iota_0^* \lambda] \in H^1(L; \mathbb{R})$. By Gromov [Gro85], there are no closed exact Lagrangians in $\mathbb{R}^{2n}$, meaning that $[\iota^* \lambda] \neq 0$ in that case. In fact, for Lagrangian product tori in $\mathbb{R}^{2n}$, it follows from work of Viterbo [Vit90, Theorem A] that $[\iota^* \lambda]$ is not a non-positive multiple of the Maslov class. Sikorav [Sik91, §1.5] conjectured that, in the standard basis, the class $[\iota^* \lambda]$ should be constrained to the positive orthant of $H^1(T^n; \mathbb{R}) = \mathbb{R}^n$. Chekanov proved that this is false; in fact, combining [Che96, Proposition 5.1 (b)] and Viterbo's result completely determines the *shape* $\text{Sh}(L)$ for any product torus $L = \iota(T^n)$, we get that

$$\text{Sh}(L) - [\iota^* \lambda] = H^1(T^n; \mathbb{R}) \setminus \mathbb{R}_{\leq 0} \cdot \mu_L.$$

Remarkably, it was found in more recent work by Entov–Ganor–Membrez [EGM18] and Shelukhin–Tonkonog–Vianna [STV19] that the so-called *star shape* is much smaller. The star shape is the subset of $H^1(L; \mathbb{R})$ that can be reached by Lagrangian isotopies moving along a straight line in $H^1(L; \mathbb{R})$. This salvages a version of Sikorav's above-mentioned conjecture: In [STV19, Theorem 6.8], it was proved that the star shape of a product torus in $\mathbb{R}^{2n}$ is indeed just the positive orthant.

Topology of the space of Lagrangians. In general, very few things are known about the topology of $\mathcal{L}(L)$, i.e. of the space of Lagrangians that are Lagrangian isotopic to a given Lagrangian L . The case of curves L in (finite type) surfaces Σ is, however, one notable exception to this. Indeed, one can compute the homotopy type of $\mathcal{L}(L)$ and $\tilde{\mathcal{L}}(L)$ using

a classical result of Gramain [Gra73, Theorem 4]. Alternatively, one can also use the fact that length shortening flow converges to either a point or a geodesic [Gra89] to construct a deformation retraction of these spaces. The table below gives all possible homotopy type, together with a description of the flux-monodromy group in these cases. In particular, $F(L)$ is always discrete in these cases.

Σ	L contractible?	$\tilde{\mathcal{L}}(L)$	$\mathcal{L}(L)$	$\Gamma(L)$
$\Sigma = S^2$	Yes	$\mathbb{RP}^3$	$\mathbb{RP}^2$	$\mathbb{Z}_2$
$\Sigma \neq S^2$	Yes	$S\Sigma$	Σ	0
$\Sigma = \mathbb{T}^2$	No	$\mathbb{T}^2$	S^1	$\mathbb{Z}$
$\Sigma \neq \mathbb{T}^2$	No	S^1	$\{*\}$	0

TABLE 1. Homotopy types of spaces of curves in finite type surfaces. Here, $S\Sigma$ is the unit tangent bundle of the surface Σ .

In higher dimensions, the picture is much less clear. Nonetheless, the homotopy type of $\tilde{\mathcal{L}}$ was computed by Coffey [Cof05] in two important 4-dimensional cases: the Lagrangian $\mathbb{RP}^2$ in $\mathbb{CP}^2$ and Lagrangian spheres S in $S^2 \times S^2$. In both cases, it is given by the rigid transformations of the ambient spaces: $\tilde{\mathcal{L}}(\mathbb{RP}^2)$ is homotopy equivalent to $\mathrm{PU}(3)$ and $\tilde{\mathcal{L}}(S)$ is homotopy equivalent to $\mathrm{SO}(3) \times \mathrm{SO}(3)$. This yields that $\mathcal{L}(\mathbb{RP}^2)$ is homotopy equivalent to $\mathrm{PU}(3)/\mathrm{PO}(3) = (\mathrm{SU}(3)/\mathbb{Z}_3)/\mathrm{SO}(3)$ and $\mathcal{L}(S)$ is homotopy equivalent to $\mathrm{SO}(3)$. In particular, $\mathcal{L}(\mathbb{RP}^2)$ and $\mathcal{L}(S)$ have respective fundamental groups $\mathbb{Z}_3$ and $\mathbb{Z}_2$. However, since all noncontractible loops in $\mathcal{L}(L)$ come from loops in $\mathrm{Ham}(X)$, these are not detected by FM.

We end with one last well-studied case: when L is the 0-section of T^*L . In this case, Arnol'd's nearby Lagrangian conjecture [Arn86] states that the Hamiltonian orbit of L should be the set of all exact Lagrangians, i.e. those Lagrangians L' such that restriction to L' of the tautological 1-form λ_0 is exact. This is only known to hold in S^1 [folklore], S^2 [Hin04], $\mathbb{RP}^2$ [HPW16], and $\mathbb{T}^2$ [DRGI16]. Moreover, in the first three cases, the authors show that the *strong* nearby Lagrangian conjecture holds: the space of exact Lagrangians is (weakly) contractible. This also determines the (weak) homotopy type of $\mathcal{L}(L)$.

Lemma 2.1. *Let L be the 0-section in T^*L . The space $\mathcal{L}(L)$ deformation retracts onto $\mathcal{L}\mathrm{Ham}(L)$. In particular, whenever the strong nearby Lagrangian conjecture holds, $\mathcal{L}(L)$ is (weakly) contractible.*

Proof. Fix a Riemannian metric on L . Note that, for any $L' \in \mathcal{L}(L)$, the canonical projection $\pi : T^*L \rightarrow L$ restricts to a homotopy equivalence $\pi|_{L'} : L' \rightarrow L$. Therefore, there is a unique harmonic 1-form $\sigma_{L'}$ on L such that $\pi|_{L'}^*[\sigma_{L'}] = [\lambda_0|_{L'}]$ in $H^1(L'; \mathbb{R})$. Let $\{\psi_{L'}^t\}$ be the flow on T^*L generated by $\pi^*\sigma_{L'}$. A direct computation shows that $\psi_{L'}^1(L')$ is exact.

Note that, by elliptic regularity, a C^∞ -continuous perturbation of L' gives a C^∞ -continuous perturbation of $\sigma_{L'}$, and thus of $\psi_{L'}^t$. Moreover, $\sigma_{L'} = 0$ — and thus $\psi_{L'}^t = id$ for all t — whenever L' is exact. Therefore, $[0, 1] \times \mathcal{L}(L) \rightarrow \mathcal{L}(L)$; $(t, L') \mapsto \psi_{L'}^t(L')$ gives a deformation retraction onto the subspace $\mathcal{L}^{ex}(L)$ of $\mathcal{L}(L)$ formed by exact Lagrangians.

But since $\mathcal{L}(L)$ is path connected, so is $\mathcal{L}^{ex}(L)$. However, a Lagrangian isotopy through exact Lagrangians is a Hamiltonian isotopy; this is well-known, but see Lemma A.2 for a self-contained proof. Therefore, $\mathcal{L}^{ex}(L) = \mathcal{L}\text{Ham}(L)$. $\square$

2.2. Basic facts on the flux-monodromy morphism. Let $L \subset X$ be a closed connected Lagrangian submanifold of a symplectic manifold (X, ω) . By $\mathcal{L} = \mathcal{L}(L)$ we will denote the space of (unparametrized) Lagrangians in X that are Lagrangian isotopic to L , which we see as a subspace of $C^\infty(L, X)/\text{Diff}(L)$. More precisely, we mean by this that L' is in $\mathcal{L}$ if there exists a smooth map $I: [0, 1] \times L \rightarrow X$ such that each $\iota_t = I(t, \cdot)$ is a Lagrangian embedding, and such that $\text{im } \iota_0 = L$, and $\text{im } \iota_1 = L'$. Such a map I is called a *Lagrangian isotopy*. Every Lagrangian isotopy can be extended to an ambient smooth isotopy $\{f_t\}_{t \in [0, 1]}$ of X , in the sense that $f_t \circ \iota_0 = \iota_t$. We call $\{f_t\}$ an *ambient Lagrangian isotopy*. For our purposes, it is often more convenient to forget about the parametrization and instead work with the corresponding path $\Pi = \{L_t\}_{t \in [0, 1]}$ in $\mathcal{L}$ of embedded Lagrangian submanifolds defined as the image $\text{im } \iota_t = L_t$. By abuse of terminology, we call both I and Π a Lagrangian isotopy. Sometimes, we refer to I as a *parametrization* of Π .

Given $\Pi = \{L_t\}_{t \in [0, 1]}$ with $L_0 = L$ and $L_1 = L'$, we can choose an ambient Lagrangian isotopy $\{f_t\}_{t \in [0, 1]}$, i.e. one that satisfies $f_t(L) = L_t$ for all $t \in [0, 1]$. This induces two maps, which we use throughout the paper:

$$\Phi^\Pi = (f_1)_*: H_2(X, L) \rightarrow H_2(X, L'), \quad \text{and} \quad \phi^\Pi = f_1|_L \in \text{Diff}(L, L').$$

One can check that the first map is well-defined and that the second one is well defined up to post-composing with an element in $\text{Diff}_0(L')$. In particular, the induced map $\phi_*^\Pi: H_1(L) \rightarrow H_1(L')$ is well-defined.

We equip the space $\mathcal{L}$ with its natural C^1 -topology unless otherwise stated; see [Ono08] or [ACLS25] for more details. Let $\pi_1(\mathcal{L}) = \pi_1(\mathcal{L}, L)$ be the fundamental group of $\mathcal{L}$ with basepoint L (which will often be suppressed in the notation). Before turning to the *flux-monodromy morphism*, we define and discuss the *monodromy morphism* and the *flux map*. Let $\text{MCG}(L) = \text{Diff}(L)/\text{Diff}_0(L) = \pi_0(\text{Diff}(L))$ be the smooth mapping class group.

Definition 2.2. *The monodromy morphism is defined by*

$$\mathcal{M}: \pi_1(\mathcal{L}) \rightarrow \text{MCG}(L), \quad \mathcal{M}(\Lambda) = [\phi^\Lambda]. \quad (7)$$

This is well-defined. Indeed,

- (1) the map ϕ^Λ is well-defined up to $\text{Diff}_0(L)$;
- (2) since the mapping class group is discrete, the element $[\phi^\Lambda]$ does not depend on the choice of loop representing the element $\Lambda \in \pi_1(\mathcal{L})$.

It is easy to check that $\mathcal{M}$ is indeed a morphism with respect to the product $(\phi, \psi) \mapsto \psi\phi$ on $\text{MCG}(L)$.

Let $\Pi = \{L_t\}_{t \in [0,1]}$ be a smooth path (which is not necessarily a loop) of Lagrangians starting at $L_0 = L$. Pick an ambient isotopy $\{f_t\}$ realizing Π . For any $\xi \in H_1(L)$, let $c : S^1 \rightarrow L$ be a loop satisfying $[c] = \xi$. Let $C_\xi^\Pi \in H_2(X, L_0 \cup L_1)$ be the relative class defined by the map $(s, t) \mapsto f_t(c(s))$. Geometrically, this is the cylinder swept out by the loop c under the Lagrangian isotopy Π . The map $H_1(L) \mapsto H_2(X, L_0 \cup L_1)$ defined by $\xi \mapsto C_\xi^\Pi$ is linear. Furthermore, it satisfies $\partial C_\xi^\Pi = \phi_*^\Pi \xi - \xi$. The equation

$$\text{Flux } \Pi \cdot \xi = \langle [\omega], C_\xi^\Pi \rangle \quad (8)$$

defines the *Lagrangian flux* $\text{Flux } \Pi \in H^1(L; \mathbb{R})$ of Π . This definition is independent of all choices we have made and is furthermore invariant under endpoint-preserving homotopies in the space of Lagrangians. Therefore, if we take $\Lambda = \{L_t\}_{t \in [0,1]}$ to be a loop in $\mathcal{L}$, then the Flux map descends to a map on $\pi_1(\mathcal{L})$.

Definition 2.3. *Let $L \subset X$ be a closed, connected Lagrangian submanifold. The flux map $\text{Flux} : \pi_1(\mathcal{L}, L) \rightarrow H^1(L; \mathbb{R})$ is defined by (8).*

This setup allows us to define the *flux-monodromy morphism* as in Definition 1.4.

Before moving on, we consider how the flux-monodromy morphism behaves under concatenation. By $\Pi * \Pi'$, we denote the concatenation of the Lagrangian path Π followed by Π' . By Π^- , we denote the inverse path of Π , defined by $\Pi_t^- = \Pi_{1-t}$.

Proposition 2.4. *The flux-monodromy map $\text{FM}_L : \pi_1(\mathcal{L}) \rightarrow H^1(L; \mathbb{R}) \rtimes \text{MCG}(L)$ has the following properties:*

(1) *It is a morphism: For any $\Lambda, \Lambda' \in \pi_1(\mathcal{L})$ it satisfies*

$$\text{FM}_L(\Lambda * \Lambda') = (\mathcal{M}(\Lambda)^* \text{Flux}(\Lambda') + \text{Flux}(\Lambda), \mathcal{M}(\Lambda') \mathcal{M}(\Lambda));$$

(2) *It has a natural change-of-basepoint formula: Let $\Lambda \in \pi_1(\mathcal{L}, L)$ and Π be a path in $\mathcal{L}$ with endpoints $\Pi(0) = L$ and $\Pi(1) = L'$. Then $\Pi^- * \Lambda * \Pi \in \pi_1(\mathcal{L}, L')$ satisfies*

$$\text{FM}_{L'}(\Pi^- * \Lambda * \Pi) = ((\phi^\Pi)^{-*} (\mathcal{M}(\Lambda)^* \text{Flux } \Pi + \text{Flux } \Lambda - \text{Flux } \Pi), [\phi^\Pi \circ \phi^\Lambda \circ (\phi^\Pi)^{-1}]). \quad (9)$$

The element $[\phi^\Pi \circ \phi^\Lambda \circ (\phi^\Pi)^{-1}] \in \text{MCG}(L')$ is well-defined because ϕ^Π is well-defined up to post-composition with an element in $\text{Diff}_0(L')$.

Proof. Let Π, Π' be paths in $\mathcal{L}$ such that $\Pi(1) = \Pi'(0)$, meaning that the concatenation $\Pi * \Pi'$ makes sense. We claim that

$$\text{Flux}(\Pi * \Pi') = (\phi^\Pi)^* \text{Flux } \Pi' + \text{Flux } \Pi. \quad (10)$$

Both properties listed in the proposition follow from a straightforward computation using (10). To prove the identity (10), recall that, from (8), $\text{Flux}(\Pi * \Pi') \cdot \xi$ is given by the symplectic area of the cylinder over ξ swept out by $\Pi * \Pi'$. Decomposing the cylinder as $C_\xi^{\Pi * \Pi'} = C_{\phi_*^\Pi \xi}^{\Pi'} \cup C_\xi^\Pi$ and applying $[\omega]$ yields the claim. $\square$

By $[\omega]_L \in H^2(X, L; \mathbb{R})$ we denote the class in relative cohomology which assigns to any class in $H_2(X, L)$ its symplectic area. Note that any element $b \in H_2(X, L)$ and any path $\Pi = \{L_t\}_{t \in [0,1]}$ starting at $L = L_0$ induces a natural element $\Phi^\Pi b \in H_2(X, L_1)$ with boundary $\partial \Phi^\Pi b = \phi_*^\Pi \partial b$ when transported using Π . Furthermore, $\Phi^\Pi b$ is homologous to $b + C_{\partial b}^\Pi$. Hence, applying $[\omega]$ yields the *area identity*:

$$(\Phi^\Pi)^*[\omega]_{L_1} - [\omega]_{L_0} = \partial^* \text{Flux } \Pi. \quad (11)$$

We denote the Maslov class (on homology) by $m_L \in H^2(X, L; \mathbb{Z})$ and call its evaluation on a surface with boundary on L the Maslov index of the surface. Since the Maslov class is a discrete invariant, it is preserved under Lagrangian isotopies. This implies an important property.

Proposition 2.5. *Let $\Lambda = \{L_t\}$ be a Lagrangian loop, i.e. $L_1 = L_0 = L$. For any loop c on L , the associated cylinder C_ξ^Λ has Maslov index 0.*

This generalizes the fact that tori defined by loops of symplectic diffeomorphisms have vanishing first Chern number, following classical work of McDuff [McD84].

Proof. Fix a parametrization $\{\iota_t : L \hookrightarrow X\}$ of Λ , and let $c : S^1 \rightarrow L$ with $[c] = \xi$, and consider $f : S^1 \times [0, 1] \rightarrow X; f(s, t) = \iota_t(c(s))$. Note that, since complex bundles over surfaces are classified by the first Chern class and $H^2(S^1 \times [0, 1]; \mathbb{Z}) = 0$, f^*TX is trivial. Moreover, given a trivialization of f^*TX , $m_L(C_\xi^\Lambda)$ is obtained as the difference of the Maslov indices of $s \mapsto T_{\iota_1(c(s))}L_1$ and $s \mapsto T_{\iota_0(c(s))}L_0$, seen, via that trivialization, as loops into the Lagrangian Grassmannian of $\mathbb{R}^{\dim X}$. But $(s, t) \mapsto T_{\iota_t(c(s))}L_t$ gives a homotopy between these loops, so that these Maslov indices are the same. $\square$

This result is particularly useful when studying *H-monotone* Lagrangians, i.e. those L such that $[\omega]_L = \frac{C}{2}m_L$ for some $C \in \mathbb{R}$. Indeed, we directly get the following.

Corollary 2.6. *Suppose that L is an H-monotone Lagrangian in X . Then, $\text{Flux}_L \Lambda = 0$ for every loop $\Lambda \in \pi_1(\mathcal{L}, L)$.*

3. DISCRETENESS AND TOPOLOGICAL PROPERTIES

We first prove Theorem 1.7, which gives necessary and sufficient conditions to ensure discreteness of the flux-monodromy group in §3.1. We then turn to the topological consequences of said discreteness, namely by proving Propositions 1.9 and 1.10. These are divided into two parts: local path connectedness in §3.2 and closedness in §3.3. Proposition 1.12, which ensures discreteness of the flux set in various settings, is proved at the end of §3.2.

3.1. Proof of Theorem 1.7. We start with the proof of the second statement in Theorem 1.7, i.e. that discreteness of $\mathcal{P}(X) = \omega(\ker c_1)$ implies discreteness of $\Gamma(L)$ for all $L \subseteq X$. In fact, we prove a slightly stronger statement that will be useful for applications.

Consider the natural projection $p^{\mathbb{Z}} : H^1(L; \mathbb{R}) \rtimes \text{MCG}(L) \rightarrow H^1(L; \mathbb{R}) \rtimes \text{Aut}(H_1(L; \mathbb{Z}))$. One can easily check that this is a morphism. We call $\text{FM}^{\mathbb{Z}} := p^{\mathbb{Z}} \circ \text{FM}$ the *homological flux-monodromy morphism* and denote its image by $\Gamma^{\mathbb{Z}}(L)$.

Proposition 3.1. *Assume that $\mathcal{P}(X)$ is discrete. Then the homological flux-monodromy group $\Gamma^{\mathbb{Z}}(L)$ is discrete for all Lagrangians $L \subset X$.*

Considering the projection $\Gamma(L) \rightarrow \Gamma^{\mathbb{Z}}(L)$ and the discreteness of $\text{MCG}(L)$, this obviously implies the second part of Theorem 1.7.

Proof. Let $\Lambda \in \pi_1(\mathcal{L}, L)$ be a loop in the space of Lagrangians of some Lagrangian submanifold $L \subset X$. If its monodromy preserves some $\xi \in H_1(L)$, i.e. if $\mathcal{M}(\Lambda)\xi = \xi$, then the cylinder C_{ξ}^{Λ} swept out by ξ under Λ satisfies $\partial C_{\xi}^{\Lambda} = 0$ and thus defines² a class $\Sigma_{\xi} = [C_{\xi}^{\Lambda}] \in H_2(X)$. This correspondance satisfies

$$\text{Flux}(\Lambda)\xi = \omega(\Sigma_{\xi}) \quad \text{for all } \xi \in H_1(L). \quad (12)$$

Moreover, the first Chern number $c_1(\Sigma_{\xi}) \in \mathbb{Z}$ of Σ_{ξ} is half the Maslov index $m_L(C_{\xi}^{\Lambda})$ of C_{ξ}^{Λ} . By Proposition 2.5, we have that $c_1(\Sigma_{\xi}) = 0$, and $\text{Flux}(\Lambda)\xi$ is in $\mathcal{P}(X) \subset \mathbb{R}$ whenever Λ has trivial homological monodromy.

Now let $\{(f_n, M_n)\}_{n \in \mathbb{N}} \in \Gamma^{\mathbb{Z}}(L) \subset H^1(L; \mathbb{R}) \rtimes \text{Aut}(H_1(L; \mathbb{Z}))$ be a sequence converging to (f, M) , and let $\Lambda_n \in \pi_1(\mathcal{L}, L)$ be a sequence of loops with $\text{FM}^{\mathbb{Z}}(\Lambda_n) = (f_n, M_n)$. Note that $\text{Aut}(H_1(L; \mathbb{Z}))$ is itself a discrete group. Therefore, the sequence M_n is constant and equal to M for all n greater than or equal to some n_0 . The sequence $\{\Lambda_n * \Lambda_{n_0}^{-}\}_{n \in \mathbb{N}}$ has flux $\text{Flux}(\Lambda_n * \Lambda_{n_0}^{-}) = f_n - M^* f_{n_0}$ and trivial monodromy for all $n \geq n_0$ by Proposition 2.4. By the above paragraph, it follows that $f_n - M^* f_{n_0} \in H^1(L; \mathcal{P}(X))$. Since the latter group is discrete, the sequence f_n must stabilize, too. $\square$

Remark 3.2. A few things need to be noted about this proof.

- (1) An approach analogous to the latter part of the proof of Proposition 3.1 also shows that $\Gamma(L)$ is discrete if and only if its subgroup

$$\Gamma_0(L) = \text{FM}(\ker(\mathcal{M}: \pi_1(\mathcal{L}, L) \rightarrow \text{MCG}(L)))$$

is. This will be used later on.

- (2) The proof of Proposition 3.1 only uses classes in $H_2(X)$ realized by tori. Therefore, we only need the subset of $\mathcal{P}(X)$ generated by Chern-zero tori to be discrete. For example, products of surfaces of genus at least 2 have this property, although the full $\mathcal{P}(X)$ might not be discrete.
- (3) When only $\mathcal{P}_S(X)$, the subset of $\mathcal{P}(X)$ formed by spherical classes, is known to be discrete, essentially the same proof as above implies that $\Gamma^{\mathbb{Z}}(L)$ is discrete for all L .

²This class is unique up to a choice of element in $H_2(L)$, but the symplectic area (12) does not depend on this choice, since L is Lagrangian.

bounding enough disks, i.e. such that the natural map $\pi_2(X, L) \rightarrow \pi_1(L) \rightarrow H_1(L)$ has finite cokernel. $\bullet$

Before proving the other part of Theorem 1.7, we note that the proof of Proposition 3.1 actually allows us to compute $\Gamma(L)$ and $F(L)$ in some important cases.

Proposition 3.3. *Let X be H -monotone, meaning monotone on $H_2(X)$, and let $L \subseteq X$ be a Lagrangian. The natural projections $\Gamma(L) \rightarrow \mathcal{M}(L)$ and $\Gamma^{\mathbb{Z}}(L) \rightarrow \mathcal{M}^{\mathbb{Z}}(L)$ are isomorphisms. If moreover the inclusion $L \hookrightarrow X$ induces an injection $H_1(L; \mathbb{R}) \hookrightarrow H_1(X; \mathbb{R})$, then $F(L) = 0$.*

As the notation suggests, $\mathcal{M}^{\mathbb{Z}}(L)$ is the image of the monodromy group $\mathcal{M}(L)$ of L in $\text{Aut}(H_1(L; \mathbb{Z}))$.

Proof. The kernel $\Gamma_0(L)$ of $\Gamma(L) \rightarrow \mathcal{M}(L)$ consists of the set of fluxes realized by loops with trivial monodromy. By the proof of Proposition 3.1, these fluxes are all in $H^1(L; \mathcal{P}(X))$. But by H -monotonicity, $\mathcal{P}(X) = 0$, so that $\Gamma_0(L)$ is trivial. The proof for $\Gamma^{\mathbb{Z}} \rightarrow \mathcal{M}^{\mathbb{Z}}(L)$ is completely analogous.

Suppose now that $H_1(L; \mathbb{R}) \hookrightarrow H_1(X; \mathbb{R})$. Then, the real homological monodromy $\mathcal{M}^{\mathbb{R}}(\Lambda)$ of any Lagrangian loop Λ is trivial. By the universal coefficient theorem, this means that $\mathcal{M}^{\mathbb{Z}}(\Lambda)$ acts trivially on the free part of $H_1(L)$. Since the torsion part of $H_1(L)$ is finitely generated, there can be no homomorphism from it to the free part, and its automorphism group is finite. Therefore, there is some N such that $\mathcal{M}^{\mathbb{Z}}(\Lambda)^N = \mathcal{M}^{\mathbb{Z}}(\Lambda^N)$ is the identity. By the first paragraph, this means that $\text{Flux}(\Lambda^N) = N \text{Flux}(\Lambda) = 0$, where the first equality follows from the fact that Flux is a homomorphism in this setting. Therefore, $F(L) = 0$ as stated. $\square$

Note that this proves Case (c) of Proposition 1.12. The rest will be proven in Section 3.2.

