Day 08 - Spin-c, parity, and Maslov gradings

Date: 2026-10-30. **Reading load:** approximately 4 hours 45 minutes. **Companion Obsidian note:** HRW_Day08_Obsidian.md.

Abstract The question The ungraded curve set $\widehat{HF}(M)$ determines the bordered type D structure, but the generators in $\widehat{CFD}(M)$ also carry relative Spin^c, $\mathbb{Z}/2$, and refined Maslov information. Which *extra geometry* remembers those gradings? Today the answer is: orientations, controlled lifts to a punctured cylinder, and (when there are multiple components in one Spin-c summand) **grading arrows**. The Maslov index is measured by puncture winding plus tangent rotation.

Warning Version and convention policy HRW23 means the 87-page, March 8, 2023 version (arXiv:1604.03466v3), with the theorem/page numbers cited below. **HRW22** means the 2022 companion paper; when referring to its numbered Definitions 28 or 35, we explicitly mean the arXiv:1810.10355v1 version, because pagination and some example numbering differ in the published article. Coefficients are $\mathbb{F} = \mathbb{F}_2$. These are **relative** gradings unless an absolute normalization is separately supplied. We will label independent algebraic deductions as such, not as HRW theorems.

Reading map

1. The two grading theorems
2. Parity and oriented curves
3. The half-lattice
4. The grading cover
5. Refined pairing
6. Phantom arrows
7. Maslov difference
8. The full grading group
9. Exercises
10. Solutions
11. Oral exam
12. Checklist

1. The two grading theorems

Let M be a compact oriented three-manifold with one torus boundary, let $z \in \partial M$ be marked, and let $T_M = \partial M \setminus \{z\}$. Choose an oriented boundary parametrization (α, β) as in HRW. The previous lectures established the ungraded dictionary

$$\widehat{CFD}(M, \alpha, \beta) \quad \longleftrightarrow \quad \widehat{HF}(M) = \{(\gamma_i, V_i, \Phi_i)\}.$$

Theorem HRW23, Theorem 1.7 (grading-enhanced invariant) For each $\mathfrak{s} \in \mathrm{Spin}^c(M)$, there is a collection of **oriented** immersed closed curves with local systems, denoted $\widehat{HF}(M, \mathfrak{s})$, in a specified covering space $\overline{T}_M \to T_M$. The collection is defined up to translation by the deck group, and its projection over all $\mathfrak{s}$ recovers the ungraded curve set. HRW23, Proposition 6.7 specifies the refined Spin^c and parity information in the gluing pairing.

Theorem HRW22, Theorem 5 (grading-enhanced pairing) The geometric pairing for $Y = M_0 \cup_h M_1$ agrees with $\widehat{HF}(Y)$ **as a relatively graded vector space**, including its Spin^c decomposition and relative Maslov grading within each summand. In the 2018 arXiv text, this is Theorem 5 and is developed in Section 2.

Logical distinction. HRW23, Section 6, builds orientations and a covering-space refinement, sufficient for its applications. HRW22, Section 2, explains the fuller Maslov grading and introduces extra **grading arrows** connecting otherwise disconnected curve components. These are not the crossover arrows used to encode local systems.

2. Parity from the torus algebra

2.1. Type D and type A have different vertex conventions

Let $N = \widehat{CFD}(M, \alpha, \beta; \mathfrak{s})$ be a reduced type D structure. Its relative mod-two grading g_D satisfies

$$g_D(\partial x) = g_D(x) + 1, \qquad g_D(\rho_I \otimes y) = g_{\mathcal{A}}(\rho_I) + g_D(y).$$

In HRW23, Section 6.1 (following Petkova), the algebra parity is

$$g_{\mathcal{A}}(\rho_1) = g_{\mathcal{A}}(\rho_3) = 0, \quad g_{\mathcal{A}}(\rho_2) = g_{\mathcal{A}}(\rho_{12}) = g_{\mathcal{A}}(\rho_{23}) = g_{\mathcal{A}}(\rho_{123}) = 1.$$

If the coefficient of $\rho_I \otimes y$ in $\delta^1(x)$ is nonzero, then

$$\boxed{g_D(y) - g_D(x) \equiv 1 + g_{\mathcal{A}}(\rho_I) \pmod 2.} \tag{2.1}$$

The associated type A generator parity is **shifted on one idempotent**. With HRW's convention, take

$$g_A(x) = g_D(x) + \epsilon(x), \qquad \epsilon(x) = \begin{cases} 1, & x \text{ has idempotent } \iota_0, \\ 0, & x \text{ has idempotent } \iota_1. \end{cases}$$

Consequently, an arrow $x \xrightarrow{I} y$ has $g_A(y) - g_A(x) = 1 + g_{\mathcal{A}}(\rho_I) + \epsilon(y) - \epsilon(x)$. Evaluating the six labels yields the following *derived check of HRW23, Section 6.1*:

Arrow label I	$g_D(y) - g_D(x)$	$g_A(y) - g_A(x)$
1	1	0
2	0	1
3	1	0
12	0	0
23	0	0
123	0	1

All entries are modulo two. This explains HRW's statement that **only labels** 2 **and** 123 **change the sign on an A-decorated graph**. Do not confuse this with the D-grading: labels 1 and 3 change g_D.

