HRW Qualifying Examination — Day 17

The full correspondence: extended type D structures, crossover arrows, decorated curves, and geometric pairing

Course date: Sunday, November 8, 2026 · **Study load:** 4 hours 50 minutes · **Coefficients:** $\mathbb{F} = \mathbb{F}_2$. This is a synthesis lecture, not an alternative proof of every diagrammatic step in HRW23. The full lecture is also attached as HRW_Day17.pdf.

1. The source-level dependency graph

Let M be a compact, connected, smooth, oriented three-manifold with $\partial M \cong T^2$, and fix the puncture $z \in \partial M$. Write $T_M = \partial M \setminus \{z\}$. The fundamental promise of the main HRW paper is that a boundary-parametrized type D module can be interpreted by a **finite collection of unobstructed immersed closed curves with finite-dimensional local systems** in T_M. Its construction is algorithmic; its classification and gluing properties require additional theorems.

> **Theorem: HRW23, Theorems 1.1, 1.2, 1.4, 1.5** If $N = \widehat{CFD}(M, \alpha, \beta)$, then N is extendable (Theorem 1.4). Every extendable type D structure is represented by decorated immersed curves (Theorem 1.5). For manifold invariants the resulting decorated curves are independent of the boundary parametrization up to regular homotopy and local-system isomorphism (Theorem 1.1). If $Y = M_0 \cup_h M_1$, the (appropriately reflected) curve sets compute $\widehat{HF}(Y)$ by immersed Lagrangian intersection Floer homology (Theorem 1.2).

The logical order is

$$\widehat{CFD}(M,\alpha,\beta) \xrightarrow[]{\text{Theorem 1.4}} \text{extendable } N$$
$$\xrightarrow[\S 3.3]{\text{Proposition 3.8}} \text{curve configuration + crossover arrows}$$
$$\xrightarrow[\S 3.6]{\text{Proposition 3.13}} \text{only parallel-strand arrows}$$
$$\xrightarrow[\S 3.7]{\text{Theorem 1.5}} \{(\gamma_i, V_i, \Phi_i)\}_i$$
$$\xrightarrow[\S 4]{\text{Theorem 4.11}} \text{classified object with pairing}$$
$$\xrightarrow[\S 5]{\text{Proposition 5.1}} \mathcal{C}(M) \text{ independent of the marking.}$$

Theorem 2.2 is a separate *chain-level bridge* from train-track bigons to the type A/type D box tensor differential. It is then used in §4 to prove the pairing-preserving classification, and in §5 to recover $\widehat{HF}$ for a glued manifold. One cannot conclude the gluing formula merely from the existence of curves.

Primary paper audit. For exact numbering use the 87-page HRW23 arXiv version 3: Theorems 1.1–1.5 (pp. 2–4); Theorem 2.2 (§2.6, p. 13); Definition 3.1 (§3.1, p. 17); Proposition 3.8 (§3.3, p. 23); Proposition 3.13 (§3.6, p. 31); Definition 3.16 and Theorem 1.5 proof (§3.7, p. 37); Theorem 4.11 (§4.2, p. 46); Proposition 4.16 (§4.3, p. 52); Proposition 5.1 (§5, pp. 53–55). The paper appeared electronically in 2023 but in *J. Amer. Math. Soc.* **37 (2024), 391–498** (DOI 10.1090/jams/1029); it is **not** the two-author *Geometry & Topology* article at pages 823–924.

2. The torus algebra and actual type D data

Fix a boundary marking (α, β) with the **HRW outside-boundary convention** $\beta \cdot \alpha = 1$. Set $\mathcal{J} = \mathbb{F}\iota_0 \oplus \mathbb{F}\iota_1$. The unital torus algebra $\mathcal{A}$ has orthogonal idempotents ι_0, ι_1 and six Reeb-chord elements

$$\rho_1,\ \rho_2,\ \rho_3,\ \rho_{12},\ \rho_{23},\ \rho_{123}.$$

The permitted source–target idempotents are

Coefficient	Start	End
$\rho_1, \rho_3, \rho_{123}$	ι_0	ι_1
ρ_2	ι_1	ι_0
ρ_{12}	ι_0	ι_0
ρ_{23}	ι_1	ι_1

Other than idempotent products, the nonzero products of two non-idempotent basis elements are **exactly**

$$\rho_1\rho_2 = \rho_{12}, \quad \rho_2\rho_3 = \rho_{23}, \quad \rho_1\rho_{23} = \rho_{123}, \quad \rho_{12}\rho_3 = \rho_{123}.$$