For the other direction in Theorem 1.7, we prove its contrapositive, i.e. under the assumption that $\mathcal{P}_S(X)$ is not discrete, we will find a Lagrangian L such that the image of FM_L is not discrete. The Lagrangian L is obtained by embedding a small product torus

$$T(a) = T(a_1, \dots, a_n) = \{\pi|z_i|^2 = a_i\} \subset \mathbb{C}^n \quad (13)$$

for $a = (a_1, \dots, a_n) \in \mathbb{R}_{>0}^n$ by a symplectic embedding $\varphi: B^{2n}(b) \hookrightarrow X$. Note that any two such tori $\varphi(T(a))$ and $\varphi(T(a'))$ are Lagrangian isotopic by a deformation through the images of product tori. We call this the *obvious* Lagrangian isotopy. Such tori were studied in great detail by Chekanov–Schlenk in [CS16]; in fact, our argument is a straightforward adaptation of [CS16, Theorem 1.5] that takes into account the monodromy of the Hamiltonian diffeomorphisms appearing in it.

Theorem 3.4. *Let X be a symplectic manifold, and consider $\varphi: B^{2n}(b) \hookrightarrow X$ a symplectic ball embedding. For every real number $c > 0$, there exists $A > 0$ such that for all $a \in (0, A]$, the following holds. If $d_1, \dots, d_k$ and $e_1, \dots, e_k$ satisfy $d_j, e_j \geq c$ for all $j \in \{1, \dots, k\}$ and*

$$d_j - e_j \in \mathcal{P}_S(X), \quad (14)$$

then the tori (provided the corresponding product tori are contained in the domain of φ),

$$L = \varphi(T(a, \dots, a, a + d_1, \dots, a + d_k)), \quad L' = \varphi(T(a, \dots, a, a + e_1, \dots, a + e_k)) \quad (15)$$

are Hamiltonian isotopic through a Hamiltonian isotopy with trivial monodromy. Here, $\text{Diff}(L, L')$ and $\text{Diff}(L)$ are identified by the obvious Lagrangian isotopy.

The case where X is *special* (in the terminology of Chekanov–Schlenk) is absent from the statement because if X is special, then, in particular, $\mathcal{P}_S(X)$ is discrete. The key difference with respect to [CS16, Theorem 1.5] is that we take the $d_j - e_j$ to lie in $\mathcal{P}_S(X)$ instead of the larger group G_a generated by the quantities $\omega(S) - ac_1(S)$ for $S \in \pi_2(X)$. Note that $\mathcal{P}_S(X)$ is obtained from G_a by restricting to elements $S \in \ker c_1$. As will become apparent in the proof, considering spheres with trivial first Chern number is precisely the restriction needed to obtain trivial monodromy of the Hamiltonian diffeomorphisms constructed by Chekanov–Schlenk.

Proof of Theorem 3.4. Following the notation of the proof of [CS16, Theorem 1.5], we define, on a suitable domain in $(\mathbb{C}^\times)^2$ and for some $m \in \mathbb{Z}$ and $s > 0$, the symplectomorphism

$$\psi_{m,s}(\rho_1, \theta_1, \rho_2, \theta_2) = (\rho_1, \theta_1 + m\theta_2, \rho_2 + s - m\rho_1, \theta_2), \quad (16)$$

which we have expressed in symplectic polar coordinates $(\rho_i, \theta_i) \in \mathbb{R}_{\geq 0} \times S^1$ defined by $z_i = \sqrt{\rho_i} e^{2\pi i \theta_i}$ for all $z_i \neq 0$. Now let $S \in \pi_2(X)$ and set $s = \omega(S)$ and $m = c_1(S)$. In the case $n = 2$, [CS16, Proposition 6.1] shows that there is a neighbourhood U of the isotropic circle $\{0\} \times S^1(y_0) \subset \mathbb{C}^2$ and a Hamiltonian isotopy χ of X which restricts to $\varphi \circ \psi_{m,s} \circ \varphi^{-1}|_{\varphi(U)}$ on $\varphi(U)$. Thus using (16), we find that for every $T(x, y) \subset U$, the Hamiltonian diffeomorphism χ maps $\varphi(T(x, y))$ to $\varphi(T(x, y + s - mx))$ with monodromy, expressed in the standard basis of product tori, given by

$$(\chi|_{T(x,y)})_* = \begin{pmatrix} 1 & m \\ 0 & 1 \end{pmatrix}. \quad (17)$$

Note that the fact that these tori are Hamiltonian isotopic relies non-trivially on the existence of the sphere S , since in $\mathbb{C}^2$ itself, the product tori $T(x, y)$ and $T(x', y')$ are Hamiltonian isotopic if and only if (x, y) and (x', y') agree up to permutation. The presence of S allows one to shift the tori by an amount depending on the area and the Maslov class of S . The key idea of our proof is restricting our attention to elements satisfying $c_1(S) = 0$, since these have trivial monodromy by (17). For general n , the corresponding Hamiltonian isotopy is obtained as a product of $\psi_{m,s}$ with the identity in the remaining components, and we obtain the following as a straightforward consequence of [CS16, Proposition 6.1].

Proposition 3.5. *Let $S \in \mathcal{P}_S(X) = \ker(c_1: \pi_2(X) \rightarrow \mathbb{Z})$, set $s = \omega(S)$, and let $d_1, \dots, d_k$ for some $k \geq 1$. If furthermore $d_1 + \dots + d_k + s < b$, then there is a neighbourhood U of $(0, \dots, 0, d_1, \dots, d_k) \in \mathbb{R}_{\geq 0}^n$ and a Hamiltonian isotopy χ such that*

$$\chi(\varphi(T(a_1, \dots, a_{n-1}, a_n))) = \varphi(T(a_1, \dots, a_{n-1}, a_n + s)) \quad (18)$$

for all $a = (a_1, \dots, a_n) \in U$. Furthermore, χ induces trivial monodromy under the identification of the mapping class groups under the obvious Lagrangian isotopy.

This is the main building block of the proof of [CS16, Theorem 1.5]. The remainder of their strategy can be implemented by replacing $\pi_2(X)$ by $\ker c_1 \subset \pi_2(X)$ and adding the conclusion about trivial monodromy everywhere. $\square$

Corollary 3.6. *If $\mathcal{P}_S(X)$ is not discrete, then there is a Lagrangian torus $L \subset X$ such that FM_L has non-discrete image.*

Proof. Let $\varphi: B^{2n}(b) \hookrightarrow X$ be a symplectic ball embedding. Choose $a > 0$, $k \geq 1$, and a sequence $d_i = (d_{i,1}, \dots, d_{i,k})$ of vectors converging componentwise to $d = (d_1, \dots, d_k)$ such that $d_j - d_{i,j} \in \mathcal{P}_S(X)$ for all $j \in \{1, \dots, k\}$. Since $\mathcal{P}_S(X) \subset \mathbb{R}$ is dense, we can assume that all d_i are pairwise distinct. Let $L = \varphi(T(a, \dots, a, a + d_1, \dots, a + d_k))$ and $\mathcal{L}$ the space of its Lagrangian deformations. Define $\Lambda_i \in \pi_1(\mathcal{L})$ to be the loop obtained by applying first the Hamiltonian isotopy from L to $\varphi(T(a, \dots, a, a + d_{i,1}, \dots, a + d_{i,k}))$ obtained from Theorem 3.4, followed by the obvious Lagrangian isotopy. The monodromy of all Λ_i is trivial, again by Theorem 3.4. Its flux is given by $\text{Flux } \Lambda_i = (0, \dots, 0, d_1 - d_{i,1}, \dots, d_k - d_{i,k})$ in the pushforward of the standard basis of the homology of the product torus. By our choice of d_i , this proves the claim. $\square$

3.2. C^1 -local path connectedness. We now move to the proofs of the part of Propositions 1.9 and 1.10 having to do with local path connectedness. As a byproduct, we also give a proof of Proposition 1.12. We start with the most straightforward direction.

Lemma 3.7. *If 0 is isolated in $F(L)$ (resp. $\Gamma(L)$), then $\mathcal{L}\text{Ham}(L)$ (resp. $\mathcal{L}\text{Ham}(\iota)$) is locally path connected in the C^1 topology.*

Again, 0 being isolated in $\Gamma(L)$ is equivalent to it being discrete, since it is a group.

Proof. Suppose that $\mathcal{L}\text{Ham}(L)$ is not C^1 -locally path connected. Then, there are a Weinstein neighbourhood $\Psi: D_r^*L \rightarrow X$ of L and a sequence of Lagrangians $\{L_k\}$ in $U = \Psi(D_r^*L)$, which are Hamiltonian isotopic to L in X , but not in U , such that $L_k \rightarrow L$ in the C^1 topology. Indeed, if L_k were Hamiltonian isotopic to L in U and C^1 -close to it, then it would need to be the graph of a (closed) 1-form σ_k that vanishes on all loops, i.e. $\sigma_k = df_k$ for some $f_k: L \rightarrow \mathbb{R}$. But then, $t \mapsto \Psi(\text{graph}(tdf_k))$ would be a C^1 -small path in $\mathcal{L}\text{Ham}(L)$ from L to L_k in U , contradicting the failure of path-connectedness.

By C^1 convergence, we have that $L_k = \Psi(\text{graph } \sigma_k)$ for $\sigma_k \rightarrow 0$ in C^1 topology. As noted above, σ_k is not exact. In fact, the flux of the path $t \mapsto \Psi(\text{graph}(t\sigma_k))$ is precisely $[\sigma_k] \neq 0$. But we know that there is an isotopy $t \mapsto \phi_H^t(L)$ from L to L_k which is Hamiltonian and thus must have vanishing flux. Therefore, the concatenation of these isotopies forms a Lagrangian loop with nonvanishing flux $[\sigma_k] \rightarrow 0$. Therefore, 0 is not isolated in $F(L)$.

For the case of $\mathcal{L}\text{Ham}(\iota)$, the argument is essentially the same. However, ι and the members $\phi_k \circ \iota$ of the sequence may now be chosen to be C^1 -close, so that the ensuing loop has no

monodromy. Therefore, we instead get that 0 is not isolated in $\Gamma_0(L)$ — and thus neither in $\Gamma(L)$. Recall that $\Gamma_0(L) = \text{FM}(\ker \mathcal{M})$ is the subgroup of $\Gamma(L)$ formed by the image under the flux-monodromy map of loops with trivial monodromy. $\square$

We now show a partial converse statement to the above lemma.

Lemma 3.8. *Suppose that 0 is not isolated in $F(L)$. There is a sequence $\{L_k\}$ such that*

- (i) $L_k \rightarrow L$ in C^1 topology;
- (ii) *there exists a Lagrangian isotopy from L to L_k with flux zero;*
- (iii) L_k *is not Hamiltonian isotopic to L in any Weinstein neighbourhood of L .*

Moreover, if $\Gamma(L)$ is not discrete, then we may take a parametrization $I_k = \{\iota_{k,t}\}$ of the Lagrangian isotopy from L to L_k in (ii) so that $\iota_{k,1} \rightarrow \iota$ in C^1 topology.

Proof. Let Λ_k be a sequence of Lagrangian loops with $\text{Flux } \Lambda_k = a_k \neq 0$ such that $a_k \rightarrow 0$. Let Ψ be a Weinstein neighbourhood of L . For large enough k , there are closed 1-forms σ_k with $[\sigma_k] = a_k$, such that the graph of σ_k is contained in the domain of Ψ . We set $L_k = \Psi(\text{graph}(\sigma_k))$ and show that this has the desired properties. Since $a_k \rightarrow 0$, we have that $L_k \rightarrow L$ in the C^1 -topology, proving (i). For (ii), consider the Lagrangian isotopy obtained by concatenating $t \mapsto \Psi(\text{graph}((1-t)\sigma_k))$ with the loop Λ_k . This has flux $-a_k + a_k = 0$ as desired. Since $a_k \neq 0$, L_k is not Hamiltonian isotopic to L in any Weinstein neighbourhood, proving (iii).

Finally, when $\Gamma(L)$ is not discrete, then $\Gamma_0(L)$ is not either; see Remark 3.2. Therefore, in that case, Λ_k may be chosen to have trivial monodromy. In other words, for a parametrization I_k of the loop, the diffeomorphism $f_k = \iota_{k,1}^{-1} \circ \iota$ is in the identity component of $\text{MCG}(L)$. But then, any isotopy from the identity to f_k^{-1} in $\text{Diff}(L)$ extends via the Weinstein neighbourhood to a Hamiltonian isotopy of X fixing L . Therefore, by concatenating with such an isotopy, we can suppose that f_k is the identity without affecting the flux. The rest of the proof then follows as before. $\square$

We can upgrade the Lagrangian isotopy from the above lemma to a Hamiltonian one in two important cases.

Proposition 3.9 ([Che96]). *Let L be a product torus in $\mathbb{C}^n$. Then, the Lagrangian isotopy of Lemma 3.8 may be chosen to be Hamiltonian.*

Note that the same results of Chekanov imply that if a product torus has non-discrete flux set, then the flux set is dense in $H^1(L; \mathbb{R})$. In particular, 0 is not isolated.

Proof. Suppose that 0 is not isolated in $F(L)$. Note that the existence — from Lemma 3.8 — of a Lagrangian isotopy from L to L_k with zero flux together with the area identity (11) implies the existence of isomorphisms $\Phi_k : \pi_2(\mathbb{C}^n, L) \rightarrow \pi_2(\mathbb{C}^n, L_k)$ such that

$$\Phi_k^*[\omega]_{L_k} = [\omega]_L \quad \text{and} \quad \Phi_k^* \mu_{L_k} = \mu_L.$$

Furthermore, for k large enough, we can assume that L_k is also an elementary torus. Therefore, the result follows from Chekanov's classification of elementary tori in $\mathbb{C}^n$ [Che96]. $\square$

Proposition 3.10. *Let $L \subseteq X$ be a Lagrangian such that the map $\iota_* : H_1(L; \mathbb{R}) \rightarrow H_1(X; \mathbb{R})$, induced by the inclusion, is injective. Then, any Lagrangian isotopy as in Lemma 3.8 is homotopic to a Hamiltonian one.*

Proof. Suppose that 0 is not isolated in $F(L)$. Since ι_* is injective, the boundary map $H_2(X, L; \mathbb{R}) \rightarrow H_1(L; \mathbb{R})$ is zero, so that every Lagrangian isotopy of L preserves the area spectrum by the area identity (11). By Proposition A.1, there is thus a symplectic isotopy $\{\psi_t\}$ generating any Lagrangian isotopy $\{L_t\}$ based at L . In particular, for a Lagrangian isotopy $\{L_t\}$ as constructed in Lemma 3.8, we have that

$$\iota^* \text{Flux}(\{\psi_t\}) = \text{Flux}(\{L_t\}) = 0.$$

But note that $\ker \iota^*$ is precisely the image of $H^1(X, L; \mathbb{R}) \rightarrow H^1(L; \mathbb{R})$, so that there is a closed 1-form σ on M that vanishes on (a neighbourhood of) L . Note that such a form generates a symplectic isotopy that fixes L and L_1 . Therefore, by applying the symplectic isotopy generated by $-\sigma$, we get a homotopy from our original Lagrangian isotopy to one that is generated by $\{\psi_t\}$ such that

$$\text{Flux}(\{\psi_t\}) = 0.$$

But then, up to a homotopy relative to endpoints, $\{\psi_t\}$ is in fact Hamiltonian (see, for example, [MS17, Section 10.2]). $\square$

Remark 3.11. In general, one cannot conclude that a Lagrangian isotopy with vanishing flux is homotopic to a Hamiltonian isotopy. For example, the Clifford and Chekanov tori are not Hamiltonian isotopic in $\mathbb{C}^n$, but are Lagrangian isotopic [Che96]. If both tori have the same monotonicity constant, such a Lagrangian isotopy must have vanishing flux by the area formula (11), since the Maslov class must be preserved.

If one looks at the classical proof that a symplectic isotopy with vanishing flux is homotopic to a Hamiltonian one [Ban97], we see that two problems arise in the Lagrangian setting:

- (1) if $\{f_t\}$ and $\{g_t\}$ are ambient Lagrangian isotopies starting at L , $\{f_t g_t\}$ is not necessarily one, and
- (2) even when we can compose ambient Lagrangian isotopies, the homotopy of the Lagrangian isotopy will result in a family of Lagrangian *immersions*.

See Appendix B for a solution to the first problem through h -principles. $\bullet$

To close this subsection, we finally prove Proposition 1.12.

Proof of Proposition 1.12. As noted before, Proposition 3.3 gives a proof of Case (c).

Suppose that we are in Case (a), i.e. L is a weakly exact Lagrangian. If $L' \in \mathcal{L}\text{Ham}(L)$ were C^1 -close to L but not Hamiltonian isotopic to L in a Weinstein neighbourhood, then it would be the graph of a nonexact closed 1-form σ in that neighbourhood. However, every

nonexact closed 1-form is cohomologous to a 1-form with at most one zero [FS06], i.e. L' is Hamiltonian isotopic in a Weinstein neighbourhood of L to some Lagrangian with only one intersection point with L . By the degenerate Lagrangian Arnol'd conjecture for weakly exact Lagrangians [Flo89, Hof88], this is not possible. Therefore, $\mathcal{L}\text{Ham}(L)$ must be C^1 -locally path connected. By Proposition 3.10, this implies the discreteness of $F(L)$ — closedness follows from Proposition 3.12 below.

For Case (b), i.e. when L is a non-displaceable torus, the proof is essentially the same: if L' were the graph of a nonexact closed 1-form, it would be displaceable from L in a Weinstein neighbourhood, contradicting the nondisplaceability of L . $\square$

3.3. C^1 -closedness. In [Sol13], it is claimed that discreteness of $F(L)$ implies that the Hamiltonian orbit of L in X is C^1 -closed in the space of Lagrangians of X diffeomorphic to L . However, [Sol13] does not take into account monodromy and erroneously claims that $F(L)$ is a subgroup of $H^1(L; \mathbb{R})$. We now fix this gap.

Proposition 3.12. *If 0 is isolated in $F(L)$, then the Hamiltonian orbit of L is C^1 -closed in $\mathcal{L}(L)$.*

Proof. Suppose that $\{L_i\} \subseteq \mathcal{L}\text{Ham}(L)$ C^1 -converges to some Lagrangian K . Then, for i large, L_i is the graph of a closed 1-form in a Weinstein neighbourhood of K . For each i , make a Lagrangian loop Λ_i based at L by concatenating

- (1) a Hamiltonian isotopy from L to L_i ;
- (2) the obvious Lagrangian isotopy from L_i to K in the Weinstein neighbourhood;
- (3) the obvious Lagrangian isotopy from K to L_{i+1} ;
- (4) a Hamiltonian isotopy from L_{i+1} to L .

On the one hand, the flux of Λ_i goes to 0 as $i \rightarrow \infty$ since $L_i \rightarrow K$ in the C^1 topology. On the other hand, 0 is isolated in $F(L)$. Therefore, $\text{Flux}(\Lambda_i) = 0$ for large i . But the flux of Λ_i is simply the flux of the two obvious isotopies from L_i to L_{i+1} , which we can compute to be $[\sigma_{i+1} - \sigma_i]$. Therefore, for i large, $[\sigma_i] \in H^1(K; \mathbb{R})$ is constant. Since $[\sigma_i] \rightarrow 0$ by C^1 -convergence, we must have that $[\sigma_i] = 0$, i.e. L_i is the graph of an exact form for i large. Therefore, L_i is Hamiltonian isotopic to K (in its Weinstein neighbourhood). $\square$

By a slight adaptation of Solomon's arguments, we can prove the parametrized form of his statement.

Proposition 3.13. *If $\Gamma(L)$ is discrete, then the Hamiltonian orbit of $\iota : L \hookrightarrow X$ is C^1 -closed in $\tilde{\mathcal{L}}(L)$, the space of Lagrangian embeddings isotopic to $\iota : L \hookrightarrow X$.*

Proof. If $\Gamma(L)$ is discrete, then so is $G = \text{Flux}(\ker \mathcal{M}) \cong \Gamma_0(L)$. But then, Flux defines a map $\overline{\text{Flux}} : \tilde{\mathcal{L}} \rightarrow H^1(L; \mathbb{R})/G$, where $H^1(L; \mathbb{R})/G$ is a manifold. It is easy to see that 0 is a regular value of $\overline{\text{Flux}}$. Therefore, $\overline{\text{Flux}}^{-1}(0)$ is a closed manifold. It follows essentially from Proposition A.1 below or Corollary 6.4 of [Sol13] that the connected component of $\overline{\text{Flux}}^{-1}(0)$

containing ι in $\tilde{\mathcal{L}}$ — and thus in the space of all Lagrangian embeddings — is its Hamiltonian orbit, which concludes the proof. $\square$

We also prove the converse, under certain topological conditions, by essentially following the proof in the symplectic case.

Lemma 3.14. *Suppose that $\iota_* : H_1(L; \mathbb{R}) \rightarrow H_1(X; \mathbb{R})$ is injective or that L is a product torus in $\mathbb{C}^n$. If $\mathcal{L}\text{Ham}(L)$ (resp. $\mathcal{L}\text{Ham}(\iota)$) is C^1 -closed in $\mathcal{L}(L)$ (resp. $\tilde{\mathcal{L}}(\iota)$), then $F(L)$ (resp. $\Gamma(L)$) is discrete.*

Proof. First, note that non-discreteness implies density in our setting. For $\Gamma(L)$ or for $F(L)$ with ι_* injective, this follows from the fact that these are subgroups of some real vector space. For the case when L is a product torus, this follows from the classification of product tori by Chekanov [Che96].

Suppose that $F(L)$ is dense, and take a C^1 -small closed 1-form σ on L such that $[\sigma] \in H^1(L; \mathbb{R}) \setminus F(L)$. Note that its graph L' is not Hamiltonian isotopic to L in X , since otherwise the concatenation of the obvious isotopy from L to L' with a Hamiltonian isotopy from L' to L would be a Lagrangian loop of flux $[\sigma]$. By an analogous construction as in Lemma 3.8 (also using Propositions 3.9 and 3.10), we can take a sequence of C^1 -small closed 1-forms σ_k with graphs L_k such that

- (i) $L_k \rightarrow L'$ in C^1 topology;
- (ii) L and L_k are Hamiltonian isotopic in X ;
- (iii) $[\sigma_k] \in F(L)$.

Since $L' \notin \mathcal{L}\text{Ham}(L)$, the Hamiltonian orbit is not closed.

When $\Gamma(L)$ is dense, we mirror the construction using the density of $\Gamma_0(L)$, just as in Lemma 3.8. $\square$

4. IN $\mathbb{C}^2$: A CASE STUDY

We review preliminary material about Arnol'd–Liouville tori in §4.1 and almost toric fibrations in §4.2. We compute the flux-monodromy group of Lagrangian tori in $\mathbb{C}^2$ in §4.3, and prove the existence of a section to the flux-monodromy morphism in §4.4. Theorem 1.17 follows from these results. We also establish a complete description of the fundamental group of the space of *immersed* Lagrangians in §4.5. In §4.6, we study the Hamiltonian action on $\mathcal{L}(L)$, when L is a Lagrangian torus in several natural 4-manifolds. In §4.7, we prove Proposition 1.12, which states a strong form of local path connectedness of the Hamiltonian orbit of any Lagrangian torus in $\mathbb{C}^2$ in the Hausdorff metric.

4.1. Preparation: isotopies of Arnol'd–Liouville tori. Although we will only use the following in the special case of an almost toric fibration on $\mathbb{C}^2$, we briefly discuss properties of the flux-monodromy morphism on lifts of curves in the base of completely integrable Hamiltonian systems. This discussion is based on [Dui80] and [Eva23]. Let $F : (X, \omega) \rightarrow \mathbb{R}^n$

be a completely integrable Hamiltonian system on a symplectic $2n$ -manifold. This means that $\{F_i, F_j\} = 0$ and that F has compact, connected fibres. Furthermore, we ask that the set of regular values $B_{\text{reg}} \subset \mathbb{R}^n$ is open and dense in the image of F . The classical Arnol'd–Liouville theorem implies that $F^{-1}(b) \subset X$ is a Lagrangian torus for all $b \in B_{\text{reg}}$. These are called *Arnol'd–Liouville tori*. Lifting paths in B_{reg} yields Lagrangian isotopies.

Definition 4.1. *Let $\gamma: [0, 1] \rightarrow B_{\text{reg}}$ be a smooth curve in the regular values of an integrable system F . We call the Lagrangian isotopy $\Pi_\gamma = \{F^{-1}(\gamma(t))\}_{t \in [0, 1]}$ the lift of γ .*

There is a strong relationship between action coordinates of the system F and the Lagrangian flux of the lift of a curve. Following Evans [Eva23, Definition 2.14], we fix a regular fibre $L = F^{-1}(b_0)$ and define the *flux map* $I: \tilde{B}_{\text{reg}} \rightarrow H^1(L; \mathbb{R})$ on the universal cover $\tilde{B}_{\text{reg}} \rightarrow B_{\text{reg}}$ as the Lagrangian flux of the lift of a path between b_0 and b . Note that we have passed to the universal cover, since the Lagrangian flux may depend on the choice of path. From that point of view, we obtain

$$\text{Flux } \Pi_\gamma = I(\tilde{\gamma}(1)) - I(\tilde{\gamma}(0)), \quad (19)$$

where $\tilde{\gamma}: [0, 1] \rightarrow \tilde{B}_{\text{reg}}$ is a lift of γ to the universal cover $\tilde{B}_{\text{reg}}$. As it turns out, the map I computes action coordinates.