> **Proof Complete proof of the A-sign rule** For labels $1, 2, 3, 123$, the initial and terminal idempotents differ. Hence $\epsilon(y) - \epsilon(x) = 1$ modulo two, so the A-parity difference is $g_{\mathcal{A}}(\rho_I)$. For labels $12, 23$, the idempotents agree, so the A-parity difference is $1 + g_{\mathcal{A}}(\rho_I) = 0$. Substituting the six algebra parities proves the table.

2.2. From signs to orientations

HRW23, Figure 45 partitions the edges incident to each switch into types I and II. Proposition 6.1 asserts that for a **reduced extendable** type D structure each switch has at least one edge of each type. Table 1 then converts a choice of vertex signs into orientations of all underlying edges. Their Lemma 6.2 says that the orientations on the two ends of every edge agree.

> **Proof Proof check for HRW23, Lemma 6.2** At a vertex, Table 1 assigns an edge direction from the vertex sign and whether the incident edge is of type I or II. Edges labeled 2 or 123 have opposite I/II types at their ends, while the remaining labels have matching types (Figure 45). Exactly the first two labels reverse the A-sign, by the table above. The two changes compensate, so the local orientations agree at both ends of each edge. This establishes consistency. The classification of incident edges is a finite check against Figure 45, not a general theorem for arbitrary graphs.

For a valence-two loop, there is exactly one type-I and one type-II half-edge at every vertex. The orientation propagates continuously around the loop. Reversing the orientation reverses the mod-two grading of the entire component. If several curves encode a connected local system, their orientations are coherent. **One orientation choice per connected graded summand** is the geometric expression of the parity shift (HRW23, Remark 6.3).

> **Warning Absolute parity** Changing all signs within an unnormalized $\mathfrak{s}$-summand changes the relative **labeling** of parity classes but not their differences. The absolute $\mathbb{Z}/2$ grading on a closed three-manifold uses an additional normalization. In a glued pairing HRW fixes compatible conventions so that the sign of an oriented intersection records the parity of the corresponding tensor generator.

3. Refined Spin-c torsors and the half-lattice

Write

$$L = H_1(\partial M; \mathbb{Z}) \cong \mathbb{Z}\alpha \oplus \mathbb{Z}\beta, \qquad j_* : L \longrightarrow H_1(M; \mathbb{Z}).$$

Let ℓ generate $\ker j_*$ (the **homological longitude**; it may be a nonprimitive multiple of the rational longitude). Put

$$H_M = \operatorname{im} j_* \cong L/\langle \ell \rangle.$$

An ordinary Spin^c structure on M is an affine point of a torsor for $H^2(M;\mathbb{Z}) \cong H_1(M, \partial M; \mathbb{Z})$. A generator of a bordered module also remembers its idempotent and has a *refined* Spin-c label in a torsor modeled on H_M. These are different affine structures.

3.1. The half-integer lattice

Define the subgroup $2S \subset L$ by

$$2S = \{m\alpha + n\beta \in L : m + n \equiv 0 \pmod 2\}, \qquad S = \tfrac{1}{2}(2S) \subset L \otimes \mathbb{R}.$$

Equivalently,

$$\boxed{S = L \sqcup \left(L + \tfrac{\alpha+\beta}{2}\right).} \tag{3.1}$$

Proposition Derived lattice lemma (complete proof) The lattice S contains L with index two. Its two cosets are L and $L + (\alpha + \beta)/2$.

Proof. Every (m, n) with $m + n$ even has coordinates both even or both odd. Dividing by two gives either an integral vector or an integral vector plus $(1/2, 1/2)$. Conversely, both types double to integral vectors with even coordinate sum. The second coset is nonzero and has order two in S/L. □

Let S_M be its image modulo the homological longitude, so $H_M \subset S_M \subset \frac{1}{2}H_M$. The refined labels of the two idempotent summands occupy the two H_M-cosets. More precisely, HRW23, Lemma 6.4 identifies the union of the refined labels within a fixed $\mathfrak{s}$ with an S_M-torsor, making the ι_0 labels one coset and ι_1 labels the other.

3.2. The six edge increments

For an edge $x \xrightarrow{I} y$, the increment $s(y) - s(x)$ is given by **HRW23, Table 2 (p. 58)**:

I	Refined Spin-c increment Δ_I
1	$-(\alpha + \beta)/2$
2	$(\alpha - \beta)/2$
3	$(\alpha + \beta)/2$
12	$-\beta$
23	α
123	$(\alpha - \beta)/2$

Do not memorize this as six unrelated entries: four algebra multiplications explain the internal consistency.