The algebra differential vanishes. Write a type D structure as a finite-dimensional $\mathcal{I}$-module $V = \iota_0 V \oplus \iota_1 V$ with coefficient map $\delta^1 : V \to \mathcal{A} \otimes_{\mathcal{I}} V$. The defining type D equation is

$$(\mu \otimes 1)(1 \otimes \delta^1)\delta^1 = 0.$$

Equivalently, for $N = \mathcal{A} \otimes_{\mathcal{I}} V$ the induced left $\mathcal{A}$-linear differential ∂ satisfies $\partial^2 = 0$. If a graph has an edge $x \xrightarrow{I} y$, it means $\rho_I \otimes y$ is a term of $\delta^1(x)$. A reduced graph contains no edge labelled by the algebra unit 1 (HRW23, §2.3).

> **Warning: Two serious traps** A type D structure is **not** simply an arbitrary labelled graph: the length-two coefficient paths must cancel modulo two after multiplication in $\mathcal{A}$. Conversely the algebra does not require all length-two labelled paths to cancel individually; different paths can contribute the same coefficient. Also, the algebraic operation $\rho_1\rho_2 = \rho_{12}$ must not be confused with concatenating *two separate geometric immersed-curve components*.

Derived check. Suppose $\delta^1(x) = \rho_1 \otimes y$ and $\delta^1(y) = \rho_2 \otimes z$, with no other contributions to the $\rho_{12} \otimes z$ coefficient. Then $(\mu \otimes 1)(1 \otimes \delta^1)\delta^1(x)$ contains $\rho_{12} \otimes z \neq 0$, violating the type D equation. The graph cannot be a type D module. This one-line calculation is a useful qualifying-exam sanity check.

3. Curvature, extendability, and the matrix-factorization checkpoint

The **extended torus algebra** $\widetilde{\mathcal{A}}$ is a quiver algebra obtained by restoring the chord ρ_0; see HRW23, §3.1, Figure 13 and Definition 3.1. One has

$$\mathcal{A} = \widetilde{\mathcal{A}}/\langle\rho_0\rangle, \qquad U = \rho_{1230} + \rho_{2301} + \rho_{3012} + \rho_{0123},$$

where U is **central**. An extended type D structure has induced differential $\tilde{\partial}$ satisfying the *curved* equation

$$\boxed{\tilde{\partial}^2 = U\,\mathrm{id}.}$$

The underlying ordinary type D structure is obtained by setting $\rho_0 = 0$. Since every summand of U contains ρ_0, its image is zero, and the curved relation reduces to $\partial^2 = 0$.

> **Theorem: HRW23, Definition 3.1 and Theorem 1.4** A type D module N is *extendable* when there exists an extended type D module $\widetilde{N}$ whose reduction modulo the ideal $\langle\rho_0\rangle$ is homotopy equivalent to N. The type D invariant of a torus-boundary three-manifold is extendable. This theorem does **not** state that every abstract type D structure is extendable.

HRW use an auxiliary ring $R = \mathbb{F}[U]/(U^2)$ and matrices obeying $M^2 = UI$ as a normal-form tool (§3.3, Lemma 3.9). These auxiliary matrices must not be silently identified with the literal

differential matrix over the ordinary torus algebra.

3.1. A complete 2×2 matrix factorization

Take

$$M = \begin{pmatrix} 0 & 1 \\ U & 0 \end{pmatrix} \quad \text{over } R = \mathbb{F}_2[U]/(U^2).$$

Direct multiplication gives

$$M^2 = \begin{pmatrix} U & 0 \\ 0 & U \end{pmatrix} = UI_2.$$

This is the basic paired-block model. Its mod-U reduction is $\left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$, which is an ordinary square-zero differential. **The curvature has disappeared upon reduction**, not through an illicit cancellation of U.

3.2. A nontrivial conjugation

Set $P = I_2 + E_{12} = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$. Since the characteristic is two, $P^{-1} = P$. Compute

$$PMP^{-1} = \begin{pmatrix} U & 1+U \\ U & U \end{pmatrix} = M'.$$

The new matrix is visibly different but still obeys

$$(M')^2 = PM^2P^{-1} = UI_2.$$

Indeed its upper-left entry is $U^2 + (1+U)U = U$ and the off-diagonal entries are zero. The conjugation illustrates why a presentation can change substantially while the curved algebraic object stays equivalent. It also models the elementary linear-algebra component of a crossover-arrow move.