Proposition 4.2. *Identifying $H^1(T^n; \mathbb{R}) = \mathbb{R}^n$, the map I computes action coordinates on small enough subsets $U \subset \tilde{B}_{\text{reg}}$. Conversely, any set of action coordinates coincides with I up to an integral affine transformation, i.e. up to post-composing with a map in $\mathbb{R}^n \rtimes \text{GL}(n; \mathbb{Z})$.*

This proposition is a combination of [Eva23, Lemma 2.15] and the uniqueness of action coordinates; see [Eva23, Section 2.5]. Note that if F is already expressed in action coordinates, then $\text{Flux } \Pi_\gamma = \gamma(1) - \gamma(0)$, where we identify $H^1(T^n; \mathbb{R}) \cong \mathbb{R}^n$ via the basis on $H_1(T^n)$ generated by the closed orbits generated by the components F_i of F . An example of such an F is any moment map generating an effective Hamiltonian T^n -action. This occurs, for example, when $F: X \rightarrow \mathbb{R}^n$ is a toric moment map.

4.2. Almost toric fibrations. Recall that by $T(a, b) \subset \mathbb{C}^2$, we denote the product torus $S^1(a) \times S^1(b)$. For any $a > 0$, the torus of the form $T(a, a)$ is monotone and is called the *Clifford torus*. We denote it by $T_{\text{Cl}}(a)$. Another torus we will discuss is the so-called *Chekanov torus* $T_{\text{Ch}}(a)$, for all $a > 0$. It was introduced in [Che96]; see also [EP97]. By now, there are several equivalent points of view on the Chekanov torus. In this paper, we will use Chekanov's original definition as well as the fact that the Chekanov torus appears as a monotone fibre in the almost toric fibration obtained from performing a so-called *nodal trade* on the standard toric fibration on $\mathbb{C}^2$. Let us briefly outline the approach using almost toric fibrations (ATFs); for more details, we recommend [Eva23]. The map $F: \mathbb{C}^2 \rightarrow \mathbb{R}_{\geq 0}^2$ defined by $F_i = \pi|z_i|^2$ for $i \in \{1, 2\}$ is the standard moment map defining a toric Hamiltonian T^2 -action. From a slightly different point of view, it can be viewed as a special type of completely integrable Hamiltonian system:

- (1) Its singular values are on the boundary $\partial\mathbb{R}_{\geq 0}^2$. They consist of two families of singularities of *elliptic-regular type* at points of the form $(z_1, 0)$ for $z_1 \neq 0$ and $(0, z_2)$ for $z_2 \neq 0$, respectively. These singular points live over the interior of the edges of $\mathbb{R}_{\geq 0}^2$. These two families meet at the point $(0, 0)$ which is of *elliptic-elliptic type* and a fixed point of the T^2 -action.
- (2) Over the regular locus $\mathbb{R}_{> 0}^2$, the map F defines action coordinates whose corresponding angle coordinates are induced by the T^2 -action. Thus, F can be viewed as an Arnol'd–Liouville chart on the open dense set of regular points of $\mathbb{C}^2$. Its Arnol'd–Liouville tori are just the product tori $T(a, b)$.

There is a smooth deformation $\{F_t: \mathbb{C}^2 \rightarrow \mathbb{R}^2\}$ of the integrable system $F = F_0$, in which the elliptic-elliptic singularity disappears for $t > 0$ and a so-called *focus-focus* singularity appears in the system. Under this deformation, the two elliptic-regular families are merged into one such family. In almost toric geometry, this deformation is called a *nodal trade*, since the image of the focus-focus singularity under F_t is called a *node*. The node is isolated and has a neighbourhood in which every point except for itself is regular. In the spirit of Duistermaat [Dui80], one can try to recover action-coordinates on the complement of the focus-focus singularity. This fails, since the focus-focus singularity has non-trivial *monodromy* and the latter is an obstruction to finding global action-coordinates; see [Dui80, Theorem 2.2]. Monodromy here means that the torus bundle over the punctured neighbourhood of a node is non-trivial. In fact, there is a basis in which it has monodromy given by the matrix

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}. \quad (20)$$

The best one can do is to remove a ray connecting the node to the boundary of $\text{im } F_t \subset \mathbb{R}^2$ (or to infinity) and compute action coordinates on the complement of the ray. This yields the so-called *almost toric base diagram*, depicted in the case of $\mathbb{C}^2$ in the middle and on the right-hand side of Figure 3. In the middle picture, the Arnol'd–Liouville tori over the dashed line are Hamiltonian isotopic to Chekanov tori.

We use the almost toric base diagram in the right hand side of Figure 3 to define a non-trivial Lagrangian loop $\Xi \in \pi_1(\mathcal{L}, L)$ for $L = T_{\text{Cl}}$ or $L = T_{\text{Ch}}$.

Proposition 4.3. *Let L be any Clifford or Chekanov torus in $\mathbb{C}^2$. Then, there is a loop $\Xi \in \pi_1(\mathcal{L}, L)$ with*

$$\text{FM}(\Xi) = \left(0, \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}\right) \quad (21)$$

in a suitable basis of $H_1(T^2)$, which yields an identification of $H^1(T^2; \mathbb{R}^2) \rtimes \text{MCG}(T^2) \cong \mathbb{R}^2 \rtimes \text{GL}(2; \mathbb{Z})$.

Proof. Let us prove the claim for the Clifford torus $L = T_{\text{Cl}}(a)$. The proof for the Chekanov torus is analogous. As illustrated in Figure 3, a nodal trade followed by a change in branch cut and a $\text{GL}(2; \mathbb{Z})$ -shear produce an almost toric base diagram of $\mathbb{C}^2$. It is obtained by




FIGURE 3. Almost toric base diagrams corresponding to the nodal trade in $\mathbb{C}^2$. The image of F is depicted on the left-hand side. The Clifford torus $T_{\text{Cl}}(a)$ is the fibre over (a, a) . In the middle is the almost toric base diagram after a nodal trade, whose node sits between the points $(1, 1)$ and $(2, 2)$. The torus over $(1, 1)$ is the Chekanov torus $T_{\text{Ch}}(1)$ and the one over $(2, 2)$ is the Clifford torus $T_{\text{Cl}}(2)$. On the right hand side is another ATF-base diagram of the same fibration obtained from the middle one by applying a change in branch cut and a shear transformation in $\text{GL}(2; \mathbb{Z})$.

decorating the upper half-plane $\mathbb{R} \times \mathbb{R}_{\geq 0}$ with a toric boundary on the edge $\mathbb{R} \times \{0\}$ and a node at $(0, c)$, with the branch cut line $\{0\} \times [c, +\infty)$. By nodal slides, we can choose $c > 0$ freely without changing the symplectomorphism type of the space. This almost toric base diagram encodes a completely integrable system on $\mathbb{C}^2$ with a fibre at $(0, c)$ containing a focus-focus singularity and whose Arnol'd–Liouville fibre over the point $(0, a)$ is the Clifford torus $T_{\text{Cl}}(a)$. The set of regular values of the corresponding almost toric fibration is given by $B_{\text{reg}} = \mathbb{R} \times \mathbb{R}_{>0} \setminus \{(0, c)\}$.




FIGURE 4. The loop ξ which lifts to the Lagrangian loop $\Xi \in \pi_1(\mathcal{L}, T_{\text{Cl}})$.

Now let $\xi: S^1 \rightarrow \mathbb{R} \times \mathbb{R}_{\geq 0}$ be a loop with $\xi(0) = \xi(1) = (0, a)$ going once around the node at $(0, c)$ in the clockwise direction; see Figure 4. The precise shape of ξ is irrelevant; only its

homotopy class in B_{reg} preserving $\xi(0) = \xi(1) = (0, a)$ matters. The lift, as in Definition 4.1, of ξ yields an element $\Xi \in \pi_1(\mathcal{L}, T_{\text{Cl}})$. Recall that the almost toric base diagram allows us to pick action coordinates $F = (F_1, F_2)$ on the complement $\mathbb{R} \times \mathbb{R}_{\geq 0} \setminus (\{0\} \times [c, +\infty))$ of the branch cut. The branch cut tells us that these action coordinates are identified by the shear $(F_1, F_2) \mapsto (F_1, F_2 + F_1)$ when traversing the branch cut from right to left; see [Eva23, Chapter 7]. Now let $f_1, f_2 \in H_1(T_{\text{Cl}})$ be the classes defined by the closed orbits generated by the action coordinates F_1, F_2 . The coordinate F_1 is well defined, but since T_{Cl} is a fibre on the branch cut line, F_2 is defined only up to adding integer multiples of F_1 . We pick one such choice for F_2 to define f_2 . In that basis, the topological monodromy of the torus bundle over the image of ξ (traversed in the clockwise direction) is precisely the matrix in (20); see [Eva23, Lemma 7.2] (note that in Evans' book, the coordinates are rotated by $\frac{\pi}{2}$). This proves that the monodromy $\mathcal{M}(\Xi)$ is the one claimed in (21). To compute the flux of Ξ , we use (19). We claim that $I(\tilde{\xi}(1)) = I(\tilde{\xi}(0))$ for any lift $\tilde{\xi}$ of ξ to the universal cover $\tilde{B}_{\text{reg}}$. Indeed, by Proposition 4.2, the flux map I coincides with F up to integral affine transformations. Since the loop ξ (in the image of F) closes up, we find that $I(\tilde{\xi}(1)) = I(\tilde{\xi}(0))$ by continuity of I . $\square$

4.3. The flux-monodromy morphism for tori. Let us now prepare and prove Theorem 1.17. Besides the loop $\Xi \in \pi_1(\mathcal{L}, T_{\text{Cl}})$ constructed in Proposition 4.3, we define one other loop based at the Clifford torus $T_{\text{Cl}} \subset \mathbb{C}^2$ needed to generate $\Gamma(T_{\text{Cl}})$. Let $\phi_t \in \text{Ham } \mathbb{C}^2$ be a Hamiltonian isotopy satisfying $\phi_1(z_1, z_2) = (z_2, z_1)$ on a compact neighbourhood of T_{Cl} . For example, one can choose $\phi_t \in U(2)$ and use a cut-off function to make it compactly supported. Now set

$$\Sigma = \{\phi_t(T_{\text{Cl}})\}_{t \in [0,1]}. \quad (22)$$

Although ϕ_t itself is not a loop, its time one map ϕ_1 maps T_{Cl} to itself, and thus we obtain an element $\Sigma \in \pi_1(\mathcal{L}, T_{\text{Cl}})$. Its flux vanishes, since it is induced by a Hamiltonian flow.

In the basis defined by $e_1 = S^1 \times \{*\}$ and $e_2 = \{*\} \times S^1$, the monodromy of Σ swaps coordinates, $\mathcal{M}(\Sigma)e_1 = e_2$ and $\mathcal{M}(\Sigma)e_2 = e_1$. As it turns out, this is the only non-trivial monodromy that can be realized by a Lagrangian loop induced by a Hamiltonian isotopy.

Proposition 4.4. *Let Λ be a loop in $\mathcal{L}$ based at T_{Cl} . If Λ can be deformed to a loop in $\mathcal{L}\text{Ham}(T_{\text{Cl}})$, then, in the basis e_1, e_2 , we have that*

$$\mathcal{M}(\Lambda) \in \left\{ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \right\}. \quad (23)$$

Proof. The monodromy is an invariant up to deformation of loops. Every loop in $\mathcal{L}\text{Ham}(T_{\text{Cl}})$ is induced by a Hamiltonian isotopy; see Proposition A.1. The conclusion then follows from Chekanov's classification of possible monodromies of symplectomorphisms mapping a product torus to a product torus [Che96, Theorem 4.5]; see also later work by Yau [Yau09]. Heuristically, the elements $e_1, e_2 \in H_1(L)$ appear as boundaries of Maslov two holomorphic

disks, and hence the set $\{e_1, e_2\}$ must be preserved by any monodromy element realized by a Hamiltonian isotopy. $\square$

Corollary 4.5. *The loop Ξ from Proposition 4.3 cannot be deformed into $\mathcal{L}\text{Ham}(T_{\text{Cl}})$.*

This follows immediately from the fact that $\mathcal{M}(\Xi)$ has infinite order. Note that this is consistent with the fact that Ξ must traverse a wall of Chekanov tori — see Figure 4 — which are not Hamiltonian isotopic to the Clifford torus.

Let us now prove that the loops $\Xi, \Sigma \in \pi_1(\mathcal{L}, T_{\text{Cl}})$ generate the image of the flux-monodromy map. This follows from the elementary fact that Lagrangian isotopies preserve the Maslov class. First, note that we have used different bases to express the monodromy of Ξ and that of Σ . In what follows, we will use the basis $f_1, f_2 \in H_1(T_{\text{Cl}})$ defined in the proof of Proposition 4.3, since it yields cleaner expressions for monodromy. The basis $e_1, e_2 \in H_1(T_{\text{Cl}})$ is related to it by the shear $f_1 = e_1 - e_2, f_2 = e_2$. Hence, in the basis f_1, f_2 , we obtain

$$\mathcal{M}(\Sigma) = \begin{pmatrix} -1 & 1 \\ 0 & 1 \end{pmatrix}. \quad (24)$$

Similarly, the Maslov class is $m_{T_{\text{Cl}}} = 2(e_1^* + e_2^*) = 2f_2^* \in H^1(T_{\text{Cl}})$ under the natural isomorphism $H^2(\mathbb{C}^2, T_{\text{Cl}}) \cong H^1(T_{\text{Cl}})$.

Proposition 4.6. *Let T_{Cl} be a Clifford torus in $\mathbb{C}^2$. In the basis $f_1 = e_1 - e_2, f_2 = e_2$ of $H_1(T_{\text{Cl}})$, we have that*

$$\Gamma(T_{\text{Cl}}) = \left\{ \left(0, \begin{pmatrix} \pm 1 & k \\ 0 & 1 \end{pmatrix} \right) \mid k \in \mathbb{Z} \right\} \cong \mathbb{Z} \rtimes \mathbb{Z}_2. \quad (25)$$

In particular, the elements Ξ, Σ generate $\Gamma(T_{\text{Cl}})$.

Proof. By Corollary 2.6, the flux is always zero, so we only need to compute the possible monodromies. The matrices $\mathcal{M}(\Xi)$ and $\mathcal{M}(\Sigma)$ generate the group on the right-hand side of (25), thus proving the inclusion $\supseteq$. To prove that $\Gamma(T_{\text{Cl}})$ does not contain any other elements, let

$$M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{GL}(2; \mathbb{Z}),$$

be a monodromy matrix expressed in the basis f_1, f_2 . Since the Maslov index is preserved by Lagrangian isotopies, we find that $m_{T_{\text{Cl}}} M = m_{T_{\text{Cl}}}$ for $m_{T_{\text{Cl}}} = (0, 2)$. It follows that $c = 0$ and $d = 1$. Since $M \in \text{GL}(2; \mathbb{Z})$, we find $ad - bc = \pm 1$, so that $a = \pm 1$. This proves the inclusion $\subseteq$. $\square$

We can now prove the first part of Theorem 1.17, i.e. everything but the existence of a section of FM.

Proof of the 1st part of Theorem 1.17. By [DRGI16], L is Lagrangian isotopic to T_{Cl} . Let Π be a path in $\mathcal{L}$ starting at T_{Cl} and ending at L . This fixes a preferred isomorphism

$\pi_1(\mathcal{L}, T_{\text{Cl}}) \cong \pi_1(\mathcal{L}, L)$, namely the one given by $\Lambda \mapsto \Pi^- * \Lambda * \Pi$. On $H_1(L)$, we fix the basis $g_1 = (\phi^\Pi)_* f_1, g_2 = (\phi^\Pi)_* f_2$, where ϕ^Π is defined as in Section 2.2. Now let $\Lambda' = \Pi^- * \Lambda * \Pi \in \pi_1(\mathcal{L}, L)$, and suppose that $\text{Flux } \Pi = (x, y) \in H^1(T_{\text{Cl}}; \mathbb{R}) \cong \mathbb{R}^2$ in the basis f_1^*, f_2^* . Then the change-of-basepoint formula from Proposition 2.4 shows that there is $k \in \mathbb{Z}$ such that

$$\begin{aligned} \text{Flux}(\Lambda') &= \mathcal{M}(\Lambda)^* \text{Flux } \Pi + \text{Flux } \Lambda - \text{Flux } \Pi \\ &= \begin{pmatrix} \pm 1 & 0 \\ k & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} - \begin{pmatrix} x \\ y \end{pmatrix} = x \begin{pmatrix} \pm 1 - 1 \\ k \end{pmatrix} \end{aligned}$$

and

$$\mathcal{M}(\Lambda') = \begin{pmatrix} \pm 1 & k \\ 0 & 1 \end{pmatrix} \quad (26)$$

in the basis $g_1, g_2 \in H_1(L)$ and its dual basis. Since $\Lambda \in \pi_1(\mathcal{L}, T_{\text{Cl}})$, we have used Proposition 4.6 to determine $\mathcal{M}(\Lambda)$ and to show that $\text{Flux}(\Lambda) = 0$. $\square$

Interestingly, we find the following characterization of monotone tori in $\mathbb{C}^2$ in terms of their flux-monodromy group, as a consequence of Theorem 1.17.

Corollary 4.7. *A Lagrangian torus $L \subset \mathbb{C}^2$ is monotone if and only if $\Gamma(L)$ contains a non-trivial element of the form $(0, M)$.*

Proof. Let $L \subset \mathbb{C}^2$ be monotone with $[\omega_0]_L = \frac{C}{2} m_L$ for $C > 0$. Let Π be a path of Lagrangians starting at $T_{\text{Cl}} = T(1, 1)$ and ending at L . As in the proof of Theorem 1.17, we define the basis $g_i = \phi^\Pi_* f_i \in H_1(L)$ for $i \in \{1, 2\}$. As we have seen in the proof of Proposition 4.6, the Maslov class is given by $m_{T_{\text{Cl}}} = 2f_2^*$ and thus $m_L = 2g_2^*$. Since $(\mathbb{C}^2, \omega_0 = d\lambda)$ is exact, we use monotonicity to find $[\lambda]_L = Cg_2^*$ and $[\lambda]_{T_{\text{Cl}}} = f_2^*$ in case of the Clifford torus. The area identity (11) for the path Π yields

$$\text{Flux } \Pi = (\phi^\Pi)^* [\lambda]_L - [\lambda]_{T_{\text{Cl}}} = (\phi^\Pi)^* Cg_2^* - f_2^* = (C - 1)f_2^*.$$

In other words, we have $\text{Flux } \Pi = (0, C - 1)$ in the basis f_1^*, f_2^* . Since $x = 0$, using the notation of Theorem 1.17, the flux-monodromy group $\Gamma(L)$ contains elements of the form $(0, M)$ with $M \neq \text{id}$ by Theorem 1.17. Conversely, suppose that there is a loop $\Lambda \in \pi_1(\mathcal{L}, L)$ with $\text{FM}(\Lambda) = (0, M)$ for $M \neq \text{id}$. We again use the area identity to find $M^*[\lambda]_L - [\lambda]_L = 0$, meaning that M^* fixes the area class. Since it fixes the Maslov class, too, this implies that $M = \text{id}$ if the area and Maslov classes are linearly independent. $\square$

4.4. A section for the flux-monodromy morphism. We now complete the proof of Theorem 1.17. More precisely, we prove the following stronger statement.

Theorem 4.8. *Let L be a Lagrangian torus in $\mathbb{C}^2$. The flux-monodromy morphism $\text{FM} : \pi_1(\mathcal{L}) \rightarrow \Gamma(L)$ admits a section, i.e. there is a morphism $s : \Gamma(L) \rightarrow \pi_1(\mathcal{L})$ such that $\text{FM} \circ s = \text{id}$.*

The proof follows from a detailed understanding of the generators Ξ, Σ of $\Gamma(L)$ above and will take up the remainder of this subsection.

We start with an explicit description of the $\mathbb{Z}_2$ factor.

Lemma 4.9. *There is a Hamiltonian Lagrangian loop exchanging the two factors of $T_{\text{Cl}} = S^1(1) \times S^1(1)$ such that the following holds. Let $\Sigma \in \pi_1(\mathcal{L}, T_{\text{Cl}})$ be the class this loop represents. We may choose the Hamiltonian isotopy so that Σ^2 is trivial in $\mathcal{L}\text{Ham}(T_{\text{Cl}})$ — and thus a fortiori in $\mathcal{L}(T_{\text{Cl}})$.*

Proof. Note that the standard action of $U(2)$ on $\mathbb{C}^2$ is Hamiltonian. Therefore, if $A : [0, 1] \rightarrow U(2)$ is a path of matrices starting at the identity and with

$$A(1) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

then we can realize Σ as the Hamiltonian isotopy $t \mapsto A(t) \cdot T_{\text{Cl}}$.

Note that $\det A^2 : [0, 1] \rightarrow S^1$ is a loop and that Σ^2 is represented by $t \mapsto A(t)^2 \cdot T_{\text{Cl}}$. Since the determinant induces an isomorphism $\pi_1(U(2)) \xrightarrow{\sim} \pi_1(S^1) = \mathbb{Z}$, the loop A^2 in $U(2)$ must be homotopic to the loop $t \mapsto \text{diag}(e^{2\pi k i t}, 1)$ for some $k \in \mathbb{Z}$. In particular, $t \mapsto A(t)^2 \cdot T_{\text{Cl}}$ is homotopic to the constant loop at T_{Cl} through Hamiltonian isotopies. $\square$

Remark 4.10. As we shall see below, the loop $t \mapsto \text{diag}(e^{2\pi i t}, 1) \cdot T_{\text{Cl}}$ is not contractible as a loop of Lagrangian immersions — and thus *a fortiori* in $\tilde{\mathcal{L}}$. In fact, it represents an element of infinite order in the fundamental group of that space. This shows an important difference between the topology of unparametrized Lagrangians and that of parametrized ones. $\bullet$

The previous lemma gives a section from the $\mathbb{Z}_2$ factor of $\Gamma(T_{\text{Cl}})$. Moreover, the fact that $\mathcal{M}(\Xi)$ has infinite order in the monodromy group of T_{Cl} ensures that Ξ also has infinite order in $\pi_1(\mathcal{L}, T_{\text{Cl}})$. Therefore, Theorem 4.8 follows directly from the following proposition.

Proposition 4.11. *The loop $\Xi\Sigma\Xi\Sigma^{-1}$ is contractible in $\mathcal{L}$.*

Proof. Let $t \mapsto A(t) \in U(2)$ be the path of unitary matrices realizing the loop Σ , as in the proof of Lemma 4.9. We fix an isotopy $\{\phi^t\} \subset \text{Diff}(\mathbb{C}^2)$ realizing Ξ , i.e. $\Xi = \{\phi^t(T_{\text{Cl}})\}$. We define parametrizations $I = \{\iota_t\}$ and $I' = \{\iota'_t\}$ of Ξ by setting respectively $\iota_t = \phi^t \circ \iota$ with $\iota_0 = \iota$, the standard parametrization of the Clifford torus, and $\iota'_t = \phi^t \circ \iota'_0$ with $\iota'_0 = A(1) \circ \iota_1$. Notice that $\iota'_t = \phi^t \circ A(1) \circ \phi^1 \circ \iota$ and that $\Xi\Sigma\Xi\Sigma^{-1}$ is parametrized by the concatenation

$$\{\iota_t\} \# \{A(t) \circ \iota_1\} \# \{\iota'_t\} \# \{A(t)^{-1} \circ \iota'_1\}.$$

Claim 4.12. This path is homotopic relative to endpoints to $\{\iota_t\} \# \{A(1) \circ \iota'_t\}$ in $\tilde{\mathcal{L}}$.

Claim 4.13. The isotopy $\{\phi^t\}$ may be chosen so that $A(1) \circ \phi^t = \phi^{1-t} \circ (\phi^1)^{-1} \circ A(1)$.

Assuming both claims for now, the proposition easily follows. Indeed,

$$A(1) \circ \iota'_t = \phi^{1-t} \circ (\phi^1)^{-1} \circ A(1) \circ A(1) \circ \phi^1 \circ \iota = \phi^{1-t} \circ \iota = \iota_{1-t},$$

so that $\{A(1) \circ \iota'_t\}$ is simply the path $\{\iota_t\}$ backwards. $\square$

Proof of Claim 4.12. First, note that if $I = \{\iota_t\}$ is the parametrization of a Lagrangian isotopy and $\{\psi_t\}$ is a symplectic isotopy starting at the identity, then

$$\{\iota_t\} \# \{\psi_t \circ \iota_1\} \approx \{\psi_t \circ \iota_t\} \approx \{\psi_t \circ \iota_0\} \# \{\psi_1 \circ \iota_t\}.$$

To see this, we use an old trick. Define $F : [0, 1]^2 \rightarrow \tilde{\mathcal{L}}$ by $F(s, t) = \psi_t \iota_s$. Going from $(0, 0)$ to $(1, 1)$ in the square via the lower edges gives the first path, via the diagonal the second path, and via the upper edges the last path; see Figure 5. Then, any homotopy in the square between these paths defines a homotopy in $\tilde{\mathcal{L}}$ between the above paths.




FIGURE 5. The square realizing the homotopy.