Proposition Additivity check (complete calculation) Whenever $\rho_I \rho_J = \rho_K \neq 0$, we have $\Delta_I + \Delta_J = \Delta_K$.

Proof. The four nonzero products give

$$\begin{aligned} \Delta_1 + \Delta_2 &= -\beta = \Delta_{12}, \\ \Delta_2 + \Delta_3 &= \alpha = \Delta_{23}, \\ \Delta_1 + \Delta_{23} &= (\alpha - \beta)/2 = \Delta_{123}, \\ \Delta_{12} + \Delta_3 &= (\alpha - \beta)/2 = \Delta_{123}. \end{aligned}$$

These are all the nontrivial products of non-idempotent basis elements in the torus algebra. □

The increments for labels $1, 2, 3, 123$ cross between the two half-lattice cosets; those for $12, 23$ lie in H_M and preserve the idempotent coset. This gives a second verification of the grading geometry.

4. The grading cover and the lift theorem

4.1. The cover is not usually the universal cover

HRW23, Definition 1.6, defines $\overline{T}_M \to T_M$ as the cover corresponding to the kernel of

$$\pi_1(T_M) \longrightarrow H_1(\partial M) \xrightarrow{j_*} H_1(M).$$

With the standard identification $T_M = (\mathbb{R}^2 \setminus \mathbb{Z}^2)/\mathbb{Z}^2$ (up to placement of the marked lattice), one can write

$$\boxed{\overline{T}_M = (\mathbb{R}^2 \setminus \mathbb{Z}^2)/\langle \ell \rangle,} \tag{4.1}$$

where the translation by the integer vector ℓ acts freely. The deck group is

$$\boxed{\mathrm{Deck}(\overline{T}_M/T_M) = L/\langle \ell \rangle = H_M.}$$

The underlying surface of the cover is a punctured cylinder (with the appropriate periodic array of removed lattice points). It is generally **not** the plane punctured at all lattice points, which is a larger cover.

Warning Homological versus rational longitude If $\ell = d\lambda$ with $d > 1$ and λ primitive, quotienting the plane by λ instead of by $d\lambda$ loses the $\mathbb{Z}/d$ deck-group information. This changes the refined torsor, even though the two vectors have the same slope.

4.2. A torsion example

Suppose $L = \mathbb{Z}\mu \oplus \mathbb{Z}\lambda$ and $\ker j_* = \langle 2\lambda \rangle$. Then

$$H_M \cong \mathbb{Z}\langle \mu \rangle \oplus (\mathbb{Z}/2)\langle \lambda \rangle.$$

The cover $\overline{T}_M$ has period 2λ, not λ. The residual translation by λ is a nontrivial deck transformation of order two, while translation by μ has infinite order. This is an algebraic model, not an assertion about a particular named manifold.

4.3. Lifting a decorated graph

Theorem HRW23, Proposition 6.5 Fix a summand $\widehat{CFD}(M, \alpha, \beta; \mathfrak{s})$ and the choice identifying its refined Spin-c labels with lifts of the parametrizing arcs in $\overline{T}_M$. Its train track admits a unique lift for which each switch representing x lies on the lift of α or β prescribed by $s(x)$.

Proof Proof, modulo HRW's local edge dictionary First choose the lift of every switch using its refined label. For an arrow $x \xrightarrow{I} y$, the standard train-track segment for label I has displacement Δ_I between its endpoint lifts; compare HRW23, Figure 7 and Table 2. Hence the chosen lifted endpoints are connected by a lift of that segment. This produces a lift of every edge and makes all vertices consistent. Any other lift with those prescribed switch positions agrees on every edge by uniqueness of path lifting. Finally, changing the initial torsor identification translates **all** lifted switches by one element of H_M. This is precisely the deck-translation ambiguity. □

There is no canonical origin for a Spin-c torsor. Therefore a cover diagram is canonically defined only up to a **simultaneous** deck translation in the given $\mathfrak{s}$ summand. Translating separate components independently would erase relative grading information.

4.4. Why the reduction can be done upstairs

HRW23, proof of Theorem 1.7 (pp. 59-60), observes that train-track simplifications from Section 3 are supported in local handle regions. The moves lift to the appropriate regions in $\overline{T}_M$, and the relevant change-of-basis formulas preserve the relative parity. The resulting oriented decorated curves remain well defined modulo deck translation and boundary-parametrization changes already addressed in Section 5.

Theorem HRW23, Corollary 6.6 (homological longitude constraint) After filling in the puncture, each component of $\widehat{HF}(M)$ represents a multiple of the homological longitude ℓ. Moreover, if $b_1(M) = 1$ the oriented sum of curves in each $\mathfrak{s}$ represents ℓ; when $b_1(M) > 1$ it is null homologous. The last two statements concern the **total oriented cycle for a Spin-c summand**, not necessarily a claim that each individual component is primitive.