> **Warning: Extendability versus uniqueness** Existence of a curved extension is Theorem 1.4 for manifold modules. Its *essential uniqueness up to homotopy* is a later theorem, HRW23 **Proposition 4.16**. The latter must not be assumed when proving the initial structure theorem; HRW prove it only after establishing the curve classification.

4. From matrices to tracks: the two reductions

A **naive train track** is obtained by placing the ι_0 and ι_1 generators at switches on the two marked arcs, and tracing the Reeb-chord terms as directed paths in the punctured torus (HRW23, §2.4, Figures 7–8). The local source/target directions are dictated by the idempotents. A switch is a tangent-line junction, not an ordinary transverse self-intersection.

An *embedded arc segment* in this construction does not imply the total train track is a multicurve: there may be branching switches and crossover arrows. The structure theorem proceeds in two reductions.

4.1. First reduction: valence two plus crossover arrows

> **Theorem: HRW23, Proposition 3.8** Every valid reduced train track is equivalent to one whose switches each have a single unoriented strand on each side, together with additional directed **crossover arrows** connecting strands. The geometric arrangement satisfies further restrictions on the peripheral rectangles and the central square.

The matrix normal-form lemma over $\mathbb{F}[U]/(U^2)$ permits simultaneous basis changes, including elementary transvections $I + E_{ij}$ and $I + UE_{ij}$. Geometrically these changes correspond to adding/sliding crossover arrows. A head sliding past the tail of another arrow can create a *composition arrow*; omitting it changes the path-counting operator (HRW23, §3.4, Figure 21).

Derived finite-dimensional illustration. For a differential matrix D with $D^2 = 0$ and an invertible change of basis P, $D' = PDP^{-1}$ again squares to zero. Yet D and D' can have different patterns of nonzero entries. The number of visible graph arrows is therefore *not invariant*. A crossover-arrow calculus must preserve the corresponding homotopy equivalence class, not a naive count of arrows.

4.2. Second reduction: finite-depth arrows disappear

Crossover arrows are assigned directional weights according to how long their strand germs agree before distinguishing themselves. HRW's *depth* is extracted from those forward/backward weights (see §§3.5–3.6). The statement to use is precise:

> **Theorem: HRW23, Proposition 3.13** A curve configuration of finite depth m is equivalent to a curve configuration with the same number of strands and depth at least $m + 1$.

Repeated depth improvement removes arrows connecting genuinely nonparallel strands. What survives connects strands that remain parallel indefinitely: these arrows encode **local-system monodromy**.

> **Proof: Proof reconstruction of Theorem 1.5 — incomplete diagrammatic steps labelled** Start from an extendable reduced type D structure. Proposition 3.8 produces a valence-two curve configuration with crossover arrows. Proposition 3.13 repeatedly increases the minimum depth; the finite number of strands gives a uniform finite-depth obstruction, so all remaining arrows join parallel strands. In each bundle, count directed paths across a transverse cut to obtain an invertible monodromy matrix. Replace the bundle by one immersed closed curve with a finite-dimensional local system. The resulting curve-like track represents a module homotopy equivalent to the original. **The four-step arrow-sliding algorithm and its noncreation-of-smaller-depth estimates are not proved here**; these are precisely HRW23, §3.6, pp. 34–37. The final extraction of monodromy is §3.7, Figure 35.

This is a constructive argument; saying merely "the Fukaya category classifies the objects" would hide the principal content of the paper.

5. Local systems and the classification theorem

A decorated curve is a triple (γ, V, Φ), where $\gamma : S^1 \looparrowright T_M$ is an unobstructed immersed loop (the universal-cover lift is embedded), V is finite-dimensional over $\mathbb{F}$, and $\Phi \in GL(V)$ is the monodromy. Orientation reversal of the parametrization replaces Φ by Φ^{-1}. Changing the basis of V conjugates Φ.

A packet of k parallel strands has a k-dimensional fiber. Given a cut transverse to the packet, the matrix entry A_{ij} counts directed paths from strand i to strand j modulo two. Sliding the cut changes the matrix by conjugation, not by arbitrary row reduction. This is HRW23, §3.7, Figure 35.