To conclude, we simply apply these homotopies to our setting:

$$\begin{aligned} \{\iota_t\} \# \{A(t) \circ \iota_1\} \# \{\iota'_t\} \# \{A(t)^{-1} \circ \iota'_1\} &\approx \{\iota_t\} \# \{A(t) \circ \iota_1\} \# \{A(t)^{-1} \circ \iota'_0\} \# \{A(1) \circ \iota'_t\} \\ &= \{\iota_t\} \# \{A(t) \circ \iota_1\} \# \{A(t)^{-1} \circ A(1) \circ \iota_1\} \# \{A(1) \circ \iota'_t\} \\ &\approx \{\iota_t\} \# \{A(t)^{-1} \circ A(t) \circ \iota_1\} \# \{A(1) \circ \iota'_t\} \\ &= \{\iota_t\} \# \{\iota_1\} \# \{A(1) \circ \iota'_t\} \\ &\approx \{\iota_t\} \# \{A(1) \circ \iota'_t\} \end{aligned}$$

Here, we have made use of the fact that $A(1)^{-1} = A(1)$ and $\iota'_0 = A(1) \circ \iota_1$. $\square$

Proof of Claim 4.13. Here, we use the explicit *Auroux (integrable) system* underlying the almost toric base diagram depicted in Figure 4. It is explicitly given by

$$F : \mathbb{C}^2 \rightarrow \mathbb{R}^2, \quad (z_1, z_2) \mapsto (\pi|z_2|^2 - \pi|z_1|^2, |z_1 z_2 - c|^2)$$

for some $c > 0$. It has image $\mathbb{R} \times \mathbb{R}_{\geq 0}$, and computing action coordinates on the complement of the ray $\{0\} \times [c^2, +\infty)$ yields the same diagram as in Figure 4; see [Eva23, §7.1] for more details. This integrable system has the following properties:

- (1) For all $d > c$, the fibre $F^{-1}(0, d^2)$ is Hamiltonian isotopic to the Clifford torus $T_{\text{Cl}}(a(d))$ for some $a(d) > 0$.
- (2) We have $F \circ A(1) = I \circ F$, where I denotes the involution $I(x, y) = (-x, y)$ on $\mathbb{R}^2$.

The second property is immediate by the definition of F . The first property follows from the fact that, under the (singular) reduction by the S^1 -action generated by $F_1 = \pi|z_2|^2 - \pi|z_1|^2$, both $T_{\text{Cl}}(a)$ and the fibre $F^{-1}(0, d^2)$ map to a circle enclosing the puncture of the reduced space. Setting $a(d) > 0$ equal to the symplectic area of the reduced circle of $F^{-1}(0, d^2)$ yields the claim.

The element Ξ can be viewed as the lift of a curve running through a circle centred at $(0, c^2)$ in the image of F in $\mathbb{R}^2$. From now on, we fix $c > 0$ such that $a(c) < 1$. This means that there is some $d > c$ with $F^{-1}(0, d^2)$ Hamiltonian isotopic to $T_{\text{Cl}}(1) \subset \mathbb{C}^2$. Therefore, up to conjugation by a Hamiltonian isotopy, we can assume that $F^{-1}(0, d^2) = T_{\text{Cl}}(1)$. Take X' to be the vector field on $\mathbb{R}^2$ whose flow $\{(\psi')^t\}_{t \in [0, 1]}$ is a full rotation around the singular value $(0, c^2)$ of the Auroux system F . Let $A_1 \Subset A_2$ be annuli centered at a point $(0, x)$ such that $(0, d^2) \in A_1$ and that A_2 is disjoint from $(0, c^2)$ and $\mathbb{R} \times \{0\}$. Take $\beta : \mathbb{R}^2 \rightarrow [0, 1]$ rotationally symmetric such that $\beta|_{A_1} \equiv 1$ and $\text{supp } \beta \subseteq A_2$. We then set $X := \beta X'$.

Since the flow $\{\psi^t\}$ of X is supported in the regular locus of F , we can lift it to a flow $\{\phi^t\}$ of $\mathbb{C}^2$. By construction, $\phi^t(0, d^2)$ is homotopic to the loop ξ of the proof of Proposition 4.3. In particular, $\phi^t(T_{\text{Cl}})$ is the loop Ξ . By construction, we have that $I \circ \psi^t = \psi^{-t} \circ I$. By point (2) above, $A(1)$ is a lift of the involution I on the base, so that we have $A(1) \circ \phi^t = \phi^{-t} \circ A(1)$. But since ϕ^t is obtained from integration by a time-independent vector field, we have that

$$(\phi^t)^{-1} = \phi^{-t} = \phi^{1-t} \circ (\phi^1)^{-1},$$

thus giving the desired relation. $\square$

4.5. Relation to immersions. We now relate our study of $\mathcal{L}$ to what happens in the space

$$\mathcal{L}_{\text{imm}} := \{f : \mathbb{T}^2 \looparrowright \mathbb{C}^2 \mid f^*\omega = 0\} / \text{Diff}(\mathbb{T}^2)$$

of unparametrized immersed Lagrangian tori in $\mathbb{C}^2$.

Note that there is an obvious inclusion $\mathcal{L} \subseteq \mathcal{L}_{\text{imm}}$. Likewise, whenever L is embedded, FM_L factors through $\pi_1(\mathcal{L}_{\text{imm}}, L)$. Therefore, Theorem 4.8 also gives a splitting $\pi_1(\mathcal{L}_{\text{imm}}, T_{\text{Cl}}) = N \rtimes (\mathbb{Z} \rtimes \mathbb{Z}_2)$. However, we can prove more.

Proposition 4.14. *The fundamental group of $\mathcal{L}_{\text{imm}}$ based at T_{Cl} is isomorphic to $\mathbb{Z} \oplus (\mathbb{Z} \rtimes \mathbb{Z}_2)$, i.e. $N = \mathbb{Z}$ sits in the centre of $\pi_1(\mathcal{L}_{\text{imm}}, T_{\text{Cl}})$.*

The proposition follows from Proposition 4.15 and Lemma 4.18 below.

Proposition 4.15. *The monodromy map $\mathcal{M}$ on $\pi_1(\mathcal{L}_{\text{imm}}, T_{\text{Cl}})$ has the same image as in the embedded case, and its kernel is infinite cyclic.*

Proof. Let $\tilde{\mathcal{L}}_{\text{imm}}$ be the space of Lagrangian immersions $\mathbb{T}^2 \looparrowright \mathbb{C}^2$. The fibration $\text{Diff}(\mathbb{T}^2) \hookrightarrow \tilde{\mathcal{L}}_{\text{imm}} \twoheadrightarrow \mathcal{L}_{\text{imm}}$ induces an exact sequence

$$\pi_1(\mathrm{Diff}(\mathbb{T}^2)) \xrightarrow{i_1} \pi_1(\tilde{\mathcal{L}}_{imm}) \xrightarrow{j_1} \pi_1(\mathcal{L}_{imm}) \xrightarrow{\partial} \pi_0(\mathrm{Diff}(\mathbb{T}^2)) \xrightarrow{i_0} \pi_0(\tilde{\mathcal{L}}_{imm}).$$

Since $\mathrm{Diff}(\mathbb{T}^2)$ deformation retracts onto $\mathbb{T}^2 \times \mathrm{GL}(2; \mathbb{Z})$ [Gra73], we have that $\pi_0(\mathrm{Diff}(\mathbb{T}^2)) = \mathrm{GL}(2; \mathbb{Z})$ and $\pi_1(\mathrm{Diff}(\mathbb{T}^2)) = \mathbb{Z}^2$.

Likewise, the Gromov–Lees h -principle [Gro70, Lee76] implies that $\tilde{\mathcal{L}}_{imm}$ is homotopy equivalent to $C^\infty(\mathbb{T}^2, \mathrm{U}(2))$. Therefore, we have that

$$\pi_0(\tilde{\mathcal{L}}_{imm}) = [\mathbb{T}^2, \mathrm{U}(2)] = [\mathbb{T}^2, S^1] = H^1(\mathbb{T}^2; \mathbb{Z}) = \mathbb{Z}^2$$

and

$$\begin{aligned} \pi_1(\tilde{\mathcal{L}}_{imm}) &= \pi_1(\Lambda^2 \mathrm{U}(2)) = \pi_1(\Lambda \mathrm{U}(2)) \oplus \pi_2(\Lambda \mathrm{U}(2)) \\ &= \pi_1(\mathrm{U}(2)) \oplus \pi_2(\mathrm{U}(2))^2 \oplus \pi_3(\mathrm{U}(2)) \\ &= \mathbb{Z}^2, \end{aligned}$$

where $[\cdot, \cdot]$ denotes homotopy of unbased maps and ΛX the free loop space of the space X . In the π_0 computation, we have made use of the fact that $\mathrm{U}(2) = S^1 \times S^3$ (topologically) and that S^1 is a $K(\mathbb{Z}, 1)$ -space. The equality $[\mathbb{T}^2, S^3] = 0$ follows from cellular approximation, because the 2-skeleton of S^3 may be chosen to be a point. In the π_1 computation, we have also used the adjunction $C^\infty(\mathbb{T}^2, \mathrm{U}(2)) = C^\infty(S^1, C^\infty(S^1, \mathrm{U}(2)))$ together with the fact that the fibration $\Omega X \rightarrow \Lambda X \rightarrow X$ admits a section via constant loops, so that the associated long exact sequence splits in each degree. Note that this only implies in general that $\pi_1(\Lambda X) = \pi_1(X) \ltimes \pi_1(\Omega X) = \pi_1(X) \ltimes \pi_2(X)$. However, since X is a topological group in our case, its fundamental group must be abelian, so that the semidirect product is trivial, i.e. it is direct. Likewise, this ensures that all connected components of ΛX are homeomorphic, so that the splitting holds whatever base point we choose for $\pi_1(\Lambda X)$.

It is well known that the identification of $\pi_0(\tilde{\mathcal{L}}_{imm})$ with $H^1(\mathbb{T}^2; \mathbb{Z})$ sends a Lagrangian immersion to (half of) its Maslov class — see e.g. [Vit90]. Through this identification, i_0 must thus send a matrix $M \in \mathrm{GL}(2; \mathbb{Z})$ to half the Maslov class $\frac{1}{2}m^*M$ of the embedding $i_0 \circ M$, where $m \in H^1(\mathbb{T}^2; \mathbb{Z})$ is the Maslov class of the standard embedding of the Clifford torus.

Moreover, recall that the homotopy long exact sequence of a fibration $F \hookrightarrow E \rightarrow B$ comes from the long exact sequence of the pair (E, F) by identifying $\pi_k(E, F)$ with $\pi_k(B)$. Under this identification, the map $\pi_k(B) \rightarrow \pi_{k-1}(F)$ corresponds to the usual boundary map $\pi_k(E, F) \rightarrow \pi_{k-1}(F)$. Therefore, we may think of a loop in $\mathcal{L}_{imm}$ as a path in $\tilde{\mathcal{L}}_{imm}$ with endpoints in the fibre over the Clifford torus, with ∂ sending this path to its endpoints. But under the identification of the fibre over the Clifford torus with $\mathrm{Diff}(\mathbb{T}^2)$, this means that ∂ is precisely the monodromy $\mathcal{M}$.

To summarize the previous steps, the monodromy map has the same image in the immersed case as in the embedded one, and there is an exact sequence

$$\mathbb{Z}^2 \xrightarrow{i_1} \mathbb{Z}^2 \xrightarrow{j_1} \pi_1(\mathcal{L}_{imm}) \xrightarrow{\mathcal{M}} \mathbb{Z} \rtimes \mathbb{Z}_2 \longrightarrow 0.$$

Following the identifications above, we see that the generators of $\mathbb{Z}^2 = \pi_1(\text{Diff}(\mathbb{T}^2))$ are sent by i_1 to the classes represented by $t \mapsto ((\theta, \phi) \mapsto \text{diag}(e^{2\pi i(t+\theta)}, e^{2\pi i\phi}))$ and $t \mapsto ((\theta, \phi) \mapsto \text{diag}(e^{2\pi i\theta}, e^{2\pi i(t+\phi)}))$ in the fundamental group $C^\infty(\mathbb{T}^2, \text{U}(2))$ based at the usual embedding of T_{Cl} . Just as noticed in the proof of Lemma 4.9, both these loops are homotopic and generate the $\pi_1(\text{U}(2)) = \mathbb{Z}$ factor of $\pi_1(\tilde{\mathcal{L}}_{imm}) = \mathbb{Z}^2$. Equivalently, as maps $\mathbb{T}^3 \rightarrow \text{U}(2) = S^1 \times S^3$, both maps pull back $d\theta$ on S^1 to dt and have zero degree on S^3 — it is not surjective — so that they both generate the first $\mathbb{Z}$ factor of $[\mathbb{T}^3, \text{U}(2)]/[\mathbb{T}^2, \text{U}(2)]$. Therefore, j_1 has kernel and image isomorphic to $\mathbb{Z}$, which proves the statement. $\square$

Remark 4.16. The fact that the monodromy has the same image in both the immersed and embedded worlds already follows from the facts that the based point T_{Cl} is *embedded* so that the monodromy factors through $\pi_1(\mathcal{L}_{imm}, L)$, and that $\mathcal{M}(\pi_1(\mathcal{L}, T_{\text{Cl}}))$ is already as large as it could be by Proposition 4.6. $\bullet$

As stated in Remark 4.10, we see from this proof that, seen as a parametrized loop, Σ has infinite order in $\pi_1(\tilde{\mathcal{L}}_{imm})$. However, that group also has another generator that survives in $\pi_1(\mathcal{L}_{imm})$; we denote it by Υ .

Remark 4.17. Using the same approach, we can actually compute all homotopy groups of $\mathcal{L}_{imm}$ — or, at least, reduce it to another hard problem:

$$\pi_k(\mathcal{L}_{imm}) \cong \begin{cases} \pi_k(S^3) \oplus \pi_{k+1}(S^3)^2 \oplus \pi_{k+2}(S^3) & \text{if } k \geq 3; \\ \mathbb{Z}^3 \oplus \mathbb{Z}_2 & \text{if } k = 2. \end{cases}$$

In particular, it is finite for $k \geq 4$.

Indeed, we have $\pi_k(\mathcal{L}_{imm}) = \pi_k(\tilde{\mathcal{L}}_{imm})$ for $k \geq 3$. For $k = 2$, we instead have a short exact sequence

$$0 \longrightarrow \pi_2(\tilde{\mathcal{L}}_{imm}) \longrightarrow \pi_2(\mathcal{L}_{imm}) \longrightarrow \mathbb{Z} \longrightarrow 0.$$

This is because $\pi_k(\text{Diff}(\mathbb{T}^2)) = \pi_k(\mathbb{T}^2) = 0$ for $k \geq 2$ and $\ker i_1 = \mathbb{Z}$. But the above sequence must split since $\mathbb{Z}$ is free abelian. The result then follows from the same computation as above, together with knowing the first four homotopy groups of S^3 . $\bullet$

Lemma 4.18. *The kernel of $\mathcal{M}$ in $\pi_1(\mathcal{L}_{imm})$ is in the centre of the group.*

Proof. From the proof of Proposition 4.15, we see that this kernel is precisely the image of the natural map $\pi_1(\tilde{\mathcal{L}}_{imm}) \rightarrow \pi_1(\mathcal{L}_{imm})$. Therefore, the action of $\pi_1(\mathcal{L}_{imm})$ on the kernel by conjugation comes from changing basepoints: If $F = \{f_t\}$ is a loop in $\tilde{\mathcal{L}}_{imm}$ based at $f = f_0$, Λ is a loop in $\mathcal{L}_{imm}$ based at $L = f(\mathbb{T}^2)$, and $G = \{g_t\}$ is a parametrization of Λ with $g_0 = f$, then

$$\Lambda \cdot j_1(F) \cdot \Lambda^{-1} = j'_1(G \# F \# \overline{G}),$$

where $\overline{G} = \{g_{1-t}\}$. But for this formula to make sense, we need to identify $\pi_1(\tilde{\mathcal{L}}_{imm}, g_1)$ with $\pi_1(\tilde{\mathcal{L}}_{imm}, f)$, hence the prime on the right-hand side above. However, since $\tilde{\mathcal{L}}_{imm}$ is homotopy equivalent to a topological group, there is a canonical way to do so: conjugation by any path from g_1 to f . Moreover, looking again at the proof of Proposition 4.15, we see that this isomorphism commutes with the identification $\pi_1(\tilde{\mathcal{L}}_{imm}) = \mathbb{Z}^2$. Therefore, by taking the path $\overline{G}$, we have that $j'_1(G \sharp F \sharp \overline{G}) = j_1(F)$, i.e. the action is trivial. $\square$

Remark 4.19. The h -principle for monotone Lagrangian immersions in $\mathbb{C}^n$ [EK15, Theorem D] implies that the connected component of $\mathcal{L}_{imm}$ containing the Clifford torus is weakly homotopically equivalent to the space of monotone immersed Lagrangians — see also the parametric h -principle Theorem B.5 below. In particular, this means that all loops in $\pi_1(\mathcal{L}_{imm}, T_{Cl})$ can be realized by *exact* Lagrangian homotopies, i.e. with flux vanishing identically along the isotopy. $\bullet$

4.6. Hamiltonian action on the fundamental group. For any Lagrangian $L \subseteq X$, there is a natural action of $\text{Ham}(X)$ on $\mathcal{L}(L)$ whose orbits are precisely the Hamiltonian orbits. In particular, this induces a map $\pi_1(\text{Ham}(X)) \rightarrow \pi_1(\mathcal{L}, L)$. Note that this map has image in $\ker \text{FM}_L$, so that they are not detected by the above methods.

Nonetheless, this map can be nontrivial: in the case $\mathbb{R}P^2 \subseteq \mathbb{C}P^2$, nontrivial loops in $\text{PU}(3)$ — which is a deformation retract of $\text{Ham}(\mathbb{C}P^2)$ [Gro85] — induce nontrivial loops in $\mathcal{L}$ (see Section 2.1 above for a review of this fact). We now show that this is not the case for many Lagrangian tori in 4-manifolds.

To state our result, we use the notion of ball-embeddability: a Lagrangian torus L is *ball-embeddable* if there are a symplectic embedding $e : B \hookrightarrow X$ of a ball B and a product torus $T \subseteq B$ such that $e(T) = L$.

Proposition 4.20. *Let X be diffeomorphic to either $S^2 \times S^2$ or $\mathbb{C}P^2 \sharp k\overline{\mathbb{C}P^2}$ for $k \leq 8$. Suppose that L is Lagrangian isotopic to a ball-embeddable Lagrangian. Then, the map $\pi_1(\text{Ham}(X)) \rightarrow \pi_1(\mathcal{L}, L)$ is trivial.*

Let $X = \mathbb{C}^2$ with its standard structure, then the result holds with $\text{Ham}(\mathbb{C}^2)$ replaced with $\text{Symp}_{cyl}(\mathbb{C}^2)$, the space of symplectomorphisms ψ of $\mathbb{C}^2$ such that $\psi^\lambda_0 = \lambda_0$ outside some compact set.*

From [DRGI16], every Lagrangian torus in $\mathbb{C}^2$, $\mathbb{C}P^2$ or the monotone $S^2 \times S^2$ is Lagrangian isotopic to a ball-embeddable one. Here, $\lambda_0 = \frac{1}{2} \sum_i q_i dp_i + p_i dq_i$ is the standard radial Liouville form.

Proof. We start with the case where X is closed. Note that, given a Hamiltonian loop $\{\phi^t\}$, a Lagrangian isotopy from L to L' defines a free homotopy from $\{\phi^t(L)\}$ to $\{\phi^t(L')\}$. Therefore, it suffices to prove the case where L is ball-embeddable. In fact, by a Lagrangian isotopy within the ball, we may suppose that L is the image of the Clifford torus under a symplectic embedding $e : B_r \hookrightarrow X$ for r arbitrarily small.

But for $r > 0$ small enough and X as above, Anjos, Li, Li, and Pinsonnault [ALLP24] showed that the space $\text{Emb}_\omega(B_r, X)$ of symplectic embeddings $B_r \hookrightarrow X$ is weakly homotopy equivalent to the space $\text{SF}(X)$ of symplectic frames of X — see [LP04, Pin08] for the original proofs for $X = S^2 \times S^2$ and $X = \mathbb{CP}^2$, respectively. Moreover, the weak homotopy equivalence is given by the evaluation of the embedding at $0 \in B_r$ and its differential at that point. But $\{\phi^t(e(0))\}$ is a trivial loop in X , since every loop in X coming from a loop in $\text{Ham}(X)$ is trivial (this is a direct consequence of Arnol'd conjecture). Therefore, $\{\phi^t(e)\}$ is homotopic in $\text{Emb}_\omega(B_r, X)$ to a loop $\{e \circ A_t\}$ for some loop of linear symplectic transformations $\{A_t\}$. Note that this induces a homotopy in $\mathcal{L}$ from $\{\phi^t(L)\}$ to $\{e(A_t \cdot T_{\text{Cl}})\}$. But the group of linear symplectic transformations deformation retracts onto $\text{U}(2)$, and the map $\pi_1(\text{U}(2)) \rightarrow \pi_1(\mathcal{L}, T_{\text{Cl}})$ is trivial: this is the content of the proof of Lemma 4.9. This thus proves the compact case.

The case $X = \mathbb{C}^2$ is proven analogously, since $\text{Symp}_{\text{cyl}}(\mathbb{C}^2)$ deformation retracts onto $\text{U}(2)$. This is a folkloric fact, but we sketch the argument and thank Emmy Murphy for pointing it out to us. Essentially, the map $\text{Symp}_{\text{cyl}}(\mathbb{C}^2) \rightarrow \text{Cont}(S^3)$ sending a symplectomorphism that is cylindrical at infinity to the contactomorphism it induces on the ideal boundary S^3 of $\mathbb{C}^2$ is a Serre fibration with fibre $\text{Symp}_c(\mathbb{C}^2)$. But $\text{Symp}_c(\mathbb{C}^2)$ is contractible by a classical result of Gromov [Gro85], so that the map is a homotopy equivalence. It follows from a multiparametric version of Eliashberg's classical result on the extendability of contactomorphisms of S^3 [Eli92] that $\text{Cont}(S^3)$ deformation retracts onto $\text{U}(2)$ — see [FMAP20] for another proof. $\square$

Remark 4.21. In terms of the map $\text{Emb}_\omega(B_r, X) \rightarrow \mathcal{L}$ defined by $e \mapsto e(T_{\text{Cl}})$, the above proof can be summarized as saying that

$$\begin{array}{ccc} \pi_1(\text{Emb}_\omega(B_r, X)) & \longrightarrow & \pi_1(\mathcal{L}) \\ \uparrow & \nearrow & \\ \pi_1(\text{Ham}(X)) & & \end{array}$$

commutes for r small enough and that the horizontal arrow is trivial. $\bullet$

4.7. Hausdorff-local path connectedness. In order to prove Theorem 1.13, we first show this important local estimate. We denote by $\mathbb{T}^2$ the 2-torus, equipped with the flat Riemannian metric, and let $D_r^* \mathbb{T}^2$ be its cotangent disk bundle of capacity $r > 0$, equipped with the canonical symplectic form.

Theorem 4.22. *Let $L \subseteq (D_r^* S^1)^2$ be a Lagrangian torus. Denote by $A_{\min}^2(L)$ the infimal area of a smooth Maslov-2 disk with boundary on L ; we set $A_{\min}^2(L) = 0$ if there are no such disks. We have that*

$$A_{\min}^2(L) \leq \frac{4r}{3}.$$

Proof. Denote by π_L the composition $L \rightarrow T^*\mathbb{T}^2 \rightarrow \mathbb{T}^2$ and by $(\pi_L)_*$ the map it induces on $H_1(\cdot; \mathbb{R})$.

Suppose that $\dim \ker(\pi_L)_* = 2$. Note that the Lagrangian embedding mapping $\mathbb{T}^2$ to the product torus $T(r, r) \subset (\mathbb{R}^4, \omega_0)$ extends to a symplectic embedding Ψ of $(D_r^*S^1)^2$ into the ball $B^4(4r)$ of capacity $4r$. By the hypothesis on $(\pi_L)_*$, this induces an isomorphism $\Psi_* : \pi_2((D_r^*S^1)^2, L) \rightarrow \pi_2(B^4(2r), \Psi(L))$ that preserves areas. By [CM18], $\Psi(L)$ bounds a Maslov-2 disk of area at most $\frac{4r}{3}$. Therefore, we have that $A_{\min}^2(L) \leq \frac{4r}{3}$.

Suppose now that $\dim \ker(\pi_L)_* = 1$. If L does not bound a Maslov-2 disk u of positive area, then there is nothing to prove. Therefore, we may suppose that it does. Note that such an element must generate $\pi_2((D_r^*S^1)^2, L) \cong \mathbb{Z}$, since u must be primitive by the Maslov condition. Furthermore, the symplectic embedding sending $(D_r^*S^1)^2$ to either $D_r^*S^1 \times B^2(2r)$ or $B^2(2r) \times D_r^*S^1$ (which extend the maps $\mathbb{T}^2 \rightarrow S^1 \times T(r)$ or $\mathbb{T}^2 \rightarrow T(r) \times S^1$, respectively) induces an isomorphism on relative π_2 — this is the argument in the proof of Conjecture C in [ACLS25]. Therefore, by the Chekanov bound on displacement energy [Che98], the area of the smallest J -disk with boundary on L must have area at most r . By the above discussion, this disk is some positive multiple of u . Therefore, we have that $A_{\min}^2(L) \leq r$.

Suppose finally that $\dim \ker(\pi_L)_* = 0$. Then, L does not bound any disk by the characterization lemma [ACLS25, Lemma 9], i.e. $A_{\min}^2(L) = 0$. $\square$

Let us now prove the Hausdorff-local path connectedness in Theorem 1.13.