Proof Proof architecture of Corollary 6.6 The lift lives in the quotient by $\langle \ell \rangle$, so its projected closed-curve homology is in the cyclic subgroup generated by ℓ. To determine the signed total coefficient, fill along a transverse slope and compare algebraic intersection with the Euler characteristic of $\widehat{HF}$ of that filling in the given Spin-c structure. That Euler characteristic equals one for rational-homology-sphere fillings and zero when the filling has positive first Betti number. This last step invokes the standard Heegaard Floer Euler-characteristic property; it is not a proof from the covering space alone.

5. Refined pairing and Dehn-filling calculation

Let $Y = M_0 \cup_h M_1$ with orientation-reversing $h : \partial M_1 \to \partial M_0$, respecting the marked points. Use h to view the homological longitudes ℓ_0, ℓ_1 in the same boundary lattice $L = H_1(\partial M_0; \mathbb{Z})$, and set

$$\boxed{H_Y = L/\langle \ell_0, h_*(\ell_1)\rangle.} \tag{5.1}$$

For fixed restrictions $(\mathfrak{s}_0, \mathfrak{s}_1)$, the set of extensions to Y is an affine torsor for the appropriate H_Y in the HRW pairing framework. The two refined curves are pushed to a common cover T_Y with deck group H_Y; one factor is reflected according to the type A/type D convention. Translating a lift by an element $a \in H_Y$ selects a relative glued Spin-c class.

> **Theorem HRW23, Proposition 6.7** For $\mathfrak{t} \in \mathrm{Spin}^c(Y)$ restricting to $\mathfrak{s}_i$ on M_i, there is a deck translate $a_\mathfrak{t} \in H_Y$ for which the Floer pairing of the corresponding lifted curves computes $\widehat{HF}(Y, \mathfrak{t})$. The relative mod-two grading is the sign of the oriented intersection.

Proof reconstruction. A generator $x_0 \otimes x_1$ of $\widehat{CFA}(M_0) \boxtimes \widehat{CFD}(M_1)$ requires matching idempotents. The difference between two tensor generators is the combined refined Spin-c difference of their factors, with the second factor reflected. By the Mayer-Vietoris identification and the quotient (5.1), this difference vanishes exactly when both geometric intersections occur between the same pair of chosen lifted curve sets. The parity table of Section 2 converts the tensor-product parity $g_A(x_0) + g_D(x_1)$ into the intersection sign. The proof of HRW23, Proposition 6.7, carries out this last sign check case by case using the four possible idempotent/sign combinations. **This is a proof outline** because we do not reproduce the complete type A/type D local-segment table.

5.1. Worked Dehn-filling model: lens spaces

Glue two solid tori along a torus so their meridian slopes, expressed in a fixed oriented homology basis, are

$$\ell_0 = (1, 0), \qquad h_*(\ell_1) = (q, p), \qquad p > 0, \quad \gcd(p, q) = 1.$$

Then

$$H_Y \cong \mathbb{Z}^2/\langle (1, 0), (q, p)\rangle \cong \mathbb{Z}/p.$$

(The resulting lens-space naming convention $L(p, q)$ may vary with orientation.) The two primitive straight curves intersect transversely in p points, since

$$\left|\det \begin{pmatrix} 1 & q \\ 0 & p \end{pmatrix}\right| = p.$$

Their lifts give **one geometric intersection per class in** H_Y. The Floer differential vanishes for these straight representatives: a lifted bigon would require two affine lines in the plane to meet twice. Thus

$$\boxed{\dim_{\mathbb{F}} \widehat{HF}(Y, \mathfrak{t}) = 1 \quad \text{for every } \mathfrak{t} \in \mathrm{Spin}^c(Y).}$$

For instance, $(q, p) = (2, 5)$ gives five generators and five Spin-c summands of dimension one. This is the expected L-space behavior of a lens space. Their relative Maslov differences **between distinct Spin-c summands** are not obtained merely by putting these five intersection points in order; that would require further absolute grading normalization.

5.2. Companion-paper example: trefoil +3 surgery

HRW22, Section 2.1, Example 30 in arXiv:1810.10355v1 (Figure 18), pairs the right-handed trefoil exterior curve with a slope-three line. The picture has three generators and requires three different lifts of the surgery line to realize all of them. Hence the three generators lie in **different Spin-c structures**. This example validates the covering-space interpretation; it should not be confused with the toy calculation above, which uses two solid tori instead of a trefoil exterior.

6. The companion paper's grading arrows

Suppose $\widehat{HF}(M, \mathfrak{s})$ has r curve components. Even if every component is oriented and individually lifted, we must know their **relative** positions in the grading cover to compare generators associated with different components.