5.1. A basis-invariance proof

If $e = (e_1, \dots, e_k)$ is a frame and Φ is represented by A, replace the frame by $e' = eP$ for $P \in GL_k(\mathbb{F})$. For a coordinate column x' in the new frame, the old coordinate is $x = Px'$. Hence

$$\Phi(x) = APx' = P(P^{-1}AP)x',$$

so Φ has matrix $A' = P^{-1}AP$ in the new frame. Thus $\operatorname{tr} A$, $\det A$, the minimal polynomial, and the characteristic polynomial are frame-independent. This derivation is elementary linear

algebra, **not** an additional HRW theorem.

5.2. Direct sums versus geometric crossings

If two parallel components on the same primitive oriented free homotopy class carry monodromies A and B, their direct sum is represented by one curve with monodromy $A \oplus B$; see the equivalences defining $\mathcal{C}$ in HRW23, §4.2. A multiple cover of a primitive curve has a further block-cyclic replacement; it is **not** generally the same as a rank-one local system on the primitive curve. Neither operation licenses deleting an actual transverse crossing between distinct immersed curves.

> **Theorem: HRW23, Theorem 4.11 and Proposition 4.16** Isomorphism classes of decorated unobstructed immersed curves, modulo the specified homotopy, orientation, direct-sum, and primitive-reduction equivalences, are in bijection with homotopy classes of extendable type D structures. The bijection preserves the relevant pairing. An extendable type D module has an extension unique up to homotopy.

Proof architecture, not a complete reproduction. Surjectivity of the curve-to-module map uses Theorem 1.5. Pairing compatibility is Proposition 4.13 and the track pairing of Theorem 2.2. Injectivity is established by choosing *test curves* to detect both underlying free homotopy classes/multiplicities (Lemma 4.14) and local-system similarity types (Proposition 4.15). The latter uses rational canonical forms. Only after injectivity and surjectivity does uniqueness of the extension follow (Proposition 4.16). See HRW23, §§4.2–4.3, pp. 46–52.

5.3. Why monodromy matters

Over $\mathbb{F}_2$, compare

$$A = I_2, \qquad B = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}.$$

Both define rank-two local systems on the same geometric curve, but

$$\chi_A(t) = (t+1)^2, \qquad \chi_B(t) = t^2 + t + 1.$$

Since similar matrices have identical characteristic polynomials, the local systems are not isomorphic. The geometric curve and its rank alone do not recover the module. Notice also $B^2 = \left(\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ and $B^3 = I_2$, while A has order one.

6. The pairing theorem at chain level

The chain-level statement comes **before** the geometric classification theorem. Let ϑ_0 be reduced and let ϑ_1 be *special bounded* in the sense of HRW23, §2.6; the boundedness condition is needed to make the box-tensor differential a finite sum.

> **Theorem: HRW23, Theorem 2.2** The intersections of the geometric type A realization of ϑ_0 with the type D realization of ϑ_1 form a complex $C(\vartheta_0, \vartheta_1)$ canonically isomorphic as a chain complex to $N_0^A \boxtimes N_1$ (under the theorem's reduced and special-bounded hypotheses).

The type A realization and the type D realization are **not identical pictures**: HRW reflect across the anti-diagonal and swap the 1 and 3 labels when extracting the type A operations (see §2.3 and Figure 9).

6.1. Why each bigon is a box-tensor differential term

An intersection generator is an idempotent-compatible tensor $x \otimes y$. A bigon from $x \otimes y$ to $x' \otimes y'$ traces boundary paths on the A and D tracks. The D path gives an iterated coefficient term $\delta^k(y) = \rho_{I_1} \otimes \cdots \otimes \rho_{I_k} \otimes y'$, while the A path gives $m_{k+1}(x, \rho_{I_1}, \ldots, \rho_{I_k}) = x'$. HRW verify that the two boundary words agree after the prescribed $1 \leftrightarrow 3$ re-labelling and maximal-subword decomposition. Small bigons represent unit-labelled edges. Thus every immersed

contributing bigon corresponds to one box-tensor term and conversely. **This is a proof sketch of the nontrivial local word matching**, fully checked in HRW23, §2.6, Figures 10–12.