Corollary 4.23. *Every Lagrangian torus L in $\mathbb{C}^2$ has a system of neighbourhoods $\{U_n\}$ with the following property. There is $n_0 \in \mathbb{N}$ such that, for all $n \geq n_0$, if L' is Hamiltonian isotopic to L in $\mathbb{C}^2$ and $L' \subseteq U_n$, then L' is Hamiltonian isotopic to L in U_n . In particular, $\mathcal{L}\text{Ham}(L)$ is locally path connected in the Hausdorff topology.*

Proof. Let $\Psi : D_{r_0}^*\mathbb{T}^2 \hookrightarrow \mathbb{C}^2$ be a Weinstein neighbourhood of L . We set $U_n = \Psi(D_{1/n}^*\mathbb{T}^2)$. Although this sequence may not be defined for a finite set of indices, it makes sense for large enough $n \in \mathbb{N}$. By [CM18], L bounds a Maslov-2 disk in $\mathbb{C}^2$ of positive area. Since Maslov-0 disks form a subspace of $\pi_2(\mathbb{C}^2, L)$ of rank 1, this means that the minimal positive area $A_{\min}^2(\mathbb{C}^2, L)$ of Maslov-2 disks in $\mathbb{C}^2$ with boundary on L is positive. Now let $n_0 \in \mathbb{N}$ be large enough so that $\frac{2}{n_0} < A_{\min}^2(\mathbb{C}^2, L)$. Take $n \geq n_0$, and suppose that $L' \subset U_n$ is Hamiltonian isotopic to L in $\mathbb{C}^2$. We need to prove that L' is Hamiltonian isotopic to L in U_n . Since $L' \subset U_n$, we have $\Psi^{-1}(L') \subset D_{1/n}^*\mathbb{T}^2$, and, by Theorem 4.22, we find that

$$A_{\min}^2(D_{1/n}^*\mathbb{T}^2, \Psi^{-1}(L')) \leq \frac{2}{n} \leq \frac{2}{n_0} < A_{\min}^2(\mathbb{C}^2, L). \quad (27)$$

Therefore, $\Psi^{-1}(L')$ does not bound any Maslov-2 disks of positive area in $D_{1/n}^*\mathbb{T}^2$. Otherwise, L' would bound a Maslov-2 disk of positive area at most $\frac{2}{n_0}$ in $\mathbb{C}^2$, which (27) does not allow for. Indeed, since L' is Hamiltonian isotopic to L , we have $A_{\min}^2(\mathbb{C}^2, L') = A_{\min}^2(\mathbb{C}^2, L)$. But since there are no Maslov-2 disks of positive area, the third point of Theorem 4.1.1 of [ALP94]

is violated. By the first point of this theorem, we find that $\Psi^{-1}(L')$ is homologous to the zero section in $D_{1/n}^* \mathbb{T}^2$.

This allows us to apply [DR19, Theorem B (1)] to find that $\Psi^{-1}(L')$ is Hamiltonian isotopic inside $D_{1/n}^* \mathbb{T}^2$ to the graph of a closed 1-form. Since cohomologous 1-forms have Hamiltonian isotopic graphs, $\Psi^{-1}(L')$ is Hamiltonian isotopic to some $\mathbb{T}^2 \times \{z\}$ in $D_{1/n}^* \mathbb{T}^2 = \mathbb{T}^2 \times B^2(1/n)$. In particular, L' is Hamiltonian isotopic in U_n to a Lagrangian whose C^1 distance (in the appropriate metric) to L is at most $\frac{1}{n}$. The result then follows from discreteness of the flux (Theorem 1.17) and the usual result on C^1 -local path connectedness (Proposition 1.10). $\square$

We can now finish the proof of Theorem 1.13, i.e. prove Hausdorff-closedness of $\mathcal{L}\text{Ham}(L)$ in $\mathcal{L}(L)$.

Proof of Theorem 1.13. Let $\{L_i\}$ be a sequence in $\mathcal{L}\text{Ham}(L)$ Hausdorff-converging to some Lagrangian torus K . For i large, L_i is in a Weinstein neighbourhood of K of size $r_i > 0$, and $r_i \rightarrow 0$. Therefore, for i large, we have that $A_{\min}^2(L_i, D_{r_i}^* \mathbb{T}^2) \leq 2r_i < A_{\min}^2(L_i, \mathbb{C}^2) = A_{\min}^2(L, \mathbb{C}^2)$, which means that L_i does not bound any Maslov-2 disk in the Weinstein neighbourhood of K . Just as in the proof of Corollary 4.23, we can thus assume that $L_i \rightarrow K$ in the C^1 topology. The result then follows from Theorem 1.17 and Proposition 1.10. $\square$

Remark 4.24. Given $\tau > 0$, a similar argument shows that the subspace of $\mathcal{L}$ formed by Lagrangians L with $A_{\min}^2(L) \geq \tau$ is Hausdorff-locally path connected and closed in $\mathcal{L}$.

However, we cannot expect the whole space $\mathcal{L}$ to have those properties. Indeed, take any submanifold N and a sequence of smooth curves $\{c_k : S^1 \rightarrow N\}$ C^0 -converging to a space-filling curve of N . By [LO95, Proposition A1.4], for each k , there is a Hamiltonian isotopy ϕ_k of $\mathbb{C}^2$ sending $S^1(|\lambda_0(c_k)|) \times \{0\}$ to $c_k(S^1)$. In particular, if $\{\epsilon_k\} \subseteq \mathbb{R}_{>0}$ converges to zero fast enough, ϕ_k sends $S^1(|\lambda_0(c_k)|) \times S^1(\epsilon_k)$ to a Lagrangian torus T_k such that $T_k \rightarrow N$ in the Hausdorff metric. However, T_k is not topologically isotopic to N in any small neighbourhood: the loop corresponding to the $S^1(\epsilon_k)$ factor is contractible in such a neighbourhood. $\bullet$

5. MONODROMY AND FLUX IN $\mathbb{CP}^n$

We now turn to complex projective spaces, where we are able to determine the monodromy group of the Clifford torus explicitly as well as (most of) its shape invariant. The computation of the monodromy group requires different techniques when $n = 2$, which we explain in § 5.1, where we prove Theorem 1.19. The case $n \geq 3$ is proven in § 5.2. Theorem 1.15 then follows directly from these results. The shape invariant is computed in § 5.3, when $n \geq 3$, and in § 5.4, when $n = 2$. This latter section is also where we prove Theorem 1.20.

5.1. Lagrangian monodromy in $\mathbb{CP}^2$. In this section, we prove Theorem 1.19 from the introduction.

First, we claim that we can restrict our attention to the Clifford torus T_{Cl} . Indeed, because any Lagrangian torus in $\mathbb{CP}^2$ is Lagrangian isotopic to T_{Cl} , by the change-of-basepoint formula (9), any two Lagrangian tori have isomorphic homological monodromy.

It is easy to see that the matrices

$$A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad M = \begin{pmatrix} 2 & 1 \\ -1 & 0 \end{pmatrix} \quad R = \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix} \quad (28)$$

belong to $\mathcal{M}(T_{\text{Cl}})$. Indeed, A and M are induced by the generators $\mathcal{M}(\Sigma)$ and $\mathcal{M}(\Xi)$ of $\mathbb{Z}_2 \ltimes \mathbb{Z}$ from Sections 4.2 and 4.3 in $\mathbb{C}^2$, while R comes from a rotation of the moment polytope of $\mathbb{CP}^2$. Note that A and R generate all coordinate permutations in homogeneous coordinates.

Remark 5.1. This shows that Corollary 4.7, which characterizes monotone Lagrangian tori in $\mathbb{C}^2$ in terms of their flux-monodromy group, fails in other spaces. The Hamiltonian isotopy of $\mathbb{CP}^2$ which swaps homogeneous coordinates induces an element $\Lambda \in \pi_1(\mathcal{L}, L)$ whose flux-monodromy is $(0, A)$, for any non-monotone toric fibre of the form $L = T(a, a) \subset \mathbb{CP}^2$. •

We denote the mod-3 reduction by $\rho : \text{GL}(2; \mathbb{Z}) \rightarrow \text{GL}(2; \mathbb{Z}/3\mathbb{Z})$, and write $\mathbb{1}$ for the vector $(1, 1)$. Here is the algebraic statement that is central to the proof of Theorem 1.19 for $\mathbb{CP}^2$.

Proposition 5.2. *The monodromy group $\mathcal{M}(T_{\text{Cl}})$ is isomorphic to a subgroup of $\text{GL}(2; \mathbb{Z})$ that can be described equivalently as*

- (1) the subgroup $G_0 = \langle A, R, M \rangle$ generated by A , R , and M ;
- (2) the congruence subgroup $G_1 = \rho^{-1}(F)$ where $F = \langle \rho(A), \rho(M) \rangle$;
- (3) the congruence subgroup $G_2 = \rho^{-1}(\text{Stab}(\mathbb{1})) = \{P \in \text{GL}(2; \mathbb{Z}) \mid \mathbb{1}P = \mathbb{1} \bmod 3\}$.

We first show that Theorem 1.19 immediately follows from the proposition.

Proof of Theorem 1.19. Since $\mathbb{CP}^2$ is H -monotone, Proposition 1.12 ensures that the flux set $F(T_{\text{Cl}})$ vanishes, so that the monodromy and flux-monodromy groups of T_{Cl} are isomorphic.

The first claim of Theorem 1.19 is simply (3) from the proposition above, with the notation H' from the introduction changed to G_2 for readability.

We now compute its index by using description (2). The matrix A has order 2, M has order 3, and $AMA = M^{-1}$. Hence, the group generated by A and M is isomorphic to the dihedral group of order 6. Since $\rho(A)$ and $\rho(M)$ are non-trivial in $\text{GL}(2; \mathbb{Z}/3\mathbb{Z})$, the same holds for F . Since G_1 is the congruence subgroup $\rho^{-1}(F)$, its index in $\text{GL}(2; \mathbb{Z})$ is

$$[\text{GL}(2; \mathbb{Z}) : G_1] = [\text{GL}(2; \mathbb{Z}/3\mathbb{Z}) : F].$$

Notice that there are 48 ordered pairs of non-proportional vectors in $(\mathbb{Z}/3\mathbb{Z})^2$. Any matrix formed of such a pair has determinant $\pm 1 \bmod 3$, and thus is invertible in $\text{M}_2(\mathbb{Z}/3\mathbb{Z})$. This shows that the order of $\text{GL}(2; \mathbb{Z}/3\mathbb{Z})$ is 48 and thus that $[\text{GL}(2; \mathbb{Z}) : G_1] = 8$.

The claim concerning the generators obviously follows from (1). □

The following two facts have been established in the proof above but are independent of Proposition 5.2: (i) F has order 6, and (ii) G_1 has index 8 in $\mathrm{GL}(2; \mathbb{Z})$. They will be used in the proof of said proposition, which we now explain.

Proof of Proposition 5.2. First, notice that the columns of any matrix in $\mathrm{GL}(2; \mathbb{Z}/3\mathbb{Z})$ that stabilizes $\mathbb{1}$ are of the form $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$, $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, or $\begin{pmatrix} 2 \\ 2 \end{pmatrix}$. Since any choice of two different columns among these three yields an invertible matrix in $\mathrm{GL}(2; \mathbb{Z}/3\mathbb{Z})$, $\mathrm{Stab}(\mathbb{1})$ is a subgroup of $\mathrm{GL}(2; \mathbb{Z}/3\mathbb{Z})$ of order 6. Because F is of order 6 and $F \subset \mathrm{Stab}(\mathbb{1})$, we deduce that these groups coincide, so that $G_1 = G_2$ and consequently $[\mathrm{GL}(2; \mathbb{Z}) : G_2] = 8$.

Second, since A , R , and M belong to $\mathcal{M}(L)$, and since any Lagrangian isotopy preserves the Maslov class, we have the inclusions $G_0 \subset \mathcal{M}(L) \subset G_2 = G_1$. Thus, we only need to show that G_0 and G_2 have the same index in $\mathrm{GL}(2; \mathbb{Z})$, *i.e.* that $[\mathrm{GL}(2; \mathbb{Z}) : G_0] = 8$. We compute this index in two steps.

Step 1. Since $G_0 \subset G_2$, the matrix $-I_2$ is not in G_0 . Let $H_0 = \langle -I_2, R, M \rangle \subset \mathrm{SL}(2; \mathbb{Z})$ and $K = \langle R, M \rangle$. Since A has order 2 and $AKA^{-1} \subset K$, $[G_0 : K] = 2$. For the same reason, $[H_0 : K] = 2$, and thus

$$[\mathrm{GL}(2; \mathbb{Z}) : G_0] = [\mathrm{GL}(2; \mathbb{Z}) : H_0] = 2 \cdot [\mathrm{SL}(2; \mathbb{Z}) : H_0]. \quad (29)$$

Step 2. The following composition of surjective morphisms preserves the index of H_0

$$\mathrm{SL}(2; \mathbb{Z}) \xrightarrow{\pi} \mathrm{PSL}(2; \mathbb{Z}) \xrightarrow{\rho} \mathrm{PSL}(2; \mathbb{Z}/3\mathbb{Z}).$$

Indeed, H_0 obviously contains the kernel $\{\pm I_2\}$ of π , so that

$$[\mathrm{SL}(2; \mathbb{Z}) : H_0] = [\mathrm{PSL}(2; \mathbb{Z}) : \pi(H_0)] \quad (30)$$

with $\pi(H_0) = \pi(K) = \langle \pi(R), \pi(M) \rangle$.

On the one hand, note that $MR = \begin{pmatrix} 1 & -3 \\ 0 & 1 \end{pmatrix}$, so that $\rho(\pi(M)) = \rho(\pi(R))^{-1}$, and that $\rho(\pi(R))$ has order 3. Hence, $\rho \circ \pi(K)$ is the cyclic group of order 3, generated by the image of R . On the other hand, the kernel of ρ is the principal congruence subgroup $\Gamma(3)$ of $\mathrm{PSL}(2; \mathbb{Z})$. We make the following claim, whose proof we postpone to Remark 5.4 below.

Claim 5.3. $\Gamma(3)$ is generated by the following elements of $\mathrm{PSL}(2; \mathbb{Z})$:

$$X = \begin{pmatrix} 1 & 3 \\ 0 & 1 \end{pmatrix}, \quad Y = \begin{pmatrix} -8 & 3 \\ -3 & 1 \end{pmatrix}, \quad \text{and } Z = \begin{pmatrix} 4 & -3 \\ 3 & -2 \end{pmatrix}. \quad (31)$$

It is easy to check that

$$X = (MR)^{-1}, \quad Y = XRXR^{-1} = (RM^2R)^{-1}, \quad \text{and } Z = R^2X^{-1}R = R^2MR^2.$$

Hence, $\ker \rho$ is a subgroup of $\pi(K)$, and

$$[\mathrm{PSL}(2; \mathbb{Z}) : \pi(H_0)] = [\mathrm{PSL}(2; \mathbb{Z}/3\mathbb{Z}) : \rho \circ \pi(K)] = \frac{12}{3} = 4,$$

which concludes the proof together with (29) and (30). $\square$

Remark 5.4. In [Kul91], Kulkarni establishes two procedures, which respectively

- (1) associate a Farey symbol to any principal congruence subgroup $\Gamma(N)$ [Kul91, §12-13],
- (2) determine an independent set of generators for any subgroup given by a Farey symbol [Kul91, Theorem 6.1].

Both procedures are implemented in Sage by the following two commands:

```
sage: f = FareySymbol(Gamma(3)); f
FareySymbol(Congruence Subgroup Gamma(3))
sage: f.generators()
[
[1 3]  [-8  3]  [4 -3]
[0 1], [-3  1], [3 -2]
]
```

These define $\mathbf{f}$ as the principal congruence subgroup $\Gamma(3)$ and give the set of generators (31). •

Remark 5.5. The main steps of the proof of Proposition 5.2 above were suggested by Yves Benoist, except for the last one. Indeed, in order to compute the index of H_0 in $\mathrm{SL}(2; \mathbb{Z})$ (Step 2 above), he suggested using the action of $\mathrm{PSL}(2; \mathbb{Z})$ on the Poincaré (hyperbolic) half-plane $\mathbb{H} = \{z \in \mathbb{C} \mid \mathrm{im}(z) > 0\}$ by Möbius transformations:

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \mapsto \left[h_A : z \mapsto \frac{az + b}{cz + d} \right]. \quad (32)$$

Because $\pi(H_0)$ acts by isometries, the index $[\mathrm{SL}(2; \mathbb{Z}) : H_0]$ is given by the ratio of the hyperbolic area of the respective fundamental domains. It is well-known that a fundamental domain of the action of $\mathrm{PSL}(2; \mathbb{Z})$ is the hyperbolic triangle of vertices $x_1 = e^{i\pi/3}$, ∞ , $x_2 = e^{2i\pi/3}$, whose area is $\frac{\pi}{3}$.

Because the Möbius transformation h_R associated to R maps the geodesic joining x_1 to ∞ to that joining 0 to x_1 , and h_M maps the geodesic joining -1 to 0 to that joining -1 to ∞ , the geodesic polygon of vertices ∞ , -1 , 0, and x_1 is a fundamental domain of the action of $\pi(H_0)$. The Gauss–Bonnet formula then gives that its area is $\frac{4\pi}{3}$, so that $[\mathrm{SL}(2; \mathbb{Z}) : H_0] = 4$. •

5.2. Monodromy in $\mathbb{CP}^n$. Recall from §1.4 the definition of the *homological Lagrangian monodromy* group

$$\mathcal{M}^{\mathbb{Z}}(L) = \{(\phi^\Lambda)_* \in \mathrm{Aut}(H_1(L; \mathbb{Z})) \mid \Lambda \in \pi_1(\mathcal{L}, L)\}$$

and the *extended homological Lagrangian monodromy* group

$$\widehat{\mathcal{M}}^{\mathbb{Z}}(L) = \{\Phi^\Lambda \in \mathrm{Aut}(H_2(X, L; \mathbb{Z})) \mid \Lambda \in \pi_1(\mathcal{L}, L)\}.$$

By $\widehat{\mathcal{M}}_{\mathrm{Ham}}^{\mathbb{Z}}(L)$ we denote the corresponding subgroup generated by Hamiltonian isotopies, which we call the *extended (homological) Hamiltonian monodromy group*.

Remark 5.6. Since $\partial : H_2(X, L; \mathbb{Z}) \rightarrow H_1(L; \mathbb{Z})$ is surjective in this context, the extended homological monodromy group $\widehat{\mathcal{M}}^{\mathbb{Z}}(L)$ fully determines the homological monodromy group $\mathcal{M}^{\mathbb{Z}}(L)$. How much this determines the monodromy group $\mathcal{M}(L) \subset \text{MCG}(L)$ depends on n . Indeed, it is known that $\text{MCG}(\mathbb{T}^n)$ is isomorphic to $\text{Aut } H_1(\mathbb{T}^n; \mathbb{Z})$ when $n = 3$ [Wal68], but not when $n \geq 5$ [Hat78] — the case $n = 4$ remains open. Therefore, when $n \geq 4$, our computations may not fully determine $\mathcal{M}(L)$. •

Since we will only consider *homological* monodromies below, we drop the superscript $\mathbb{Z}$ from the notation throughout this subsection.

Any element $\Psi \in \widehat{\mathcal{M}}$ preserves the Maslov class and acts by the identity on ambient elements, i.e. on the image $J = \text{im}(j : H_2(X) \rightarrow H_2(X, L))$. Question 1.14 asks whether $\widehat{\mathcal{M}}$ contains precisely *all* such elements. The goal of this section is to answer the question positively for tori that are Lagrangian isotopic to the Clifford torus in $\mathbb{CP}^n$.

Moreover, we determine the corresponding homological Lagrangian monodromy group. Let $\mathbb{1} = (1, \dots, 1)$, and define the following subgroups of $\text{GL}(n; \mathbb{Z})$:

$$H = \{A \in \text{GL}(n; \mathbb{Z}) \mid \mathbb{1}A = \mathbb{1}\}, \quad (33)$$

$$H' = \{A \in \text{GL}(n; \mathbb{Z}) \mid \mathbb{1}A \equiv \mathbb{1} \pmod{n+1}\}. \quad (34)$$

We obviously have $H \subset H'$. The group H is isomorphic to $\mathbb{Z}^{n-1} \rtimes \text{GL}(n-1; \mathbb{Z})$. The group H' is the preimage of the mod- $(n+1)$ reduction of H and has finite index in $\text{GL}(n; \mathbb{Z})$.

Theorem 5.7. *Let $n \geq 3$. The answer to Question 1.14 is positive for*

- (1) *any Lagrangian torus $L \subset \mathbb{C}^n$ that is Lagrangian isotopic to a product torus T_0 ;*
- (2) *any Lagrangian torus $L \subset \mathbb{CP}^n$ that is Lagrangian isotopic to a toric fibre.*

More precisely, in the natural bases induced by $T_0 \subset \mathbb{C}^n$, and $T_{\text{Cl}} \subset \mathbb{CP}^n$, the corresponding (non-extended) homological Lagrangian monodromy groups are given by

$$\mathcal{M}(T_0) = H, \quad \mathcal{M}(T_{\text{Cl}}) = H'.$$

For the case of $T_0 \subset \mathbb{C}^n$, this result goes back to Chekanov [Che96, Proposition 5.1 (a)]. We reprove it for the reader's convenience and because our proof of (2) will be based on it. In both cases, the strategy of the proof is as follows. Let $L_0, L_1 \subset X$ be Lagrangian isotopic Lagrangians. Then,

$$\widehat{\mathcal{M}}_{\text{Ham}}(L_1) \subset \widehat{\mathcal{M}}(L_1) \quad \text{and} \quad \Phi^{\Pi} \widehat{\mathcal{M}}_{\text{Ham}}(L_0) (\Phi^{\Pi})^{-1} \subset \widehat{\mathcal{M}}(L_1), \quad (35)$$

where $\Pi = \{L_t\}_{t \in [0,1]}$ is a Lagrangian isotopy from L_0 to L_1 and $\Phi^{\Pi} : H_2(X, L_0) \rightarrow H_2(X, L_1)$ denotes the isomorphism it induces. In both cases appearing in Theorem 5.7, there is a convenient choice of L_0 and L_1 such that the union of the two subgroups (35) generates all elements preserving the Maslov class and acting by the identity on J .

Before proving Theorem 5.7, we first need some algebraic preparations. These are surely known to experts on discrete groups, but we were unable to find a suitable reference.

Lemma 5.8. *Let $n \geq 3$. The group H' is generated by the set of all permutation matrices, together with the matrices*

$$S_1 = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 \\ 1 & 1 & 0 & \dots & 0 \\ -1 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \end{pmatrix}, \quad T = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 \\ n+1 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \end{pmatrix}.$$

Similarly, the group H is generated by simply adding S_1 to the set of all permutations.

The proof works by a modification of the Gauß algorithm. For the full group $\mathrm{GL}(n; \mathbb{Z})$, this algorithm shows that a set of generators is given by all permutations and an elementary shear, i.e. a matrix differing from the identity by adding a 1 somewhere off the diagonal. It works by transforming a given matrix in the group into an upper triangular one by taking repeated linear combinations of rows and permuting rows and columns. This corresponds to multiplying by permutation matrices and the elementary shear and thus proves that these generate all of $\mathrm{GL}(n; \mathbb{Z})$.

Proof of Lemma 5.8. Let $K \subset H$ be the subgroup generated by permutations and S_1 . Similarly, let $K' \subset H'$ be the subgroup generated by permutations, S_1 , and T . Our goal is to show that $K = H$, and $K' = H'$ by proving that any element in H and H' , respectively, can be reduced to the identity $I \in \mathrm{GL}(n; \mathbb{Z})$ through multiplication by elements in K and K' , respectively. Let us denote by $E_{i,j}$ the elementary matrix whose only non-zero entry is a 1 at the intersection of the i -th row and the j -th column. Now let $l, i, j \in \{1, \dots, n\}$ be pairwise distinct indices. Throughout the proof, we will denote

$$S_{l,i,j} = I + E_{il} - E_{jl}.$$

These matrices are obviously in K, K' and satisfy $S_{l,i,j}^k = I + kE_{il} - kE_{jl}$. For any $(n \times n)$ -matrix, we use the following notation

$$B = \begin{pmatrix} R_1 \\ \vdots \\ R_n \end{pmatrix} = (C_1 \dots C_n)$$

for its rows and columns. Left-multiplication by $S_{l,i,j}^k$ yields the new rows R'_i given by

$$R'_i = R_i + kR_l, \quad R'_j = R_j - kR_l, \quad \text{and} \quad R'_h = R_h \text{ for all } h \notin \{i, j\}. \quad (36)$$

Now let $B \in H' \supset H$. We use the transformation on rows to show the following.

Step 1. We can bring B to the form

$$B' = \begin{pmatrix} 1 & c \\ 0 & D \end{pmatrix},$$

for $c \in \mathbb{Z}^{n-1}$ and $D \in \mathrm{GL}(n-1; \mathbb{Z})$, by multiplying B with elements in

- (1) K if $B \in H$;
- (2) K' if $B \in H'$.

To that end, denote (the transpose of) the first column of $B = B^{(0)}$ by $(a_1^{(0)}, \dots, a_n^{(0)}) \in \mathbb{Z}^n$. Up to applying transformations of the form (36) and permutations, we can assume that $a_1^{(0)} \in \mathbb{Z}_{>0}$. The idea is to subsequently multiply B on the left by elements of the form $S_{l,i,j}^k$ such that the first entry in the first column is positive and strictly decreasing with each step of the algorithm. This means that it will be equal to 1 after a finite number of steps.