> **Definition HRW22, Section 2.1, Definition 28 (arXiv v1)** A set of grading arrows on r curve components consists of $r - 1$ auxiliary arrows joining them so that the union is connected. HRW calls them **phantom arrows**: they store relative grading information but are ignored when counting the immersed bigons defining the ordinary Floer differential.

Elementary graph lemma (derived). Any connected graph on r vertices has a spanning tree with $r-1$ edges. Choosing its edges as phantom arrows connects the components without adding redundant relative-position loops. After fixing a lift of one component, the arrow displacements determine a lift of every other component by propagation along the tree. Uniqueness follows because a tree has a unique path between vertices.

This is the algebraic reason that $r-1$ arrows suffice. If $r=1$, none is required. For $r=2$, one arrow records the relative offset of the two components. This point matters for HRW22's figure-eight $+1$ surgery calculation (Figure 19): the two underlying curves belong to the same Spin-c summand, and the phantom arrow is required to compare the Maslov grading of their intersection generators.

> **Warning Two distinct kinds of crossover arrow Crossover arrows** from HRW23, Section 3, are part of a train-track presentation and can encode actual local-system transport. **Phantom grading arrows** from HRW22, Definition 28, encode relative grading offsets; they contribute **nothing** to the ungraded Floer differential. Erasing a phantom arrow loses grading data; treating it as a genuine differential arrow changes the ungraded complex. Neither mistake is harmless.

6.1. Relative Spin-c grading of two intersections

Let Γ_0, Γ_1 be connected as graded curve sets by phantom arrows. For intersection generators x, y, choose a path from x to y on Γ_0 (possibly using phantom arrows) and a path from y back to x on Γ_1. Their concatenation determines a class in

$$\boxed{Q = H_1(T^2;\mathbb{Z})/\langle [\ell_0],[\ell_1]\rangle.} \tag{6.1}$$

This is the relative Spin-c difference of x and y (HRW22, Section 2.1). Why is it independent of the chosen paths? Changing a path by a full circuit of a component adds a multiple of its longitude class, which becomes zero in Q. A grading-arrow system additionally fixes the relative choices among components.

An equivalent test is: x, y have the same relative Spin-c grading precisely when one can choose lifts of both curve sets to the punctured plane so that both x and y lift to intersections of the **same pair of lifted components**.

7. Maslov difference from a lifted closed path

7.1. The geometric formula

Assume x, y belong to the **same Spin-c summand** of the intersection Floer complex. Choose paths p_0, p_1 on the two graded curve sets from x to y whose difference lifts to a closed piecewise smooth path

$$C = \widetilde{p}_0 - \widetilde{p}_1 \subset \mathbb{R}^2 \setminus \mathbb{Z}^2.$$

Assume the curves meet orthogonally at the corners and use HRW's rightward (clockwise) turning convention. Define

$$N(C) = \sum_{z\in\mathbb{Z}^2} \mathrm{wind}_C(z), \qquad R(C) = \text{total clockwise tangent rotation along the smooth portions.}$$

The sum has finite support for a compact lifted loop. Here $R(C)$ **excludes** the corner jumps at x and y.

> **Definition HRW22, Section 2.1 / Definition 35 in arXiv:1810.10355v1** The relative Maslov difference is
>
> $$\boxed{m(y) - m(x) = 2N(C) + \frac{R(C)}{\pi}.} \tag{7.1}$$
>
> This is the intrinsic winding-and-turning formula for oriented curves with grading arrows; the relevant periodic-domain ambiguity is that of the relative Floer grading.

The **lattice points** in (7.1) represent punctures. They are counted with winding number, so negative orientations contribute negative multiplicities. It is not Euclidean area in square units, even though the formula arises by simplifying an area calculation in the torus-algebra grading group.

7.2. Why a convex bigon gives $-1+2N$

> **Proposition Derived local calculation (complete)** Suppose C bounds an oriented embedded bigon with two orthogonal convex corners and with positive (counterclockwise) boundary orientation. Then the rightward rotation along its smooth sides is $R(C) = -\pi$. Therefore

$$\boxed{m(y) - m(x) = 2N(C) - 1.} \tag{7.2}$$

Proof. The total signed counterclockwise tangent turning of the positively oriented boundary is 2π. Each convex right-angle corner contributes an exterior turning angle $\pi/2$. The smooth portions therefore contribute $2\pi - 2(\pi/2) = \pi$ of counterclockwise turning. Rightward turning is its negative, $-\pi$. Substitute into (7.1). □

For a bigon that avoids punctures, $N(C) = 0$ and the difference is -1, exactly the grading behavior of a Floer differential. If one lattice point lies inside with positive winding, the relative difference becomes $+1$. Such a puncture-containing bigon **does not contribute** to the hat differential; it is nevertheless an admissible geometric region for comparing grades. This is the key to HRW22's examples.