6.2. Why ordinary Floer homology appears later

Theorem 2.2 matches *train-track complexes* and box tensor products. Those train tracks need not be unobstructed immersed curves. Theorem 4.11 and its proof identify curve-like tracks with immersed-curve complexes equipped with parallel transport and handle admissibility for annuli. HRW23, §5 then uses the bordered pairing theorem and parametrization independence to obtain

$$\boxed{\widehat{HF}(M_0 \cup_h M_1) \cong HF(\mathcal{C}(M_0), \overline{h}(\mathcal{C}(M_1))).}$$

The original $\overline{h}$ includes the elliptic involution. HRW22's conjugation/elliptic symmetry (Theorem 7) lets one replace $\overline{h}$ with h **for actual manifold curve invariants**, not automatically for all abstract extendable modules.

> **Warning: Admissibility and self/parallel pairings** Do not equate the Floer dimension with raw intersections for arbitrary immersed curves. Periodic annular domains can force an admissible perturbation, and monodromy can contribute even when minimal geometric intersection is zero. HRW23, Definition 4.7, Theorem 4.9, and Corollary 4.10 handle these cases.

7. Computation laboratory: matrix factorization and local systems

7.1. Check the four algebra products

The torus algebra's non-idempotent products are $\rho_1\rho_2 = \rho_{12}$, $\rho_2\rho_3 = \rho_{23}$, $\rho_1\rho_{23} = \rho_{123}$, $\rho_{12}\rho_3 = \rho_{123}$. In particular,

$$(\rho_1\rho_2)\rho_3 = \rho_{123} = \rho_1(\rho_2\rho_3).$$

This is an explicit associativity check for the one nontrivial length-three word. Attempting to extend it by ρ_1 on the right fails because the corresponding concatenation is not an allowed increasing Reeb-chord string.

7.2. Detect similarity through a test curve

Let $A = I_2$ and $B = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 1 \end{smallmatrix}\right)$ as above. Pair each with the *oppositely oriented* rank-one primitive curve carrying trivial monodromy. HRW23, Corollary 4.10 (after the orientation convention) yields the local-system term

$$2\dim\ker(I + A) \quad \text{or} \quad 2\dim\ker(I + B).$$

For $A = I_2$, $I + A = 0$, so the contribution has dimension $2 \cdot 2 = 4$. For B, $I + B = \left(\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ has determinant one and trivial kernel, so the corresponding contribution is zero. Thus a Floer *test pairing* distinguishes the two monodromies even though they have the same underlying primitive curve and the same fiber rank. No standalone manifold realization is claimed for either object.

Important qualification. The formula used is the simple primitive embedded-curve case of the corollary. For an immersed primitive curve with minimal self-intersections $i(\gamma) > 0$, the full formula has an additional $2kk'i(\gamma)$ term. For noncommensurable curves, the corresponding formula is different.

8. Computation laboratory: fillings and parallel curves

8.1. Two straight slopes

Let $u = (1, 0)$ and $v = (2, 5)$ in an oriented integral basis of the boundary torus. Their minimal geometric intersection number is

$$\Delta(u, v) = \left|\det \begin{pmatrix} 1 & 2 \\ 0 & 5 \end{pmatrix}\right| = 5.$$

For embedded geodesics representing these primitive slopes, there are no immersed bigons. For trivial rank-one local systems,

$$\dim HF(u, v) = 5.$$

When the curves are meridian invariants of solid tori with an appropriate gluing identification, this recovers the rank of $\widehat{HF}$ of a lens-space filling with $|H_1| = 5$. The calculation is a derived special case of HRW23, Corollary 4.10; **it is not a general rank formula for immersed multicurves**.

If the first straight curve has a rank-two local system and the second rank three, the transverse intersection complex has dimension $5 \cdot 2 \cdot 3 = 30$, and still has zero differential because the slopes are noncommensurable geodesics. This does not assert these particular local systems come from solid tori.

8.2. Parallel embedded primitive curves

Let γ and γ' be the same free homotopy class of embedded primitive circles, with ranks k, k' and monodromies A, A'. After a small admissible perturbation, there are two geometric intersections for each fiber pair; bigon parallel transport gives a linear map $\Phi = I \otimes I + A \otimes A'$ (after orienting the second curve oppositely as in the exact statement of Corollary 4.10). The resulting two-term differential contributes homology dimension

$$\boxed{\dim HF(\gamma_A, \gamma'_{A'}) = 2 \dim \ker \Phi.}$$

The formula shows **minimal geometric intersection zero does not imply Floer homology zero**. For $k = k' = 1$ and trivial monodromies over $\mathbb{F}_2$, $\Phi = 1 + 1 = 0$, so the Floer dimension is two. This is the local-system/annular exception used repeatedly in HRW23's rank bounds.