Here is the inductive algorithm. Let $B^{(m)}$ be a matrix with (the transpose of) the first column equal to $(a_1^{(m)}, \dots, a_n^{(m)}) \in \mathbb{Z}^n$ with $a_1^{(m)} \in \mathbb{Z}_{>0}$. If $a_1^{(m)} = 1$, then terminate the algorithm. If not, then note that multiplying $B^{(m)}$ on the left by $S_{1,i,j}^k$ followed by a permutation of the j -th and first columns yield a matrix $B^{(m+1)}$ with first row $a^{(m+1)}$ given by

$$\begin{aligned} a_1^{(m+1)} &= a_j^{(m)} - ka_1^{(m)} \\ a_i^{(m+1)} &= a_i^{(m)} + ka_1^{(m)} \\ a_j^{(m+1)} &= a_1^{(m)} \\ a_h^{(m+1)} &= a_h^{(m)} \quad \text{for all } h \notin \{1, i, j\}. \end{aligned}$$

We claim that there exist $k = k_m$ and $j = j_m$ in $\mathbb{Z}$ for which

$$a_1^{(m+1)} \in \{1, \dots, a_1^{(m)} - 1\}. \quad (37)$$

To that end, first note that, for any $j \in \{2, \dots, n\}$, Euclid's division lemma provides an integer $k \in \mathbb{Z}$ such that $a_j^{(m)} - ka_1^{(m)} \in \{0, \dots, a_1^{(m)} - 1\}$. This means that we need to show that there is one $j \in \{2, \dots, n\}$ such that this number is different from 0. Assume, by contradiction, that this is not the case. Then, for any $j \in \{2, \dots, n\}$ there exists k_j such that $a_j^{(m)} = k_j a_1^{(m)}$. But then we can write

$$(a_1^{(m)}, \dots, a_n^{(m)}) = a_1^{(m)}(1, k_2, \dots, k_n),$$

which means that the determinant $\det B^{(m)} = \pm \det B$ is divisible by $a_1^{(m)}$. Since the matrices are in $\mathrm{GL}(n; \mathbb{Z})$ and $a_1^{(m)} > 0$ by hypothesis, we have $a_1^{(m)} = 1$, contradicting the induction hypothesis. We have proven the existence of k, j for which (37) holds. Take any $i \notin \{k, j\}$ and multiply on the left by $S_{1,i,j}^k$. This proves that the sequence $(a_1^{(m)})_{m \in \mathbb{N}}$ is positive and strictly decreasing, allowing us to conclude that there exists $m_* \in \mathbb{N}$ such that $a_1^{(m_*)} = 1$ and the algorithm terminates.

By multiplying $B^{(m_*)}$ on the left by $S_{1,2,n}^{a_n^{(m_*)}}$, we can bring its first column to the form $(1, a'_2, \dots, a'_{n-1}, 0)$ for some $a'_2, \dots, a'_{n-1} \in \mathbb{Z}$. Further left-multiplication by matrices of

the form $S_{1,2,j}^a$, for j from 3 to $n-1$ and appropriate integers a , brings it to the form $(1, a_2'', 0, \dots, 0)$.

Note that up until now, we have only multiplied B by elements in the subgroup $K \subset H$. In the case we have started with $B \in H$, then the matrix we end up with belongs to H , too. In that case, we find that the first column $(1, a_2'', 0, \dots, 0)$ actually satisfies $a_2'' = 0$, since we have $1 + a_2'' = 1$ by definition of H . This finishes step 1 for the case $B \in H$. In the case $B \in H'$, we find that $a_2'' = (n+1)k'$ for some $k' \in \mathbb{Z}$. Applying $T^{-k'}$ then proves the claim of step 1 for the case $B \in H'$.

Step 2. We consider the following subgroup:

$$G = \left\{ \begin{pmatrix} 1 & c \\ 0 & D \end{pmatrix} \middle| c = (c_2, \dots, c_n) \in \mathbb{Z}^{n-1}, D \in \mathrm{GL}(n-1; \mathbb{Z}) \right. \\ \left. \middle| \forall j \in \{2, \dots, n\}, c_j + \sum_i D_{ij} = 1 \right\} \subset H. \quad (38)$$

Let $G' \subset H'$ be the corresponding subgroup, where the latter condition on G is replaced by $c_i + \sum_j D_{ji} = 1 \pmod{n+1}$. In the first step, we have shown that, up to applying elements in K or K' , we can assume that $B \in G$ or $B \in G'$, respectively. Therefore, to prove the lemma, it suffices to show that any element in G is contained in H and that any element in G' is contained in H' . We start with the case of G . Note that $G \cong \mathrm{GL}(n-1; \mathbb{Z})$ via the natural isomorphism

$$\chi: \mathrm{GL}(n-1; \mathbb{Z}) \rightarrow G, \quad \chi(D) = \begin{pmatrix} 1 & c \\ 0 & D \end{pmatrix}, \quad c_i = 1 - \sum_j D_{ji}.$$

Now let $\chi(D_0)$ be the matrix obtained from $B \in H$ after the first step. Then, we know that D_0 can be written as a product $D_0 = F_1 \cdots F_N$ of permutations $\Pi_{i,j}$ and one elementary shear Σ — this is just a consequence of the ordinary Gauß algorithm proving that permutations and one shear generate all of $\mathrm{GL}(n-1; \mathbb{Z})$. Now note that both $\chi(\Pi_{i,j})$ and $\chi(\Sigma)$ are in K ; indeed, the image of a permutation in $\mathrm{GL}(n-1; \mathbb{Z})$ is a permutation in $\mathrm{GL}(n; \mathbb{Z})$, and the image of the shear Σ is, up to further coordinate permutations, equal to S_1 . Therefore, the decomposition

$$\chi(D_0) = \chi(F_1) \cdots \chi(F_N), \quad F_l \in \{\Pi_{1,2}, \dots, \Pi_{n-1,n}, \Sigma\}$$

proves that $\chi(D_0) \in K$.

The case of G' is similar, with one small difference: Since the equation $c_i + \sum_j D_{ji} = 1$ only holds modulo $n+1$ in this case, we cannot assume that the result B' of the first step lies in the image of the map χ . However, we can still decompose its lower right part $D_0 \in \mathrm{GL}(n-1; \mathbb{Z})$ as $D_0 = F_1 \cdots F_N$ as above. By a simple computation, we find that $\chi(F_N)^{-1} \cdots \chi(F_1)^{-1} \cdot B'$ is, in the notation of (38), of the form $D = I_{n-1}$ (the unit matrix in $\mathrm{GL}(n-1; \mathbb{Z})$) and $c_i \equiv 0 \pmod{n+1}$. This latter matrix can be written as a product of permutations and powers of T . $\square$

Proof of Theorem 5.7. We start by proving (1). Let $L_0 = T(1, \dots, 1) \subset \mathbb{C}^n$ be the monotone product torus. Since $\mathbb{C}^n$ is contractible, we have that $H_2(\mathbb{C}^n, L) \cong H_1(L)$. Let $e_1, \dots, e_n \in H_1(L_0)$ be the standard basis generated by elements of the form $e_i = [\{*\} \times \dots \times S^1 \times \dots \times \{*\}]$ on product tori. All coordinate permutations on the e_i are contained in $\widehat{\mathcal{M}}_{\text{Ham}}(L_0)$, since coordinate permutations on $\mathbb{C}^n$ can be realized by Hamiltonian diffeomorphisms and map L_0 to itself³.

Here we have $H_2(\mathbb{C}^n) = 0$, and thus Question 1.14 asks whether, in the basis $\{e_i\}$, the group of Lagrangian monodromies is equal to H as defined in (33). By Lemma 5.8, it is sufficient to show that S_1 is in $\widehat{\mathcal{M}}(L_0)$. To that end, let L_1 be the *fully twisted* Chekanov torus. In Chekanov's notation in [Che96], it corresponds to $L_1 = T_1^n(1) = \Theta_0^{n-1}(S^1(1))$, where Θ_0^{n-1} denotes the $(n-1)$ -fold twisting operation. However, we will follow the exposition in [Bre23, Section 5], where L_1 is constructed as the lift of a curve from a symplectic T^{n-1} -quotient. This yields the same torus L_1 as defined by Chekanov, but it allows us to understand the Lagrangian isotopy to L_0 . As in the proof of [Bre23, Proposition 5.5], we pick a basis $\xi, \zeta_2, \dots, \zeta_n \in H_1(L_1)$, where ξ is a lift of the class of the circle in the two-dimensional symplectic quotient and $\zeta_2, \dots, \zeta_n$ are in the orbit of the T^{n-1} -action by which we reduce. Perturbing the reduced space as in the proof of [Bre23, Proposition 5.5] yields a Lagrangian isotopy $\Pi = \{L_t\}_{t \in [0,1]}$ for which

$$\Phi^\Pi(e_n) = \xi, \quad \Phi^\Pi(e_i) = \zeta_i + \xi \quad \text{for } i \in \{1, \dots, n-1\}. \quad (39)$$

The distinguished class is $\{\xi\}$. Any monodromy matrix preserving the Maslov class (by monotonicity, the area class is then automatically preserved) and the distinguished class is in $\widehat{\mathcal{M}}_{\text{Ham}}(L_1)$ by [Che96, Theorem 4.5]. In particular, since

$$\mathbb{1}S_1 = \mathbb{1}, \quad \text{and} \quad S_1 e_n = e_n,$$

we find that $S_1 \in \Phi^\Pi \widehat{\mathcal{M}}_{\text{Ham}}(L_1) (\Phi^\Pi)^{-1} \subset \widehat{\mathcal{M}}(L_0)$, which, using Lemma 5.8, proves (1).

Let us now move to the case of $\mathbb{CP}^n$. Let $L = T_{\text{Cl}}$ be the monotone Clifford torus. It coincides with the image of a product torus $T(\frac{1}{n+1}, \dots, \frac{1}{n+1}) \subset \mathbb{C}^n$ under a symplectic ball embedding $\varphi: B^{2n}(1) \hookrightarrow \mathbb{CP}^n$. Through this embedding, we find a natural basis $\varphi_* e_i \in H_1(L)$. Capping these off by disks contained in the image φ (where these disks are homotopically unique), we obtain elements in $H_2(\mathbb{CP}^n, L)$. By abuse of notation, we will call these $e_1, \dots, e_n$, too. Together with the ambient spherical element $S = j([\mathbb{CP}^1]) \in H_2(\mathbb{CP}^n, L)$, which generates $J = \text{im}(j: H_2(\mathbb{CP}^n) \hookrightarrow H_2(\mathbb{CP}^n, L))$, we obtain a basis

$$S, e_1, \dots, e_n \in H_2(\mathbb{CP}^n, L). \quad (40)$$

This can be seen by extracting a short exact sequence from the long exact sequence for relative homology and using the above capping to obtain a splitting of the relative second

³Although we will not need this, $\widehat{\mathcal{M}}_{\text{Ham}}(L_0)$ is actually *equal* to the set of these permutations. This follows from [Che96, Theorem 4.5] and the fact that the e_i are *distinguished classes* in Chekanov's language.

homology. Question 1.14 asks whether the group of Lagrangian monodromies is given by

$$\widehat{G} = \{\phi \in \text{Aut } H_2(\mathbb{CP}^n, L) \mid \phi^* m_L = m_L, \phi(S) = S\}.$$

In the dual basis to (40), the Maslov class is $m_L = (2(n+1), 2, \dots, 2)$, and we find

$$\begin{aligned} \widehat{G} &= \left\{ \begin{pmatrix} 1 & a \\ 0 & B \end{pmatrix} \middle| (n+1)a = \mathbb{1} - \mathbb{1}B \right\} \\ &= \left\{ \begin{pmatrix} 1 & a \\ 0 & B \end{pmatrix} \middle| \forall j \in \{1, \dots, n\}, (n+1)a_j + \sum_i b_{ij} = 1 \right\}, \end{aligned}$$

where $a = (a_1 \dots a_n) \in \mathbb{Z}^n$ and $B = (b_{ij})_{i,j} \in \text{GL}(n; \mathbb{Z})$. Note that this implies that $B \in H'$. Hence, we obtain that $\widehat{G}$ is isomorphic to H' as defined in (34). As an isomorphism, we choose the obvious inclusion $H' \hookrightarrow \widehat{G}$. Now note that every Lagrangian loop in $\mathbb{C}^n$ with endpoints on L_0 can be included into $\mathbb{CP}^n$ by scaling appropriately and then composing with φ . Since L is Lagrangian isotopic to the image of L_0 under φ , we find that

$$\widehat{\mathcal{M}}(L_0 \subset \mathbb{C}^n) \hookrightarrow \widehat{\mathcal{M}}(L \subset \mathbb{CP}^n). \quad (41)$$

By our choice of basis, this corresponds to the natural inclusion $H \subset H' \hookrightarrow \widehat{G}$. In other words, we can generate all of $H \subset \widehat{G}$ just by Lagrangian isotopies in the image of the ball embedding. By Lemma 5.8, it is thus enough to show that $T \in \text{GL}(n+1; \mathbb{Z})$ can be realized by a Lagrangian loop based at L . To that end, recall that permutations of homogeneous coordinates of $\mathbb{CP}^n$ can be realized by Hamiltonian (and therefore also Lagrangian) isotopies. In our choice of basis, this yields all permutations of the classes

$$S - e_1 - \dots - e_n, e_1, \dots, e_n \in H_2(\mathbb{CP}^n, L).$$

The matrix of the elementary permutation exchanging the first two elements is given by

$$\begin{pmatrix} 1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & -1 & 0 & \dots & 1 \end{pmatrix} \in \widehat{G} \subset \text{GL}(n+1; \mathbb{Z})$$

Denote the lower-right $n \times n$ submatrix of this matrix by $Q \in \text{GL}(n; \mathbb{Z})$. We have $Q \in H'$ as expected. Now, note that $\mathbb{1}QT = \mathbb{1}$, i.e. QT is in the subgroup H . Since all elements of that subgroup are realizable as monodromies (even in $\mathbb{C}^n$), so is T , which concludes the proof. $\square$

5.3. Flux in $\mathbb{CP}^n$. Understanding the homological monodromy group is sometimes enough to understand $\Gamma^{\mathbb{Z}}(L)$ and the flux set $F(L)$. Indeed, by Proposition 3.3, we find that $\Gamma^{\mathbb{Z}}(T_{\text{Cl}}) \cong \mathcal{M}^{\mathbb{Z}}(T_{\text{Cl}})$ in $\mathbb{CP}^n$. As for the flux set, we find the following.

Corollary 5.9. *Let $L \subset \mathbb{CP}^n$ be the endpoint of a Lagrangian isotopy that starts at T_{Cl} and has flux $\eta \in H^1(T_{\text{Cl}}; \mathbb{R})$. In the basis induced by the natural basis on T_{Cl} , we have that*

$$F(L) = \{\eta(A - \text{id}) \mid A \in H'\}.$$

Proof. Let Π be a path with $\Pi_0 = T_{\text{Cl}}$ and $\text{Flux } \Pi = \eta$. Loops in $\pi_1(\mathcal{L}, L)$ are in bijection with loops in $\pi_1(\mathcal{L}, T_{\text{Cl}})$ by conjugating with Π . Write $\Lambda = \Pi^{-1} \circ \Lambda_0 \circ \Pi$ for $\Lambda_0 \in \pi_1(\mathcal{L}, T_{\text{Cl}})$. Recall that $\text{Flux } \Lambda_0 = 0$, since $T_{\text{Cl}} \subset \mathbb{CP}^n$ is monotone, see Corollary 2.6. By the change of base point formula in Proposition 2.4, we find that

$$\text{Flux}(\Lambda) = (\phi^\Pi)^{-*}(\mathcal{M}(\Lambda_0)^*\eta - \eta).$$

Since $\mathcal{M}(\Lambda_0) \in H'$ in the natural basis, by Theorem 5.7, the claim follows. $\square$

Let us now turn to the flux of Lagrangian *paths*. In other words, we now try to determine the *shape invariant*

$$\text{Sh}(T_{\text{Cl}}) = \{\text{Flux}\{\Pi\} \in H^1(T_{\text{Cl}}; \mathbb{R}) \mid \Pi_0 = T_{\text{Cl}} \subset \mathbb{CP}^n\}$$

for the Clifford torus in $\mathbb{CP}^n$.

Corollary 5.10. *Let $n \geq 3$. Identifying $H^1(T_{\text{Cl}}; \mathbb{R}) = \mathbb{R}^n$ via the standard basis, the set*

$$\mathbb{R}^n \setminus (\mathbb{R} \cdot \mathbb{Z}^n) = \mathbb{R}^n \setminus \{ry \in \mathbb{R}^n \mid r \in \mathbb{R}, y \in \mathbb{Z}^n\} \quad (42)$$

is contained in $\text{Sh}(T_{\text{Cl}})$.

Proof. We only consider Lagrangian isotopies of the form $\Pi = \Lambda * \Pi'$, where $\Lambda \in \pi_1(\mathcal{L}, T_{\text{Cl}})$ and Π' is an isotopy among toric fibres. Then we have

$$\text{Flux } \Pi = \mathcal{M}(\Lambda)^* \text{Flux } \Pi'.$$

Indeed, by Corollary 2.6, the flux of Λ vanishes. Note that

- (1) for any η in the moment polytope $\Delta_{\mathbb{CP}^n}$, there is an isotopy Π' with $\text{Flux } \Pi' = \eta$; see also Proposition 4.2;
- (2) for any $A \in H'$, there is $\Lambda \in \pi_1(\mathcal{L}, T_{\text{Cl}})$ with $\mathcal{M}^\mathbb{Z}(\Lambda) = A$ in the standard basis.

Therefore, we find that

$$\{\eta A \in \mathbb{R}^n \mid A \in H', \eta \in \Delta_{\mathbb{CP}^n}\} \subset \text{Sh}(T_{\text{Cl}}). \quad (43)$$

Let us prove that this set contains $\mathbb{R}^n \setminus (\mathbb{R} \cdot \mathbb{Z}^n)$. To that end, let us first point out that $x = (x_1, \dots, x_n)$ belongs to $\mathbb{R} \cdot \mathbb{Z}^n$ if and only if, for all pairs $i \neq j \in \{1, \dots, n\}$, there exist $k, l \in \mathbb{Z}$ such that $kx_i + lx_j = 0$. Let $\varepsilon > 0$. We show that we can map $x \in \mathbb{R}^n \setminus (\mathbb{R} \cdot \mathbb{Z}^n)$ into $B^n(\varepsilon) \subseteq \Delta_{\mathbb{CP}^n}$ using right-multiplication by elements in H' . By the above discussion, there exists a pair i, j for which $kx_i + lx_j \neq 0$ for all $k, l \in \mathbb{Z}$. Assume, without loss of generality, that $i = 1$ and $j = 2$. By a standard division algorithm argument, this implies that there is a sequence $(k_r, l_r) \in \mathbb{Z}^2$ such that $k_r x_1 + l_r x_2 \rightarrow 0$ for $r \rightarrow \infty$. By Lemma 5.8, we can map x to any element of the form

$$(x_1, x_2, \dots, x_{n-1}, \hat{x}_n = x_n + c(n+1)(kx_1 + lx_2))$$

for any $c, k, l \in \mathbb{Z}$. Note that here we use that $n \geq 3$. By the above, we can make the term $kx_1 + lx_2$ arbitrarily small, so that we can make $\widehat{x}_n$ arbitrarily small. Let us therefore assume that $|\widehat{x}_n| < \frac{\varepsilon}{\sqrt{n(n+1)}}$. Again by Lemma 5.8, we can map $(x_1, \dots, x_{n-1}, \widehat{x}_n)$ to a vector having components $\widehat{x}_i = x_i - s_i(n+1)\widehat{x}_n$ for $i \in \{1, \dots, n-1\}$. For a suitable choice of s_i , we therefore have $|\widehat{x}_i| < \frac{\varepsilon}{\sqrt{n}}$ for all i , and thus $\widehat{x} \in B^n(\varepsilon)$, proving the claim. $\square$

Remark 5.11. The set in (43) is strictly larger than the one in (42). In particular, the former contains $\Delta_{\mathbb{CP}^n}$. Trying to determine (43) runs into the question of trying to understand the orbits of H' , which is tricky for $n \geq 3$, as we will see below. $\bullet$

Let us now turn to obstructions. Based on [CM18], we find that some values in $H^1(T_{\text{Cl}}; \mathbb{R})$ cannot be realized as the flux of some Lagrangian isotopy. For simplicity, we again identify $H^1(T_{\text{Cl}}; \mathbb{R}) = \mathbb{R}^n$ via the standard basis of $H_1(T_{\text{Cl}})$. Under this identification, the Maslov class induces a class $m = 2\mathbb{1} = (2, \dots, 2) \in H^1(T_{\text{Cl}}; \mathbb{Z}) = \mathbb{Z}^n$, which must be preserved modulo $2(n+1)$ by Lagrangian isotopies. Let

$$\Lambda = \{\lambda \in H^1(T_{\text{Cl}}; \mathbb{Z}) \mid \lambda \cdot \xi \neq 0 \text{ for all } \xi \in H_1(T_{\text{Cl}}; \mathbb{Z}) \text{ with } \mathbb{1} \cdot \xi \equiv 1 \pmod{n+1}\}. \quad (44)$$

In other words, Λ is the set-theoretic complement in $H^1(T_{\text{Cl}}; \mathbb{Z})$ of the orthogonal complement of the affine lattice $\{\xi \in H_1(T_{\text{Cl}}) \mid \mathbb{1} \cdot \xi \equiv 1 \pmod{n+1}\}$.

Proposition 5.12. *The set*

$$B = \{\eta = r\lambda \in H^1(T_{\text{Cl}}; \mathbb{R}) \mid \lambda \in \Lambda, r \geq 1\},$$

where Λ is defined as in (44), is not in $\text{Sh}(T_{\text{Cl}})$. In particular, no $\eta = r\lambda$ where $r \geq 0$, and where $\lambda \equiv 1 \pmod{p}$ for some p dividing $n+1$ is in $\text{Sh}(T_{\text{Cl}})$.

Proof. Let Π be a Lagrangian isotopy in $\mathbb{CP}^n$ with starting point T_{Cl} and endpoint L . Suppose that Π has symplectic flux $\eta = r\lambda$ for $r \geq 1$ and $\lambda \in H^1(T_{\text{Cl}}; \mathbb{Z})$. We will show that λ must satisfy

$$\lambda \cdot \xi = 0 \text{ for some } \xi \in H_1(T_{\text{Cl}}; \mathbb{Z}) \text{ with } \mathbb{1} \cdot \xi \equiv 1 \pmod{n+1}; \quad (45)$$

c.f. (44). By [CM18, Theorem 1.1], the endpoint L of the Lagrangian isotopy bounds a topological disk whose class $\beta \in H_2(\mathbb{CP}^n, L)$

(1) has Maslov index $m(\beta) = 2$, and

(2) has area $\langle [\omega]_L, \beta \rangle \in (0, 1]$ (in our normalization of the symplectic form).

Let $\beta_0 \in H_2(\mathbb{CP}^n, T_{\text{Cl}})$ be the class such that (in the notation of §2.2) $\Phi^\Pi \beta_0 = \beta$, meaning that β_0 is pushed to β by the Lagrangian isotopy Π . Since Lagrangian isotopies preserve the

⁴Note that we have shown that all elements that are not in $\mathbb{R} \cdot H^1(T_{\text{Cl}}; \mathbb{Z}) \subset H^1(T_{\text{Cl}}; \mathbb{R})$ are actually contained in $\text{Sh}(T_{\text{Cl}})$. Therefore, supposing that $\eta = r\lambda$ with λ in the lattice $H^1(T_{\text{Cl}}; \mathbb{Z})$ is not that much of a restriction.

Maslov index, we have $m(\beta_0) = 2$. Let us now compute its area $\langle [\omega]_L, \beta \rangle$:

$$\begin{aligned} \langle [\omega]_L, \beta \rangle &= \langle [\omega]_{T_{\text{Cl}}}, \beta_0 \rangle + \text{Flux } \Pi \cdot \partial \beta_0 \\ &= 1 + \text{Flux } \Pi \cdot \partial \beta_0. \end{aligned}$$

In the first line, we have used the area formula (11). In the second line, we have used the fact that $[\omega]_{T_{\text{Cl}}} = \frac{m}{2}$, i.e. T_{Cl} is monotone. Now setting $\xi = \partial \beta_0$, we can prove (45). Indeed, since $m(\beta_0) = 2$, we find that $\mathbb{1} \cdot \xi = 1 \pmod{n+1}$. By point (2) and the area computation above, we find $\text{Flux } \Pi \cdot \xi \in (-1, 0]$. Since we have supposed that $\text{Flux } \Pi = r\lambda$ for an integral class $\lambda \in H^1(T_{\text{Cl}}; \mathbb{Z})$ and $r \geq 1$, we deduce that $\lambda \cdot \xi \in (-\frac{1}{r}, 0]$. Since $\lambda \cdot \xi \in \mathbb{Z}$, we deduce that it vanishes, as claimed in (45). This proves that B is in the complement of $\text{Sh}(T_{\text{Cl}})$.