Warning Formula and homotopy choices Formula (7.1) is evaluated on a path whose chosen lifts close. It is not an arbitrary-loop formula for pairs of generators in different Spin-c structures. In the presence of non-torsion Spin-c structures, relative Maslov data can be periodic modulo the divisibility of c_1; do not claim a canonical absolute integer grading from (7.1) alone.

7.3. A tiny worked calculation

Suppose a lifted convex bigon from x to y encloses no puncture and has the positive orientation above. Then $N = 0$, $R = -\pi$, so $m(y) - m(x) = -1$. If the same kind of bigon encloses one lattice point, $N = 1$ and the difference is 1. In both cases the difference is odd, agreeing with the fact that the two corner signs are opposite. Only the puncture-free case enters the Floer differential.

7.4. Grading-enhanced pairing: what is proved analytically?

In HRW22, Section 2.3 (arXiv v1, pp. 29-31), the paired bordered grading decomposes into a Spin-c component (coming from displacement in the cover) and a Maslov component. After reflecting the type D train track, the two Spin-c displacements cancel for generators in one closed-manifold Spin-c class. The Maslov components add. HRW deforms the train tracks to paths near the central horizontal and vertical lines; their algebraic area terms become $2N(C)$, and the corner terms become $R(C)/\pi$. This proves the grading comparison in Theorem 5.

Proof status: The local turning calculation above is complete, but the equality with LOT's full refined G-grading relies on the detailed graded-box-tensor construction in HRW22, Sections 2.2-2.3 and LOT, Chapter 11. That comparison is not reproduced from first principles here.

8. The full noncommutative grading group

HRW22, Section 2.2, recalls the torus-algebra grading group G. It is represented by triples $(m; i, j)$, with m, i, j half-integers subject to an integrality/parity constraint. The multiplication has the form

$$\boxed{(m; v)(m'; v') = (m + m' + \det(v, v');\ v + v'), \quad v = (i, j).} \tag{8.1}$$

The element $\lambda = (1; 0, 0)$ is central and records one Maslov step. The group is a central extension of the boundary homology grading by a Maslov coordinate.

Proposition Associativity and the commutator (complete algebraic derivation) The law (8.1) is associative. If $a = (m; v)$ and $b = (m'; w)$, then

$$\boxed{aba^{-1}b^{-1} = (2\det(v, w); 0).}$$

Proof. The central coordinate of $(ab)c$ is $m + m' + m'' + \det(v, w) + \det(v + w, u)$, while that of $a(bc)$ is $m + m' + m'' + \det(w, u) + \det(v, w + u)$. Bilinearity of the determinant makes these equal; the vector coordinates coincide. The inverse is $(-m; -v)$ because $\det(v, -v) = 0$. Substituting in the commutator gives $\det(v, w) - \det(w, v) = 2\det(v, w)$. □

Concrete noncommutativity. The allowed half-integer Maslov coordinates accommodate elements $a = (1/2; 1, 0)$ and $b = (1/2; 0, 1)$. Then $ab = (2; 1, 1)$ while $ba = (0; 1, 1)$. Their difference is the central element λ^2. This noncommutativity is why one cannot add the two bordered grading triples as if they lived in an ordinary abelian group.

A type D grading is a right G-set (up to its periodic-domain subgroup), while a type A grading uses the opposite-sided convention; the box product combines the two into a double-coset grading. HRW22, Section 2.2, Tables 2-3, records the exact coefficient-map shifts. The geometric formulation via lifted paths packages the same information in a much more readable way.

9. Exercises with hints

Exercise 1. The six parity shifts Starting from $g_{\mathcal{A}}(\rho_I)$ and $g_A = g_D + \epsilon$, derive **all six rows** of the parity table in Section 2. Why do only 2 and 123 change the signs in HRW's A-decorated graph? **Hint:** Compare the idempotents on both ends.

Exercise 2. The half-lattice Prove (3.1), then compute S/L and the two cosets for $L = \mathbb{Z}^2$. Explain why a ρ_1 arrow must cross between these cosets while a ρ_{12} arrow cannot. **Hint:** Use the displacement table.

Exercise 3. Four multiplication checks Verify $\Delta_1 + \Delta_2 = \Delta_{12}$, $\Delta_2 + \Delta_3 = \Delta_{23}$, $\Delta_1 + \Delta_{23} = \Delta_{123}$ and $\Delta_{12} + \Delta_3 = \Delta_{123}$. Explain how this relates to the existence of the lift in Proposition 6.5. **Hint:** The two sides of a decomposed train-track path must have the same total displacement.

Exercise 4. A cover with torsion Take $\ker j_* = \langle 3\lambda \rangle$ in $L = \mathbb{Z}\mu \oplus \mathbb{Z}\lambda$. Compute H_M, describe the deck action on $\overline{T}_M$, and explain why quotienting instead by λ is incorrect. **Hint:** Smith normal form of a one-row relation.