9. What the 2022 companion adds

The companion HRW22 is **not** merely a source of examples. It establishes geometric structures needed to recover the refinements and prove special manifold constraints.

> **Theorem: HRW22, Theorems 5 and 7, Proposition 10** Theorem 5 identifies the relative Maslov grading in the immersed-curve pairing with the relative grading of closed-manifold Heegaard Floer homology. Theorem 7 proves elliptic-involution invariance for *actual* three-manifold invariants via Spin^c conjugation. Proposition 10 excludes peripheral components: the invariant of a torus-boundary manifold has **no component isotopic to a small circle around the puncture**.

The last fact is decisive. HRW23, Theorem 1.5 classifies all *extendable* modules, but not every such abstract module comes from a three-manifold. A peripheral small circle is an algebraically meaningful object but is forbidden for manifold invariants (HRW22, §3, Proposition 10). The proof uses an additional minus-type enhancement and a four-step coefficient relation; the ordinary curved equation $\tilde{\partial}^2 = U$ alone does **not** exclude it.

HRW22, Theorem 11 relates the knot exterior's curves to $CFK^-(K)$; Proposition 12 recovers $\widehat{HFK}$ of a filling core via a *proper* slope arc, as distinct from the closed slope curve used for $\widehat{HF}$ of the filling. HRW22, §5 relates orientations and lifts to Turaev torsion, the Thurston norm, and L-space intervals. Thus the papers have different tasks:

HRW23 :	existence, classification, pairing, gluing;
HRW22 :	gradings, symmetry, knots, torsion, norm, examples.

A grading warning. The mod-two sign of an oriented intersection does not by itself give an absolute Maslov integer. For two generators in a fixed relative Spin^c summand, HRW22 uses a closed lifted path C and the winding/turning formula

$$m(y) - m(x) = 2N(C) + \frac{R(C)}{\pi}.$$

The winding number $N(C)$ counts lattice punctures with signed multiplicity, and $R(C)$ is the clockwise rotation along smooth portions, excluding corners. The formula must not be applied across unrelated Spin^c summands or without its grading conventions.

10. Oral-defense proof checkpoints

A fully defensible proof of Theorem 1.1 separates four logically independent tasks: (i) *existence* of an extension for manifold modules (Theorem 1.4); (ii) *normal form* of an extendable module as a curve with local systems (Theorem 1.5, via Propositions 3.8/3.13); (iii) *injectivity and pairing* of the classification (Theorem 4.11, via test pairings and Theorem 2.2); (iv) *independence of the marking* using Dehn-twist bimodules (Proposition 5.1). The gluing statement then uses the bordered pairing and the geometry of admissible immersed-curve complexes.

The six proof-checkpoint questions are: **(P1)** Where exactly is $\tilde{\partial}^2 = U$ used? **(P2)** Why does the reduction not terminate at arbitrary train tracks? **(P3)** What survives an infinite-depth arrow elimination? **(P4)** How does one reconstruct an abstract module from Φ? **(P5)** Why is Theorem 2.2 not already Theorem 1.2? **(P6)** What additional manifold input rules out peripheral circles?

11. Exercises

Exercise: Exercise 1 — Torus-algebra consistency (15 min) List all nonzero products of two non-idempotent torus-algebra generators. Construct a three-generator proposed type D graph $x \xrightarrow{1} y \xrightarrow{2} z$ and show why it fails unless an additional path cancels the ρ_{12} term. **Hint:** calculate $\delta^2(x)$ coefficientwise.

Exercise: Exercise 2 — Curved reduction and conjugation (20 min) Over $R = \mathbb{F}_2[U]/U^2$, compute M^2, PMP^{-1}, and $(PMP^{-1})^2$ for $M = \left(\begin{smallmatrix} 0 & 1 \\ U & 0 \end{smallmatrix}\right)$, $P = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$. Describe precisely what survives after $U = 0$. **Hint:** $P = P^{-1}$.

Exercise: Exercise 3 — Local-system equivalence (25 min) Show that a basis change conjugates monodromy. For $A = I_2$ and $B = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 1 \end{smallmatrix}\right)$, compute characteristic/minimal polynomials and prove nonisomorphism. Explain what orientation reversal does. **Hint:** compare the irreducible polynomial $t^2 + t + 1$ with $(t+1)^2$.