For the second claim, let $\lambda \equiv 1 \pmod{p}$ for some p dividing $n+1$. We show that $\lambda \in \Lambda$ as defined in (44). By the first part of the proof, this implies the claim. Let $\xi \in H_1(T_{\text{Cl}})$ with $\mathbb{1} \cdot \xi \equiv 1 \pmod{n+1}$. Then we find

$$\lambda \cdot \xi \equiv \mathbb{1} \cdot \xi \equiv 1 \pmod{p},$$

and therefore $\lambda \cdot \xi \neq 0$. □

Example 5.13. Every vector $r\mathbb{1} \in H^1(T_{\text{Cl}}; \mathbb{Z})$ for $|r| \geq 1$ is contained in B . This follows directly from the fact that $\pm\mathbb{1}$ are in Λ , which can either be seen by hand or by using the second claim in Proposition 5.12. Actually, we can do slightly better in this particular case. Namely, we find that $r\mathbb{1} \in B$ for all $r \in \mathbb{R} \setminus (-1, \frac{1}{n})$, meaning that, additionally to the above, the segment $r \in [\frac{1}{n}, 1)$ is obstructed, too. Indeed, taking $\text{Flux } \Pi = r\mathbb{1}$ as in the proof of Proposition 5.12 (with $\lambda = \mathbb{1}$), we find that $\mathbb{1} \cdot \xi \in (-\frac{1}{r}, 0]$, meaning that $\mathbb{1} \cdot \xi \in \{-(n-1), \dots, -1, 0\}$, since it is an integer. Hence there is no such ξ with $\mathbb{1} \cdot \xi \equiv 1 \pmod{n+1}$. In particular, we deduce that the obvious Lagrangian isotopies which can be realized in the moment polytope $\Delta_{\mathbb{CP}^n}$ are sharp on the diagonal $\mathbb{R}\mathbb{1} = \{(r, \dots, r) \mid r \in \mathbb{R}\} \subset H^1(T_{\text{Cl}}; \mathbb{R})$:

$$\text{Sh}(T_{\text{Cl}}) \cap \mathbb{R}\mathbb{1} = \Delta_{\mathbb{CP}^n} \cap \mathbb{R}\mathbb{1}. \quad \bullet$$

We suspect that the second part of Proposition 5.12 captures all directions which can be obstructed by the above method, but were unable to prove it; see also Question 6.7.

5.4. Flux in $\mathbb{CP}^2$. As stated in Theorem 1.17, the flux set of any Lagrangian torus in $\mathbb{C}^2$ is discrete. We now show that it is very much not the case when we move to $\mathbb{CP}^2$.

Proposition 5.14. *The flux set of a Lagrangian torus L in $\mathbb{CP}^2$ is dense if and only if there is a Lagrangian isotopy Π from T_{Cl} to L such that $\text{Flux } \Pi = (x, y)$ satisfies $\frac{x}{y} \notin \mathbb{Q}$. Moreover, whenever $F(L)$ is not dense, it is discrete.*

Note that all L with $\frac{x}{y} \notin \mathbb{Q}$ are irrational, i.e. $\omega(\pi_2(\mathbb{CP}^2, L))$ is dense in $\mathbb{R}$. However, there are irrational tori with $\frac{x}{y} \in \mathbb{Q}$, e.g. rational tori in the open ball B^4 , seen as a subset of $\mathbb{CP}^2$.

Remark 5.15. By the classification of toric fibres in $\mathbb{CP}^2$ up to Hamiltonian isotopies [STV19, Proposition 7.1], $\mathcal{L}\text{Ham}(L)$ is C^1 -locally path connected for all Lagrangian fibres $L = T_{(x,y)} \subseteq \mathbb{CP}^2$. Indeed, for any such L , there is a C^1 -neighbourhood composed of Lagrangians L' that are Hamiltonian isotopic (in that neighbourhood) to a toric fibre. But, by the classification, we can ensure that no toric fibre in a (potentially smaller) neighbourhood is Hamiltonian isotopic to L , hence local path connectedness. Note that this holds whether $\frac{x}{y}$ is rational or not. Therefore, there are some toric fibres for which $F(L)$ is dense, but where $\mathcal{L}\text{Ham}(L)$ is locally path connected. This shows the stark contrast between isotopies with vanishing flux and Hamiltonian isotopies in the Lagrangian context. •

Proof of Proposition 5.14. We consider the group $\mathcal{M}(L)$ of Lagrangian monodromies of L and its intersection $H' = \langle R, M \rangle$ with $H = \text{SL}(2; \mathbb{Z})$. Note that when $\frac{x}{y} \in \mathbb{Q}$, then $F(L)$ is a subset of an affine lattice by the change-of-basepoint formula (9), so that it is discrete.

We thus suppose that $\frac{x}{y} \notin \mathbb{Q}$. Clearly, it suffices to prove the density of the subset $F_+(L)$ of fluxes associated with orientation-preserving Lagrangian loops. Recall that the flux of all loops based at T_{Cl} vanishes. Therefore, by the change-of-basepoint formula (9), this set takes on the form

$$F_+(L) = \{(x, y)(A - id) \mid A \in H'\}.$$

Moreover, this set is dense if and only if the H' -orbit of $v = (x, y)$ is dense in $\mathbb{R}^2 - \{0\}$.

But note that $H_{\mathbb{R}} = \text{SL}(2; \mathbb{R})$ acts on $\mathbb{R}^2 - \{0\}$ via right multiplication. Since every nonzero vector of $\mathbb{R}^2$ can be completed to a frame of oriented area 1, this action is transitive. Furthermore, the stabilizer at $(1, 0)$ is the subgroup T of lower unitriangular matrices. Therefore, we can identify $\mathbb{R}^2 - \{0\}$ with the quotient under left-multiplication $T \backslash H_{\mathbb{R}}$, and it is easy to see that this identification is H -equivariant on the right.

Note that the H' -orbit of $T \cdot g_v$ is dense if and only if for every g there is a sequence $(A_k, B_k) \in H' \times T$ such that $B_k g_v A_k \rightarrow g$, which is itself equivalent to the T -orbit of $g_v \cdot H'$ being dense. But by the proof of Proposition 5.2, we know that H' has finite index in H . In particular, H' defines a lattice in $H_{\mathbb{R}}$. By classical work of Hedlund [Hed36] on the horocyclic flow, the T -orbits in $H_{\mathbb{R}}/H'$ are thus either periodic or dense. In fact, almost all orbits are dense. We are thus left with characterizing what precisely these dense orbits are.

From the proof of Proposition 5.2, we know that $[H : H'] = 8$. Therefore, there is an 8-to-1 cover $H_{\mathbb{R}}/H' \rightarrow H_{\mathbb{R}}/H$ that is T -equivariant. But since covers are local homeomorphisms, this means that, if $g' \cdot H$ is in the closure of $T(g \cdot H)$, then there are some lifts $\tilde{g} \cdot H'$ and $\tilde{g}' \cdot H'$ of $g \cdot H$ and $g' \cdot H$, respectively, such that $\tilde{g}' \cdot H'$ is in the closure of $T(\tilde{g} \cdot H')$. Moreover, if the original point was not in the orbit, then neither is its lift. From the above discussion, this means that the H' -orbit of v is dense if and only if its H -orbit is.

By irrational approximation, for every $\epsilon > 0$, there are $a, b \in \mathbb{Z}$ such that $|ax + by| < \epsilon$. Without loss of generality, we may suppose that $\gcd(a, b) = 1$. By Bézout's identity, there exist $c_0, d_0 \in \mathbb{Z}$ such that $ad_0 - bc_0 = 1$, and every other solution takes on the form $c = c_0 + ka$

and $d = d_0 + kb$ for some $k \in \mathbb{Z}$. In particular, we have that

$$(x \ y) \begin{pmatrix} a & c \\ b & d \end{pmatrix} = (ax + by \quad cx + dy) = (ax + by \quad c_0x + d_0y + k(ax + by)).$$

On the one hand, the above matrix is in H for any k by construction. On the other hand, we can ensure that both coordinates on the ensuing vector are smaller than ϵ for an appropriate k , since $|ax + by| < \epsilon$. Therefore, 0 is not isolated in $H \cdot v$, and the orbit is dense. $\square$

Remark 5.16. The proof mostly relies on the fact that $\mathcal{M}(L)$ defines a sublattice in $\mathrm{SL}(2; \mathbb{Z})$. For example, if Question 6.4 has a positive answer in the monotone $S^2 \times S^2$, then the same result will apply to Lagrangian tori in that symplectic manifold. Note that we already know that $F(L)$ is dense for some product tori in nonmonotone $S^2 \times S^2$'s [BK25]. $\bullet$

Let us now move to determining the values flux can take for Lagrangian isotopies that do not necessarily close up to loops. This is a continuation of the discussion from §5.3 for general $\mathbb{CP}^n$ to the case $n = 2$. The difference is that, in that case, we can classify H' -orbits.

Lemma 5.17. *The action of H' on the set of primitive vectors $\mathcal{P}^2$ of $\mathbb{Z}^2$ has three orbits:*

$$\begin{aligned} \mathcal{O}_+ &:= \{\lambda \in \mathcal{P}^2 \mid \lambda \equiv \mathbb{1} \pmod{3}\} = \mathbb{1} \cdot H' \\ \mathcal{O}_- &:= \{\lambda \in \mathcal{P}^2 \mid \lambda \equiv -\mathbb{1} \pmod{3}\} = -\mathbb{1} \cdot H' \\ \mathcal{O}_R &:= \mathcal{P}^2 \setminus (\mathcal{O}_+ \cup \mathcal{O}_-) = (1, 0) \cdot H' \end{aligned}$$

Proof. Since the action of H' preserves the set $\mathcal{O}_+$, the orbit of $\mathbb{1}$ is included in it. Moreover, if $\lambda \in \mathcal{P}^2$, then there is $A \in \mathrm{GL}(2; \mathbb{Z})$ such that $\mathbb{1}A = \lambda$. But if $\lambda \equiv \mathbb{1} \pmod{3}$, i.e. if $\lambda \in \mathcal{O}_+$, then $A \in H'$, which proves equality between the sets. The proof that $\mathcal{O}_- = -\mathbb{1} \cdot H'$ is completely analogous.

To conclude that $\mathcal{O}_R$ consists of one orbit, note that if $v \in \mathcal{P}^2$ is integrally transverse to some $\lambda \in \mathcal{O}_\pm$, then there is some $A \in \mathrm{GL}(2; \mathbb{Z})$ such that $(1, 0)A = v$ and $\pm \mathbb{1}A = \lambda$. This is because $\mathrm{GL}(2; \mathbb{Z})$ acts transitively on integral bases. In particular, $A \in H'$, so that $v \in (1, 0) \cdot H'$. Therefore, we only need to show that every vector in $\mathcal{O}_R$ is integrally transverse to some vector in $\mathcal{O}_\pm$.

Let $v = (v_1, v_2) \in \mathcal{O}_R$. By Bézout's identity, every integrally transverse vector to v is of the form $(\mp q, \pm p) + kv$ for some $k \in \mathbb{Z}$, where p and q form a solution to $pv_1 + qv_2 = 1$. We show that there is a vector of this form in $\mathcal{O}_\pm$.

First of all, note that $v \notin \mathcal{O}_\pm$ implies that, up to permutations, either $v_2 \equiv 0 \pmod{3}$ or $v \equiv (1, -1) \pmod{3}$. In the first case, $pv_1 \equiv 1 \pmod{3}$, so that $p, v_1 \equiv \pm 1 \pmod{3}$. In the positive case, $(-q, p) + kv \equiv (-q + k, 1) \pmod{3}$, which is in $\mathcal{O}_+$ for $k = q + 1$. In the negative case, $(-q, p) + kv \equiv (-q - k, -1) \pmod{3}$, which is in $\mathcal{O}_-$ for the same k . In the second case, $p - q \equiv 1 \pmod{3}$, so that $(q, -p) + kv \equiv (q + k, -1 - q - k) \pmod{3}$, which is in $\mathcal{O}_+$ for $k = -q + 1$. $\square$

Remark 5.18. The situation in higher dimensions is more complex. In general, there are more orbits than the three appearing in the case $n = 2$. For example, in the case $n = 3$, the vectors $(1, 0, 0)$ and $(1, -1, 1)$ are neither in the orbits $\mathbb{1} \cdot H'$, $-\mathbb{1} \cdot H'$ nor are they in the same orbit. Indeed, since 2 divides $n + 1 = 4$, their mod-2 reductions, which are invariants of the H' -orbits, differ. •

We can now determine most of $\text{Sh}(T_{\text{Cl}})$ for the Clifford torus $T_{\text{Cl}} \subset \mathbb{CP}^2$.

Proof of Theorem 1.20. As we have seen in the proof of Proposition 5.14, the H' -orbit of $\eta \in H^1(T_{\text{Cl}}; \mathbb{R}) = \mathbb{R}^2$ is dense if and only if $\eta \notin \mathbb{R} \cdot \mathbb{Z}^2$. Therefore, there is $h \in H'$ such that $\eta h^{-1} = \eta_0 = (x_0, y_0) \in \Delta_{\mathbb{CP}^2}$ in that case. Since H' is the Lagrangian monodromy group of T_{Cl} , we can pick a Lagrangian loop Λ with $\mathcal{M}(\Lambda) = h$. Furthermore, let Π be a Lagrangian isotopy through toric fibres which starts at T_{Cl} and ends at $T_{(x_0, y_0)}$. Then, $\text{Flux } \Pi = \eta_0$, and we find that

$$\text{Flux}(\Lambda * \Pi) = \text{Flux } \Lambda + \mathcal{M}(\Lambda)^* \text{Flux } \Pi = 0 + \eta_0 h = \eta.$$

This proves the claim for any irrational η .

Now let $\eta \in \mathbb{R} \cdot \mathbb{Z}^2$, and decompose it as $\eta = r\lambda$ as above. As above, we can act by the Lagrangian monodromy group H' on λ to send it

- (1) to $(1, 1)$ in case $\lambda \in \mathcal{O}_+$;
- (2) to $(-1, -1)$ in case $\lambda \in \mathcal{O}_-$;
- (3) to $(0, 1)$ in case $\lambda \in \mathcal{O}_R$.

The obstructions in the cases $\lambda \in \mathcal{O}_\pm$ follow from Proposition 5.12 and Example 5.13. Note that $(1, 0)$, $(-1, -1) \in \partial\Delta_{\mathbb{CP}^2}$, whereas, in the $(1, 1)$ direction, the critical value is $(\frac{1}{2}, \frac{1}{2}) \in \partial\Delta_{\mathbb{CP}^2}$. □

6. QUESTIONS AND OUTLOOK

A given symplectic isotopy can be deformed, by a homotopy preserving its endpoints, to a Hamiltonian isotopy if and only if its *symplectic* flux vanishes. This is obviously not true in the Lagrangian case, as there is an isotopy from the Clifford to the Chekanov torus in $\mathbb{C}^2$ which has vanishing Lagrangian flux.

Question 6.1. *Under which conditions can a Lagrangian isotopy be deformed into a Hamiltonian one?*

Since we have mostly considered loops in this paper, it may be appropriate to ask which conditions on $\Lambda \in \pi_1(\mathcal{L})$ ensures that it is deformable into a loop in $\mathcal{L}\text{Ham}$. See Proposition 4.4 for a complete answer for tori in $\mathbb{C}^2$. A related question is:

Question 6.2. *Can every loop $\Lambda \in \pi_1(\mathcal{L})$ with $\text{FM}(\Lambda) = 0$ be deformed to a loop in $\mathcal{L}\text{Ham}$?*

Here are a few questions surrounding Lagrangian tori in $\mathbb{C}^2$. Theorem 1.17 shows that the flux-monodromy group $\Gamma(L)$ of any Lagrangian torus is generated by the images of the loops Ξ and Σ (after a change of basepoint) as defined in Section 4. Although we have studied π_1 of the space of immersed Lagrangian tori in §4.4, the embedded case is open.

Question 6.3. *Is $\pi_1(\mathcal{L}, T_{\text{Cl}})$, for the Clifford torus $T_{\text{Cl}} \subset \mathbb{C}^2$, generated by Ξ and Σ , as defined in Section 4?*

As we have seen, preserving the Maslov class is the only obstruction for a given element in $H^1(L; \mathbb{R}) \rtimes \text{MCG}(L)$ to be realized in $\Gamma(L)$, for $L \subset \mathbb{C}^2$ a Lagrangian torus. For general $L \subset X$, the Maslov class is an element $m_L \in H^2(X, L; \mathbb{Z})$.

Question 6.4. *Let $\Psi \in \text{Aut } H_2(X, L)$ preserve the Maslov class, $\Psi^* m_L = m_L$, and the subgroup of ambient classes $j(H_2(X))$. Does there exist $\Lambda \in \pi_1(\mathcal{L}, L)$ that has extended monodromy $\widehat{\mathcal{M}}(\Lambda) = \Psi$?*

We define the extended flux-monodromy group $\widehat{\Gamma}(L)$ of L as the image of the extended flux-monodromy morphism $\widehat{\text{FM}} : \pi_1(\mathcal{L}) \rightarrow H^1(L; \mathbb{R}) \rtimes \text{MCG}(X, L); \Lambda \mapsto (\text{Flux}(\Lambda), \widehat{\mathcal{M}}(\Lambda))$, where $\text{MCG}(X, L) := \pi_0(\text{Diff}(X, L))$ and $\text{Diff}(X, L) := \{\phi \in \text{Diff}(X) | \phi(L) = L\}$.

Question 6.5. *Is $\widehat{\Gamma}(L)$ discrete?*

If $\partial : H_2(X, L; \mathbb{R}) \rightarrow H_1(L; \mathbb{R})$ is surjective, then yes: Just like in Lemma 3.3, (11) implies that $\widehat{\Gamma}(L)$ is topologically isomorphic to $\widehat{\mathcal{M}}(L)$. At the other end of the spectrum, if $\partial = 0$, then every Lagrangian isotopy comes from a symplectic one (Proposition A.1). Therefore, the question becomes one of the relative topology of $\text{Symp}(X, L)$, the group of symplectomorphisms fixing L , in the full group $\text{Symp}(L)$, which goes beyond the scope of this paper.

In §5.3, we have discussed the set of elements in $H^1(T_{\text{Cl}}; \mathbb{R})$ that can be realized as the flux of a Lagrangian isotopy starting at the Clifford torus T_{Cl} . In the case $n = 2$, discussed in §5.4, we have found (Theorem 1.20) a complete answer up to the following values in $H^1(T_{\text{Cl}}; \mathbb{R})$.

Question 6.6. *Let $\eta \in H^1(T_{\text{Cl}}; \mathbb{R})$ be of the form*

$$\eta = r\lambda \text{ with } \lambda \in \mathcal{O}_R = (1, 0) \cdot H' \text{ and } r \geq 0,$$

where we have used the standard basis on $H_1(T_{\text{Cl}})$, and where H' is defined in (34). Is there a Lagrangian isotopy Π with $\Pi_0 = T_{\text{Cl}}$ and $\text{Flux } \Pi = \eta$?

If the answer is yes, then Π does not come from the method of construction we have used, i.e. it cannot be written as the concatenation of a loop $\Lambda \in \pi_1(\mathcal{L}, T_{\text{Cl}})$ and an isotopy through toric fibres. If the answer is no, this should come from additional information about the disk with boundary on T_{Cl} constructed in [CM18], e.g. a constraint to which class in $H_1(T_{\text{Cl}})$ its boundary can lie. Since all of the vectors $\lambda \in \mathcal{O}_R$ are in the same H' -orbit, it suffices to settle the question for one such λ .

In the case $n \geq 3$, we have a less complete picture: On the one hand, all irrational vectors $\mathbb{R}^n \setminus (\mathbb{R} \cdot \mathbb{Z}^n)$ can be realized as flux. On the other hand, values of the form $\eta = r\lambda$ for $r \geq 0$ and $\lambda \in \Lambda$ (where Λ is defined in (44)) are obstructed. As in the case $n = 2$, we do not know whether vectors of the form $\eta = r\lambda$ for $r \geq 1$ and $\lambda \in \mathbb{Z}^n \setminus \Lambda$ are obstructed, meaning that we can ask the same question as Question 6.6. However, in the case $n \geq 3$, we do even not know that the complement of Λ forms a single H' -orbit. We suspect this to be true; see Question 6.7. We cannot completely settle the question for $\eta = r\lambda$ with $0 < r < 1$, either. Of course, for small enough r , the vector η lies in the moment polytope $\Delta_{\mathbb{CP}^n}$ and can thus be realized as flux. This question depends on how close to the moment polytope $\Delta_{\mathbb{CP}^n}$ we can map λ by an element in H' . In other words, this again boils down to understanding the orbit structure of the action of H' on $\mathbb{Z}^n$.

Question 6.7. *Consider the orbits of the action of H' on primitive vectors in $\mathbb{Z}^n$, for $n \geq 3$. Does the direct generalization of Lemma 5.17 hold, i.e. is it true that*

- (1) Λ is given by all vectors $\lambda \in H_1(T_{\text{Cl}})$ that satisfy $\lambda \equiv \mathbb{1} \pmod{p}$ for some p dividing $n + 1$?
- (2) the complement $\mathcal{P} \setminus \Lambda$ of Λ among primitive vectors is given by the orbit $(1, 0, \dots, 0) \cdot H'$?

As we have seen in Corollary 2.6, $\text{Flux } \Lambda = 0$ for any loop $\Lambda \in \pi_1(\mathcal{L}, L)$ based at an H -monotone Lagrangian L . For product tori in $\mathbb{C}^n$ and the Clifford torus in $\mathbb{CP}^n$, we have seen that any element in $\text{Sh}(L)$ can be realized by the concatenation of a loop with a straight path.

Question 6.8. *Let $L \subset (X, \omega)$ be an H -monotone Lagrangian. Can any $\eta \in \text{Sh}(L)$ be written as $\text{Flux}(\Lambda * \Pi) = \text{Flux } \Pi \circ \mathcal{M}(\Lambda)$, where $\Lambda \in \pi_1(\mathcal{L}, L)$ and Π is a Lagrangian isotopy whose flux is contained in a line segment in $H^1(L, \mathbb{R})$ for all times?*

Recall that the set $\text{Sh}^*(L) \subset H^1(L; \mathbb{R})$ of all fluxes that can be reached by a straight line is called the *star shape* [STV19]. It was also studied in [EGM18]. This question asks whether every flux is a star flux up to base change by a Lagrangian monodromy element.

Finally, we have seen that $\mathcal{L}\text{Ham}(L)$ is C^1 -locally path connected when it is a Lagrangian torus in $\mathbb{C}^2$ (Theorem 1.17) or a toric fibre in $\mathbb{CP}^2$ (Remark 5.15) — despite the fact that the flux set is often dense in the latter case. Moreover, the examples constructed in [BK25] where $\mathcal{L}\text{Ham}(L)$ is not C^1 -locally path connected are in *non-monotone* $S^2 \times S^2$. This thus begs the question:

Question 6.9. *Let L be a Lagrangian torus in a 4-dimensional H -monotone symplectic manifold X . Must $\mathcal{L}\text{Ham}(L)$ be C^1 -locally path connected?*

Recall that we know this is false when $\dim X \geq 6$ [ACLS25].

APPENDIX A. PATHS OF HAMILTONIAN ISOTOPIC LAGRANGIANS

The first appendix concerns a result that is known to experts but hard to find in the literature.

Proposition A.1. *Let $\Pi : [0, 1] \rightarrow \mathcal{L}$ be a smooth path from L_0 to L_1 . The following statements are equivalent:*

- (a) $\text{Flux}(\{\Pi(\tau)\}_{0 \leq \tau \leq t}) = 0, \forall t;$
- (b) $\Pi(t) \in \mathcal{L}\text{Ham}(L), \forall t;$
- (c) *there is some $H : [0, 1] \times X \rightarrow \mathbb{R}$ such that $\Pi(t) = \phi_H^t(L_0), \forall t.$*

Moreover, if we only have that $\omega(H_2(X, \Pi(t)))$ is constant for all t , then we still have that $\Pi(t) = \psi_t(L_0)$ for some $\{\psi_t\} \in \text{Symp}_{c,0}(X)$.

By a *smooth path* of Lagrangians, we mean that there is a smooth map $F : [0, 1] \times L \rightarrow M$ such that $f_t = F(t, \cdot)$ parametrizes $\Pi(t)$. By the Whitney approximation theorem, every path in $\mathcal{L}$ is homotopic relative to its endpoints to such a path.

To prove the statement, we need the following technical lemma.

Lemma A.2. *Suppose that $F : [0, 1] \times L \rightarrow X$ has the property that, for every $t_0 \in [0, 1]$, there is some $\epsilon > 0$ and some open U of X containing $f_{t_0}(L)$ with $\omega|_U = d\lambda$ and such that $f_t^*\lambda = dh_t$ for all $t \in (t_0 - \epsilon, t_0 + \epsilon)$. Then, there exists a Hamiltonian H on X such that $f_t(L) = \phi_H^t(L)$ for all $t \in [0, 1]$.*

Proof. Consider the closed time-dependent 1-form $\alpha_t := f_t^*(\iota_{v_t}\omega)$ on L , where $v_t := \frac{df_t}{dt}$. By [AS01, Lemma 2.3], it suffices to prove that α_t is exact for all $t \in [0, 1]$. For some $t_0 \in [0, 1]$, take ϵ, U , and λ as above. Then, for every $t \in (t_0 - \epsilon, t_0 + \epsilon)$, a direct computation gives that

$$\alpha_t = d\left(\frac{d}{dt}h_t - \lambda(v_t)\right), \quad (46)$$

where $f_t^*\lambda = dh_t$ as above. This proves the statement. $\square$

Proof of Proposition A.1. It is trivial that (c) implies (a) and (b). Before proving the other implications, we make the following general observation. Let $F : [0, 1] \times L \rightarrow X$ parametrize the path Π . For $t_0 \in [0, 1]$, we fix a Weinstein neighbourhood U of $f_{t_0}(L)$. Then, by the smoothness of F , there is some $\epsilon > 0$ such that $f_t(L)$ can be seen as the graph of a closed 1-form σ_t of L through that Weinstein neighborhood for all $t \in (t_0 - \epsilon, t_0 + \epsilon)$.

Take a loop $\ell : S^1 \rightarrow L$. Note that $\sigma_t(\ell)$ is precisely the area of the cylinder $C_t = \cup_{s \in [t_0, t]} \sigma_s(\ell) = \cup_{s \in [t_0, t]} f_s(\ell)$. Furthermore, the area of C_t is continuous in t :

$$|\omega(C_t) - \omega(C_s)| = |\sigma_t(\ell) - \sigma_s(\ell)| \leq \|\sigma_t - \sigma_s\| \text{length}(\ell),$$

where the norm and the length are computed using the same Riemannian metric on L .