Exercise 5. A lens-space gluing The meridian slopes of two solid tori are $(1, 0)$ and $(3, 7)$. Compute the geometric intersections, H_Y, and the dimension of each $\widehat{HF}(Y, \mathfrak{t})$. Explain why the seven intersection points have different Spin-c grades. **Hint:** Use the determinant and the quotient lattice.

Exercise 6. Phantom arrows versus local systems There are three immersed-curve components in one Spin-c summand. Give a minimal system of grading arrows. Prove that their placements determine relative lift positions up to one global deck translation. Explain why these arrows are omitted from the ungraded Floer differential. **Hint:** Use a spanning tree.

Exercise 7. Maslov winding Compute (7.1) for three closed paths with $(N, R/\pi) = (0, -1), (1, -1), (-1, 1)$. Which case can be the boundary of a positively oriented, puncture-free differential bigon? Why does being a valid grading comparison not imply contributing to the differential? **Hint:** Separate puncture counting from bigon admissibility.

Exercise 8. Noncommutative grading group Prove associativity of (8.1) and compute the commutator. Using $a = (1/2; 1, 0)$ and $b = (1/2; 0, 1)$, compare ab with ba. Explain why the central Maslov coordinate cannot generally be discarded before the type A/type D tensor product. **Hint:** $\det(v, w) = -\det(w, v)$.

10. Selected solutions

Solution to Exercise 2

For $m + n$ even, the integers m, n have the same parity. If both are even, $(m/2, n/2) \in L$; if both are odd, $(m/2, n/2) \in L + (1/2, 1/2)$. Thus $S/L \cong \mathbb{Z}/2$. The label 1 has displacement $-(\alpha + \beta)/2$, which lies outside L; the label 12 has displacement $-\beta \in L$. Therefore 1 changes the refined idempotent coset and 12 preserves it.

Solution to Exercise 4

The quotient is

$$H_M = (\mathbb{Z}\mu \oplus \mathbb{Z}\lambda)/\langle 3\lambda \rangle \cong \mathbb{Z} \oplus \mathbb{Z}/3.$$

The geometric cover is $(\mathbb{R}^2 \setminus \mathbb{Z}^2)/\langle 3\lambda \rangle$. Translation by λ gives a deck transformation of order three, and translation by μ gives one of infinite order. Quotienting by λ would force the order-three translation to become the identity, replacing the correct deck group by $\mathbb{Z}$ and erasing torsion grading information.

Solution to Exercise 5

We have

$$\left|\det \begin{pmatrix} 1 & 3 \\ 0 & 7 \end{pmatrix}\right| = 7, \qquad H_Y \cong \mathbb{Z}/7.$$

Straight representatives yield seven transverse intersections and no immersed bigons. The quotient lattice distinguishes their relative positions: each point represents a different deck class in H_Y. Hence the seven intersection generators lie in different Spin-c summands, with one generator in each, giving rank one in each $\widehat{HF}(Y, \mathfrak{t})$.

Solution to Exercise 6

Choose components $\gamma_1, \gamma_2, \gamma_3$ and arrows $\gamma_1 \to \gamma_2$, $\gamma_2 \to \gamma_3$. This is a tree. Choose any lift of γ_1; the first arrow determines the lift of γ_2, and the second determines that of γ_3. Another choice for the initial lift translates all three by the same deck element. No additional ambiguity remains because a tree has unique paths. By HRW22, Definition 28, these arrows are phantom: they change paths used for grading comparisons, but they are deliberately ignored in the bigon count defining the ungraded differential.

Solution to Exercise 7

The values are, respectively,

$$2(0) - 1 = -1, \qquad 2(1) - 1 = 1, \qquad 2(-1) + 1 = -1.$$

The first is the standard puncture-free positive bigon whose boundary contributes to a hat Floer differential when the remaining local conditions hold. The second encloses a puncture and therefore cannot contribute to the hat differential, although it provides a grading comparison. The third is not the standard positively oriented convex-right-angle bigon used in the derivation of (7.2), so its numerical value alone is insufficient to claim a differential term.

Solution to Exercise 8

Bilinearity of det gives associativity exactly as proved in Section 8. For $a = (1/2; 1, 0)$ and $b = (1/2; 0, 1)$,

$$ab = (2; 1, 1), \qquad ba = (0; 1, 1).$$

Thus $ab = \lambda^2 ba$. Multiplying the bordered grading factors in the wrong order can shift the Maslov coordinate by two; although this vanishes modulo two, it does not vanish for the refined Maslov grading.

11. Oral-exam questions and model answers

Question Q1. Why do HRW23's graph signs change across labels 2 and 123, but not across 1 and 3? The signs are the A-parities, not the original D-parities. The conversion $g_A = g_D + \epsilon$ shifts precisely one idempotent. Labels 1, 2, 3 and 123 exchange idempotents; their extra shift cancels the D-parity change for 1 and 3 but creates one for 2 and 123. Labels 12 and 23 preserve idempotents and sign.