Exercise: Exercise 4 — A Floer test for monodromy (25 min) Apply the primitive parallel-curve formula to the local systems A, B of Exercise 3 and a rank-one trivial test curve. Compute both kernels of $I + A$ and $I + B$. Explain why geometric intersection number alone cannot distinguish the local systems. **Hint:** over $\mathbb{F}_2$, $I + I = 0$.

Exercise: Exercise 5 — The exact proof dependency (30 min) State the hypotheses of HRW23 Theorems 2.2 and 4.11. Explain why "N has a train-track presentation" does **not** yet imply "N is determined by a decorated multicurve." Identify where surjectivity, injectivity, local-system identification, and parametrization independence enter. **Hint:** consult HRW23, §2.6, §§3.3–3.7, §§4.2–4.3, and §5.

Exercise: Exercise 6 — Pairing and a lens space (20 min) For primitive slopes $(1, 0)$ and $(3, 7)$, calculate the geometric intersection number and the rank-one Floer dimension. Derive $H_1(Y) \cong \mathbb{Z}/7$ for a gluing of solid tori with these meridians. Explain why two arbitrary nonprimitive immersed curves require more care. **Hint:** use the index of the sublattice and the absence of bigons between nonparallel straight

lines.

Exercise: Exercise 7 — Abstract versus realizable (25 min) Explain why the peripheral small circle cannot appear in $\mathcal{C}(M)$ for a three-manifold, although Theorem 1.5 is stated for abstract extendable modules. Identify the additional source theorem excluding it, and distinguish it from extendability. **Hint:** HRW22, Proposition 10 and its proof in §3.

Exercise: Exercise 8 — Maslov and annotation audit (20 min) For a positively oriented convex right-angle bigon enclosing no puncture, show that $N(C) = 0$, $R(C) = -\pi$, and therefore $m(y) - m(x) = -1$. Repeat with one enclosed puncture. State why the second region does not enter the hat differential and why either result may fail to compare distinct Spin^c summands. **Hint:** count the two omitted corner turns.

12. Hints and worked answers

Answer 1. The four nontrivial products are $\rho_1\rho_2 = \rho_{12}$, $\rho_2\rho_3 = \rho_{23}$, $\rho_1\rho_{23} = \rho_{123}$, and $\rho_{12}\rho_3 = \rho_{123}$. For the displayed chain, $(\mu \otimes 1)(1 \otimes \delta^1)\delta^1(x) = \rho_{12} \otimes z$. A different length-two path giving the *same coefficient* is required to cancel this over $\mathbb{F}_2$.

Answer 2. $M^2 = UI_2$. Since $P^{-1} = P$, direct multiplication gives $PMP^{-1} = \left(\begin{smallmatrix} U & 1+U \\ U & U \end{smallmatrix}\right)$. Its square is UI_2. Modulo U the first matrix becomes $\left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$ and the second $\left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$; both square to zero. The two presentations differ over R but not at this reduced displayed level.

Answer 3. $A' = P^{-1}AP$ under a basis change. The characteristic and minimal polynomials of $A = I_2$ are $(t+1)^2$ and $t+1$; for B both equal t^2+t+1. Hence they are not similar. Orientation reversal changes the monodromy to its inverse; it does not allow arbitrary replacement by an unrelated matrix.

Answer 4. $I + A = 0$ has kernel dimension 2 and yields parallel primitive Floer dimension 4. $I+B = \left(\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ has determinant one and zero kernel, giving dimension zero for the indicated test pairing. Both geometric curve pairs have the same underlying slopes, so intersection number alone cannot distinguish them.

Answer 5. Theorem 2.2 assumes a *reduced* first train track and *special bounded* second track, and identifies intersection/bigon **chain complexes** with a box tensor product. Theorem 4.11 classifies homotopy classes of **extendable** type D modules via unobstructed decorated curve sets; it uses Theorem 1.5 for surjectivity and probe pairings for injectivity. Proposition 5.1 then proves boundary marking independence for actual manifold invariants.

Answer 6. $\Delta((1,0),(3,7)) = |7| = 7$. Straight representatives have no bigons and yield Floer rank 7. The Dehn gluing relations kill a subgroup of index 7 in the boundary lattice, giving $H_1(Y) \cong \mathbb{Z}/7$ for the resulting lens-space gluing. Nonprimitive curves may be multiple covers and carry nontrivial monodromies, so Corollary 4.10 must be applied through its primitive-reduction rules.