Suppose now that we are in (a). Then, for $t \geq t_0$, we have that

$$\sigma_t(\ell) = \omega(C_t) = \text{Flux}(\{\gamma(\tau)\}_{0 \leq \tau \leq t})(\ell) - \text{Flux}(\{\gamma(\tau)\}_{0 \leq \tau \leq t_0})(\ell) = 0.$$

Likewise, if $t \leq t_0$, we have that

$$\sigma_t(\ell) = \omega(C_t) = -(\text{Flux}(\{\gamma(\tau)\}_{0 \leq \tau \leq t})(\ell) - \text{Flux}(\{\gamma(\tau)\}_{0 \leq \tau \leq t_0})(\ell)) = 0.$$

Going through all possible loops ℓ , we conclude that σ_t is exact for all $t \in (t_0 - \epsilon, t_0 + \epsilon)$. We then get (c) from Lemma A.2.

Suppose that we are in (b), i.e. $f_t(L) \in \mathcal{L}\text{Ham}$. Then, there is some Hamiltonian isotopy $\{\phi_s^{t,t_0}\}_{s \in [0,1]}$ such that $\phi_1^{t,t_0}(f_t(L)) = f_{t_0}(L)$. Therefore, $C'_t := \cup_{s \in [0,1]} \phi_s^{t,t_0}(f_t(\ell))$ is a cylinder with vanishing area and one boundary agreeing with C_t , so that

$$\omega(C_t) = \omega(C_t \# C'_t) \in \omega(H_2(M, f_{t_0}(L))).$$

But $\omega(H_2(M, L))$ is totally disconnected since $H_2(X, f_{t_0}(L))$ is finitely generated. Therefore, $\omega(C_t)$ must be constant in t , so that $\sigma_t(\ell) = 0$. We again get (c) from Lemma A.2.

Finally, we deal with the case when we only know that $\omega(H_2(X, f_t(L)))$ is constant. Suppose there is some $A \in H_2(X, f_{t_0}(L))$ such that $\partial A = \ell$. Then, we again have that

$$\omega(C_t) = \omega(A \# C_t) - \omega(A) \in \omega(H_2(M, f_{t_0}(L))),$$

so that $\sigma_t(\ell) = 0$ as in the above paragraph. But (46) becomes $\alpha_t = \frac{d}{dt}\sigma_t - d(\lambda(v_t))$ in the non-exact case. Therefore, α_t must vanish on the image of the boundary map $H_2(X, f_t(L)) \rightarrow H_1(f_t(L)) = H_1(L)$. By the long exact sequence in cohomology (with compact coefficients), this is equivalent to the existence of a compactly supported time-dependent closed 1-form θ_t on X such that $[\alpha_t] = f_t^*[\theta_t]$. But then, there is a function $K : [0, 1] \times L \rightarrow \mathbb{R}$ such that $\alpha_t = f_t^*\theta_t + dk_t$. It then suffices to pick a compactly supported function $\tilde{K} : [0, 1] \times X \rightarrow \mathbb{R}$ such that $f_t^*\tilde{k}_t = k_t$ and to let $\{\psi_t\}$ be the symplectic isotopy generated by $\theta_t + d\tilde{k}_t$. $\square$

Remark A.3. The proof of Proposition A.1 implies that, whenever $H_1(L_0; \mathbb{R}) \rightarrow H_1(X; \mathbb{R})$ vanishes, any Lagrangian isotopy in $\mathcal{L}\text{Symp}$ is a Hamiltonian isotopy. This had already appeared under the additional hypothesis that $\omega(H_2(X, L))$ be discrete in [ACLS25].

At the other end of the spectrum, it also implies that every Lagrangian isotopy is generated by a symplectic one whenever $H_1(L_0; \mathbb{R}) \rightarrow H_1(X; \mathbb{R})$ is injective. $\bullet$

APPENDIX B. LAGRANGIAN HOMOTOPIES WITH VANISHING FLUX

A symplectic isotopy with vanishing *symplectic* flux is homotopic (relative to its endpoints) to a Hamiltonian isotopy. This classical fact is due to Banyaga [Ban97]; see for example [MS17, Theorem 10.2.5]. The corresponding statement is false for Lagrangian isotopies: there is a Lagrangian isotopy between the Clifford and the Chekanov torus in $\mathbb{C}^2$ which has vanishing flux. In this appendix, we discuss a Lagrangian analogue of this fact in the setting of Lagrangian *immersions* under a rationality assumption. More precisely, we want to prove the following.

Theorem B.1. *Let $\{f_t : Q \looparrowright X\}_{t \in [0,1]}$ be a Lagrangian homotopy such that f_0 is H -rational. The flux $\text{Flux}(\{f_t\})$ vanishes if and only if $\{f_t\}$ is Lagrangian homotopic with fixed endpoints to an exact homotopy.*

We recall that a *regular homotopy* is a homotopy through immersions. It is *Lagrangian* if these immersions are all Lagrangian. A Lagrangian homotopy $\{f_t\}$ is *exact* if the associated map $F : [0, 1] \times Q \rightarrow X; (t, x) \mapsto f_t(x)$ has the property that $F^*\omega = dt \wedge dh_t$ for some smooth family $\{h_t : Q \rightarrow \mathbb{R}\}$. When all f_t 's are embeddings, being exact is equivalent to being Hamiltonian; see Lemma A.2.

We will give the precise definition of H -rationality for immersions below, but we note for now that, when f_0 is an embedding, it is equivalent to $\omega(H_2(X, f_0(Q)))$ being discrete.

The proof of Theorem B.1 relies on an extension of Evans and Kędra's h -principle for monotone immersions in $\mathbb{C}^n$ [EK15, Theorem D].

B.1. Mapping cone of an immersion. We first recall the construction of a mapping cone in the homological setting. Let A be a PID. Given a map $f : Q \rightarrow X$, we consider $\text{Cone}_k(f; A)$ to be the product A -module $C_k(X; A) \oplus C_{k-1}(Q; A)$. We take the convention that $C_{-1}(Q; A) = 0$. This comes with a differential

$$\partial_k(s, r) := (f_\#(r) - \partial s, \partial r),$$

where ∂ is the usual differential of singular homology and $f_\# : C_\bullet(Q; A) \rightarrow C_\bullet(X; A)$ is the chain map induced by f . Essentially, cycles in this differential correspond to chains in X with boundary in Q , together with a lift of its boundary to Q via f . Boundaries are pairs of boundaries in X and Q , but where the boundary in X is twisted by the chain in Q .

There is also a dual construction with the cohomological cone $\text{Cone}^k(f; A) := C^k(X; A) \oplus C^{k-1}(Q; A)$ and the differential $d^k(\theta, \tau) = (d\theta, f^\#(\theta) - d\tau)$.

As we will be using the homology of these complexes throughout this appendix, we introduce a convenient name and notation.

Definition B.2. *The homology group of X relative to the map f in degree k is*

$$H_k(X, f; A) := H_k(\text{Cone}_\bullet(f; A)) = \frac{\ker \partial_k}{\text{im } \partial_{k+1}}.$$

Dually, the cohomology group of X relative to f in degree k is

$$H^k(X, f; A) := H^k(\text{Cone}^\bullet(f; A)) = \frac{\ker d^k}{\text{im } d^{k-1}}.$$

It follows from standard homological algebra that these (co)homology groups are homotopy invariants of f . Likewise, the homologies and cohomologies are related through the usual universal coefficient theorems, as they come from an actual topological space: the mapping cone C_f . Moreover, we have an exact sequence

$$\cdots \longrightarrow H_{k+1}(X, f) \xrightarrow{\partial} H_k(Q) \xrightarrow{f_*} H_k(X) \xrightarrow{j} H_k(X, f) \longrightarrow \cdots, \quad (47)$$

where ∂ and j are induced by the obvious chain maps $C_{k+1}(X) \times C_k(Q) \rightarrow C_k(Q)$ and $C_k(X) \rightarrow C_k(X) \times C_{k-1}(Q)$, respectively, and a dual one in cohomology.

It is a straightforward exercise to check the following.

Lemma B.3. *If f is an embedding, then the obvious chain map $C_k(X) \times C_{k-1}(Q) \rightarrow C_k(X)/C_k(f(Q))$ is actually a quasi-isomorphism. It induces a commutative diagram between the above exact sequence and the homology exact sequence of the pair $(X, f(Q))$.*

Note also that we have flexibility in the choice of model of (co)chain complex in the construction of the cone. In particular, if we instead take differential forms, we recover the Bott–Tu construction of the de Rham cohomology relative to a map [BT82, I.§6]. The usual arguments show that this is isomorphic to the singular cohomology with coefficients in $\mathbb{R}$. Note that, by the cohomology version of (47), this also means that any closed k -form θ such that $f^*\theta = 0$ defines a cohomology class $[\theta]_f$ in $H^k(X, f; \mathbb{R})$.

B.2. A h -principle for H -rational Lagrangian immersions. We first define the class of Lagrangian immersions that we will be interested in studying.

Definition B.4. *A Lagrangian immersion $f : Q \looparrowright X$ is called H -rational if there is some $\tau \geq 0$ such that $\omega(H_2(X, f; \mathbb{Z})) = \tau\mathbb{Z}$.*

This definition makes sense since $f^*\omega = 0$. Note that, up to some normalization of ω , this is equivalent to ω defining a class $[\omega]_f$ in $H^2(X, f; \mathbb{Z})$.

We now follow the logic of Evans and Kędra’s h -principle [EK15, Theorem D] to prove an h -principle for H -rational Lagrangian immersions. To state it, we note that the Lagrangian flux is well-defined for any path of maps $\{g_t\}$ as long as g_0 and g_1 are Lagrangian immersions.

Theorem B.5. *Let X be a symplectic manifold, and take $\tau \geq 0$. Suppose that we are given a smooth family of maps*

$$\begin{array}{ccc} TQ & \xrightarrow{G_t} & TX \\ \downarrow & & \downarrow \\ Q & \xrightarrow{g_t} & X \end{array} \tag{48}$$

such that

- (i) G_t is fibrewise isotropic;
- (ii) $g_i = f_i$ is a H -rational Lagrangian immersion with constant τ , $i = 0, 1$;
- (iii) $G_i = df_i$, $i = 0, 1$;
- (iv) $\text{Flux}(\{g_t\}) = 0$.

Then, the family can be deformed through fibrewise isotropic maps to a homotopy of H -rational Lagrangian immersions $\{f_t\}$ such that $\text{Flux}(\{f_t\}_{0 \leq t \leq T}) = 0$ for all $T \in [0, 1]$.

Proof. By the Gromov–Lees h -principle for Lagrangian immersions [Gro70, Lee76], we may suppose that g_t is a Lagrangian immersion and that $G_t = dg_t$ for all t . Up to rescaling, we may also suppose that τ is an integer.

Then, $f : Q \looparrowright X$ being an H -rational Lagrangian immersion of constant τ is equivalent to the existence of a Hermitian line bundle $L \rightarrow X$ with curvature form $-2\pi\omega$, whose pullback by f is flat and trivial. Indeed, it suffices to replace the exact sequence appearing in [Vér20, Lemma 6] by (47) to get this result for immersions. But it is well-known that the circle bundle $\pi : E \rightarrow X$ with curvature form $-2\pi\omega$ admits a contact form α with $d\alpha = \pi^*\omega$. Then, f^*E being flat and trivial is equivalent to f admitting a Legendrian lift.

Moreover, for any line bundle $L \rightarrow X$ with connection form $-2\pi\omega$, f_0^*L and f_1^*L have the same holonomy representation, so that one is trivial if and only if the other one is. Indeed, this is a direct application of the Gauss–Bonnet theorem (see [Vér20, Lemma 7]), applied to the cylinders swept by $\{g_t\}$, which have vanishing area by hypothesis. Therefore, f_0 and f_1 admit Legendrian lifts $\tilde{f}_0$ and $\tilde{f}_1$, respectively, in the same circle bundle.

In fact, the Gauss–Bonnet theorem ensures that the two lifts are homotopic. Indeed, it implies that the holonomy representation $\rho_t : H_1(Q; \mathbb{Z}) \rightarrow \mathrm{U}(1)$ of g_t^*E is given by

$$\rho_t = \rho_0 \cdot (\exp \circ \mathrm{Flux}(\{g_\tau\}_{0 \leq \tau \leq t})). \quad (49)$$

But since $t \mapsto \mathrm{Flux}(\{g_\tau\}_{0 \leq \tau \leq t})$ is a loop in $H^1(Q; \mathbb{R}) = \mathrm{Hom}(H_1(Q; \mathbb{Z}), \mathbb{R})$, it is contractible. Therefore, the family $\{\rho_t\}$ is homotopic to the constant one, i.e. the family $\{g_t^*E\}$ is isotopic through flat bundles to the constant one. Since the lift $\tilde{f}_i$ is given by a parallel section of f_i^*E , the deformation gives a family of sections in g_t^*E between them, i.e. lifts $\tilde{g}_t$ of g_t interpolating between $\tilde{f}_0$ and $\tilde{f}_1$. Moreover, since $E \rightarrow X$ induces a fibrewise isomorphism $\ker \alpha \rightarrow TX$, then $G_t = dg_t$ lifts to a fibrewise isotropic map $\tilde{G}_t$ over $\tilde{g}_t$.

By the Gromov–Duchamp h -principle for Legendrian immersions [Gro70, Duc84], the family of lifts $\{\tilde{G}_t\}$ is homotopic relative $\tilde{f}_0$ and $\tilde{f}_1$ through fibrewise isotropic maps to a family $\{\tilde{f}_t\}$ of Legendrian immersions. Projecting back onto X gives the required family of H -rational Lagrangian immersions. Note that, since the holonomy representation of f_t^*L is trivial for all t , any cylinder swept from f_0 to some f_T has vanishing area by (49), i.e. the flux vanishes. $\square$

Remark B.6. The same argument holds for higher parameter families. Therefore, each component of the space of maps as in (48) that contains an H -rational Lagrangian immersion of constant τ is weakly homotopy equivalent to the corresponding component of such immersions. In particular, this ensures that, if we start with a path of Lagrangian immersions, then the deformation to a path of H -rational ones can be made through Lagrangian immersions. $\bullet$

The h -principle for H -rational immersions, together with the above remark, gives directly Theorem B.1.

REFERENCES

- [ACLS25] Marcelo S. Atallah, Jean-Philippe Chassé, Rémi Leclercq, and Egor Shelukhin. Weinstein exactness of nearby lagrangians and related questions. *arXiv:2410.04158*, 2025.
- [ALLP24] Sílvia Anjos, Jun Li, Tian-Jun Li, and Martin Pinsonnault. Stability of the symplectomorphism groups of rational surfaces. *Mathematische Annalen*, 389(1):85–119, 2024.
- [ALP94] Michèle Audin, François Lalonde, and Leonid Polterovich. Symplectic rigidity: Lagrangian submanifolds. In Jacques Lafontaine and Michèle Audin, editors, *Holomorphic curves in symplectic geometry*, volume 117 of *Progr. Math.*, pages 271–322. Birkhäuser, Basel, 1994.
- [Arn86] Vladimir I. Arnol’d. First steps in symplectic topology. *Uspekhi Mat. Nauk*, 41(6(252)):3–18, 1986.
- [AS01] Meike Akveld and Dietmar Salamon. Loops of Lagrangian submanifolds and pseudoholomorphic discs. *Geom. Funct. Anal.*, 11(4):609–650, 2001.
- [ASW22] Marcin Augustynowicz, Jack Smith, and Jakub Wornbard. Homological Lagrangian monodromy for some monotone tori. *arXiv:2201.10507*, 2022.
- [Ban97] Augustin Banyaga. *The structure of classical diffeomorphism groups*, volume 400 of *Mathematics and Its Applications*. Springer, 1997.
- [BK25] Joé Brendel and Joontae Kim. Lagrangian split tori in $S^2 \times S^2$ and billiards. *Selecta Math. (N.S.)*, 31(4):Paper No. 68, 39, 2025.
- [Bre23] Joé Brendel. Real Lagrangian tori and versal deformations. *J. Symplectic Geom.*, 21(3):463–507, 2023.
- [Bre25] Joé Brendel. Hamiltonian classification of toric fibres and symmetric probes. *Algebr. Geom. Topol.*, 25(3):1839–1876, 2025.
- [BT82] Raoul Bott and Loring W. Tu. *Differential forms in algebraic topology*, volume 82 of *Graduate Texts in Mathematics*. Springer, 1982.
- [Che96] Yuri Chekanov. Lagrangian tori in a symplectic vector space and global symplectomorphisms. *Math. Z.*, 223(4):547–559, 1996.
- [Che98] Yuri Chekanov. Lagrangian intersections, symplectic energy, and areas of holomorphic curves. *Duke Math. J.*, 95(1):213–226, 1998.
- [CM18] Kai Cieliebak and Klaus Mohnke. Punctured holomorphic curves and Lagrangian embeddings. *Inventiones Mathematicae*, 212:213–295, 2018.
- [Cof05] Joseph Coffey. Symplectomorphism groups and isotropic skeletons. *Geometry & Topology*, 9(2):935–970, 2005.
- [CS16] Yuri Chekanov and Felix Schlenk. Lagrangian product tori in symplectic manifolds. *Comment. Math. Helv.*, 91(3):445–475, 2016.
- [DR19] Georgios Dimitroglou Rizell. The classification of Lagrangians nearby the Whitney immersion. *Geom. Topol.*, 23(7):3367–3458, 2019.
- [DRGI16] Georgios Dimitroglou Rizell, Elizabeth Goodman, and Alexander Ivrii. Lagrangian isotopy of tori in $S^2 \times S^2$ and $\mathbb{C}P^2$. *Geom. Funct. Anal.*, 26(5):1297–1358, 2016.
- [Duc84] Tom Duchamp. The classification of Legendre immersions. preprint, 1984.
- [Dui80] Johannes J. Duistermaat. On global action-angle coordinates. *Comm. Pure Appl. Math.*, 33(6):687–706, 1980.
- [EGM18] Michael Entov, Yaniv Ganor, and Cedric Membrez. Lagrangian isotopies and symplectic function theory. *Comment. Math. Helv.*, 93(4):829–882, 2018.
- [EK15] Jonathan D. Evans and Jarek Kędra. Remarks on monotone Lagrangians in $\mathbb{C}^n$. *Mathematical Research Letters*, 21(6):1241–1255, 2015.

- [Eli92] Yakov Eliashberg. Contact 3-manifolds twenty years since J. Martinet’s work. *Annales de l’institut Fourier*, 42(1-2):165–192, 1992.
- [EP97] Yakov Eliashberg and Leonid Polterovich. The problem of Lagrangian knots in four-manifolds. In *Geometric topology (Athens, GA, 1993)*, volume 2 of *AMS/IP Stud. Adv. Math.*, pages 313–327. Amer. Math. Soc., Providence, RI, 1997.
- [Eva23] Jonathan D. Evans. *Lectures on Lagrangian torus fibrations*, volume 105 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 2023.
- [Flo89] Andreas Floer. Cuplength estimates on Lagrangian intersections. *Communications on pure and applied mathematics*, 42(4):335–356, 1989.
- [FMAP20] Eduardo Fernández, Javier Martínez-Aguinaga, and Francisco Presas. The homotopy type of the contactomorphism groups of tight contact 3-manifolds, part I. arXiv:2012.14948, 2020.
- [FS06] Michael Farber and Dirk Schütz. Closed 1-forms with at most one zero. *Topology*, 45(3):465–473, 2006.
- [Gra73] André Gramain. Le type d’homotopie du groupe des difféomorphismes d’une surface compacte. *Annales scientifiques de l’École Normale Supérieure*, 6(1):53–66, 1973.
- [Gra89] Matthew A Grayson. Shortening embedded curves. *Annals of Mathematics*, 129(1):71–111, 1989.
- [Gro70] Misha Gromov. A topological technique for the construction of solutions of differential equations and inequalities. In *Actes du congrès international des mathématiciens (Nice 1970)*, volume 2, pages 221–225. Gauthier-Villars, 1970.
- [Gro85] Misha Gromov. Pseudo holomorphic curves in symplectic manifolds. *Invent. Math.*, 82(2):307–347, 1985.
- [Hat78] Allen Hatcher. Concordance spaces, higher simple-homotopy theory, and applications. In R. James Milgram, editor, *Algebraic and Geometric Topology*, volume XXXII of *Proceedings of Symposia in Pure Mathematics*, pages 3–21. American Mathematical Society, 1978.
- [Hed36] Gustav A. Hedlund. Fuchsian groups and transitive horocycles. *Duke Mathematical Journal*, 2(1):530–542, 1936.
- [Hin04] Richard Hind. Lagrangian spheres in $S^2 \times S^2$. *Geom. Funct. Anal.*, 14(2):303–318, 2004.
- [HLL11] Shengda Hu, François Lalonde, and Rémi Leclercq. Homological Lagrangian monodromy. *Geom. Topol.*, 15(3):1617–1650, 2011.
- [Hof88] Helmut Hofer. Lusternik–Schnirelman theory for Lagrangian intersections. *Annales de l’Institut Henri Poincaré C, Analyse non linéaire*, 5(5):465–499, 1988.
- [HPW16] Richard Hind, Martin Pinsonnault, and Weiwei Wu. Symplectomorphism groups of non-compact manifolds, orbifold balls, and a space of Lagrangians. *Journal of Symplectic Geometry*, 14(1):203–226, 2016.
- [Kul91] Ravi S. Kulkarni. An arithmetic-geometric method in the study of the subgroups of the modular group. *Am. J. Math.*, 113(6):1053–1133, 1991.
- [Lee76] J. Alexander Lees. On the classification of Lagrange immersions. *Duke Mathematical Journal*, 2(43):217–224, 1976.
- [LO95] Lê Hồng Vân and Kaoru Ono. Symplectic fixed points, Calabi invariant and Novikov homology. *Topology*, 34(1):155–176, 1995.
- [LP04] François Lalonde and Martin Pinsonnault. The topology of the space of symplectic balls in rational 4-manifolds. *Duke Mathematical Journal*, 122(2), 2004.
- [McD84] Dusa McDuff. Symplectic diffeomorphisms and the flux homomorphism. *Inventiones mathematicae*, 77(2):353–366, 1984.
- [MS17] Dusa McDuff and Dietmar Salamon. *Introduction to symplectic topology*. Oxford Graduate Texts in Mathematics. Oxford University Press, Oxford, third edition, 2017.

- [Ono06] Kaoru Ono. Floer-Novikov cohomology and the flux conjecture. *Geom. Funct. Anal.*, 16(5):981–1020, 2006.
- [Ono08] Kaoru Ono. A question analogous to the flux conjecture concerning Lagrangian submanifolds. In *Proceedings of Gökova Geometry-Topology Conference 2007*, pages 1–14. Gökova Geometry/-Topology Conference (GGT), Gökova, 2008.
- [Pin08] Martin Pinsonnault. Symplectomorphism groups and embeddings of balls into rational ruled 4-manifolds. *Compositio Mathematica*, 144(3):787–810, 2008.
- [Por24] Noah Porcelli. Families of relatively exact Lagrangians, free loop spaces and generalised homology. *Selecta Math. (N.S.)*, 30(2):Paper No. 21, 53, 2024.
- [Por25] Noah Porcelli. C^0 Lagrangian monodromy. *Bull. Lond. Math. Soc.*, 57(11):3441–3453, 2025.
- [Sik91] J.-C. Sikorav. Quelques propriétés des plongements Lagrangiens. In *Analyse globale et physique mathématique (Colloque à la mémoire d’Edmond Combet)*, number 46 in Mémoires de la Société Mathématique de France, pages 151–167. Société mathématique de France, 1991.
- [Sol13] Jake P. Solomon. The Calabi homomorphism, Lagrangian paths and special Lagrangians. *Math. Ann.*, 357(4):1389–1424, 2013.
- [STV19] Egor Shelukhin, Dmitry Tonkonog, and Renato Vianna. Geometry of symplectic flux and lagrangian torus fibrations, 2019.
- [Vér20] Alexandre Vérine. Bohr–Sommerfeld Lagrangian submanifolds as minima of convex functions. *Journal of Symplectic Geometry*, 18(1):333–353, 2020.
- [Vit90] Claude Viterbo. A new obstruction to embedding Lagrangian tori. *Inventiones mathematicae*, 100(1):301–320, 1990.
- [Wal68] Friedhelm Waldhausen. On irreducible 3-manifolds which are sufficiently large. *Annals of Mathematics*, 87(1):56–88, 1968.
- [Yau09] Mei-Lin Yau. Monodromy and isotopy of monotone Lagrangian tori. *Math. Res. Lett.*, 16(3):531–541, 2009.
- [Yau12] Mei-Lin Yau. Monodromy groups of Lagrangian tori in $\mathbb{R}^4$. *Michigan Mathematical Journal*, 61(2):431–446, 2012.

D-MATH, ETH ZÜRICH, RÄMISTRASSE 101, 8092 ZÜRICH, SWITZERLAND

Email address: `joe.brendel@math.ethz.ch`

CRM, UNIVERSITÉ DE MONTRÉAL, 2900 BOUL. ÉDOUARD-MONTPETIT, H3T 1J4, CANADA

Email address: `jean-philippe.chasse@umontreal.ca`

LABORATOIRE DE MATHÉMATIQUES JEAN LERAY, 2 RUE DE LA HOUSSINIÈRE, BP 92208, F-44322 NANTES CEDEX 3, FRANCE

Email address: `remi.leclercq@univ-nantes.fr`