Question Q2. What is the exact cover carrying $\widehat{HF}(M, \mathfrak{s})$? It corresponds to the kernel of $\pi_1(T_M) \to H_1(\partial M) \to H_1(M)$, equivalently $(\mathbb{R}^2 \setminus \mathbb{Z}^2)/\langle \ell \rangle$. Its deck group is $H_M = H_1(\partial M)/\langle \ell \rangle$. The homological longitude ℓ may not be primitive; replacing it by the rational longitude can discard torsion.

Question Q3. What is the logical role of Table 2 in proving Theorem 1.7? The table identifies the refined Spin-c difference across each edge with the displacement of its standard immersed segment. Therefore the separately chosen lifted switches are joined by lifted edges. This is the essential compatibility needed in Proposition 6.5. Section 3's local simplifications then lift, yielding the graded curve set.

Question Q4. What do phantom arrows encode, and why do they not change the ungraded Floer complex? They connect distinct curve components to fix relative choices of lifts and Maslov references. They are auxiliary paths allowed in grading comparisons but expressly ignored in the immersed-bigon differential. They differ from genuine local-system crossover arrows.

Question Q5. When is the formula $m(y) - m(x) = 2N(C) + R(C)/\pi$ applicable? The intersections must have the same relative Spin-c grading, so that the chosen pair of paths lifts to a closed path C in the punctured plane. The tangent rotations are measured along smooth portions with orthogonal corners, while $N(C)$ counts lattice punctures with winding. Periodic-domain ambiguities and absolute grading choices must also be respected.

Question Q6. State the difference between HRW23, Proposition 6.7 and HRW22, Theorem 5. Proposition 6.7 refines the pairing by Spin-c classes and gives the relative mod-two grading from intersection signs. Theorem 5 of the companion additionally identifies the full **relative Maslov grading** using the winding-plus-rotation construction and the refined noncommutative bordered grading group.

12. Reading assignment and proof checkpoints

Time	Task	Deliverable
45 min	HRW23, Section 6.1, pp. 56-57; Figures 45-46	Reproduce the six parity shifts and prove Lemma 6.2
55 min	HRW23, Section 6.2, pp. 58-60; Table 2	Derive four additive shifts and the lifted-switch construction
45 min	HRW23, Section 6.3, pp. 60-61	Explain the deck-translation criterion in Proposition 6.7
60 min	HRW22, Section 2.1; Defs. 28 and 35 in arXiv v1	Explain phantom arrows and compute Maslov bigons
45 min	HRW22, Sections 2.2-2.3	Verify the noncommutative group law and grading pairing architecture
35 min	Exercises 4-8 and oral review	Solve four exercises without notes

Total: 4 h 45 min.

- ☐ Distinguish g_D from g_A and derive the six parity shifts.
- ☐ Explain why the orientation on a loop records its relative $\mathbb{Z}/2$ grading.
- ☐ Prove $S = L \sqcup (L + (\alpha + \beta)/2)$ and explain H_M.
- ☐ Reproduce HRW23, Table 2 and verify all four multiplication identities.
- ☐ State and prove Proposition 6.5 up to the standard-edge local models.
- ☐ Explain why the homological longitude, not necessarily a primitive slope, defines the grading cover.
- ☐ State Proposition 6.7 and compute the lens-space example.
- ☐ Distinguish phantom grading arrows from local-system crossover arrows.
- ☐ Derive the Maslov value $2N - 1$ for a positive convex right-angle bigon.
- ☐ Prove associativity and noncommutativity of the refined group G.

13. References and next reading

[HRW23] J. Hanselman, J. Rasmussen, L. Watson, *Bordered Floer homology for manifolds with torus boundary via immersed curves*, arXiv:1604.03466v3 (2023), Theorem 1.7 (p. 5), Section 6 (pp. 56-61), especially Proposition 6.1, Lemma 6.2, Lemma 6.4, Proposition 6.5, Corollary 6.6 and Proposition 6.7. Primary PDF.

[HRW22] The same authors, *Heegaard Floer homology for manifolds with torus boundary: properties and examples*, *Proc. London Math. Soc.* **125** (2022), 879-967, Section 2 and Theorem 5. For numbering and figures used in this note see arXiv:1810.10355v1: Definition 28 (phantom grading arrows), Definition 35 (Maslov formula), Figures 17-19 and 23-24. Published paper.

[LOT] R. Lipshitz, P. Ozsvath, D. Thurston, *Bordered Heegaard Floer Homology*, *Memoirs AMS* **254** (2018), no. 1216, Chapter 11, especially its grading formalism. This is supporting background, not the primary reading.

Next: HRW23, Section 7.1 (Dehn-filling and L-space applications) together with HRW22, Section 1.4 (products $S^1 \times \Sigma$) or Section 3 (elliptic involution). The next lecture will keep the algebra-to-geometry proof as the organizing principle rather than turning into an application-only survey.