Answer 7. Theorem 1.5 concerns all extendable algebraic objects; it is not a realization theorem asserting every such object is $\widehat{CFD}(M)$. HRW22, Proposition 10 uses the minus-type four-step relation to rule out a peripheral circle with any nonzero invertible local system. This is genuinely stronger than the existence of an ordinary curved extension.

Answer 8. A positively oriented convex right-angle bigon has total smooth counterclockwise turning π after removing two $\pi/2$ corner jumps, hence clockwise $R = -\pi$. Thus $m(y)-m(x) = 2(0) - 1 = -1$. One enclosed puncture gives $2(1) - 1 = 1$; it is not counted by the hat differential because the region contains a puncture. The relative grading comparison requires the appropriate fixed Spin^c summand and lift.

13. Oral questions with model answers

Question: Q1. What exactly is the structure theorem, and what is its hypothesis? Model answer. HRW23, Theorem 1.5 says an *extendable* type D structure over the torus algebra can be represented by unobstructed immersed closed curves equipped with finite-dimensional local systems. Extendability is required; Theorem 1.4 establishes it for manifold-derived $\widehat{CFD}$ modules.

Question: Q2. Why can the matrix factorization not be read as an arbitrary differential on $\mathcal{A}$? Model answer. The $M^2 = UI$ matrices occur in an auxiliary normal-form reduction for the extended curved algebra. U is the central curvature. Passing to the ordinary algebra kills ρ_0 and hence U, giving a square-zero ordinary differential; the auxiliary 2×2 blocks are not themselves a literal definition of all ordinary type D modules.

Question: Q3. What is the mathematical role of finite-depth arrow elimination? Model answer. Proposition 3.8 first makes the train track a union of immersed strands with crossover arrows. Proposition 3.13 raises the minimal depth. Finite-depth arrows that join distinguishable strands can be removed, leaving only arrows joining globally parallel strands; those become the monodromy of local systems.

Question: Q4. How is the monodromy encoded and why is it not just a rank? Model answer. A cut through a packet of parallel strands gives a vector space spanned by strands. Directed-path counts define an invertible matrix A; basis changes conjugate it. The rank is the dimension, but the similarity class of A is independent information. Pairing with carefully selected test curves can distinguish similarity types.

Question: Q5. Where is the bigon-to-box-tensor identification actually proved? Model answer. In HRW23, Theorem 2.2, under reduced/special-bounded track hypotheses. The proof matches intersection generators with idempotent-compatible tensor generators and matches immersed bigons with paired type A operations and iterated type D coefficients. Theorem 1.2 additionally requires the curve classification, admissibility, marking invariance, and the bordered pairing theorem.

Question: Q6. Explain the elliptic involution caveat. Model answer. The original geometric pairing formula uses the gluing map composed with an elliptic involution because of the A/D reflection convention. HRW22, Theorem 7 shows actual manifold curve invariants are elliptically symmetric, so one may write h alone for them. That symmetry is not a property of every abstract extendable module.

Question: Q7. What prevents every algebraic curve from coming from a three-manifold? Model answer. Extendability is necessary but not a complete geometric realization criterion. HRW22, Proposition 10 excludes a peripheral small circle from a manifold invariant using extra minus-type data. HRW23, Proposition 7.13 also forbids an entirely loose invariant when the boundary is incompressible. These are genuinely manifold-level restrictions.

Question: Q8. Why is minimal geometric intersection not always the Floer dimension? Model answer. For noncommensurable primitive unobstructed curves in geodesic position, Corollary 4.10 gives rank equal to geometric intersection times the local-system ranks. For parallel or commensurable components, periodic annuli and monodromy create additional Floer homology: two parallel trivial circles have zero geometric minimum but two-dimensional Floer homology.

14. Reading assignments and verified references

Four-hour-fifty-minute schedule (290 minutes): 35 min HRW23 §§2.1–2.6 (algebra, tracks, Theorem 2.2); 60 min HRW23 §§3.1–3.3 (curvature and Proposition 3.8); 55 min HRW23 §§3.4–3.7 (arrow weights, Proposition 3.13, Theorem 1.5); 45 min HRW23 §§4–5 (Theorem 4.11, Proposition 5.1); 35 min HRW22 §§2–3 (gradings, elliptic symmetry, peripheral obstruction); 60 min selected exercises, written proof reconstruction, and oral defense.

Source-checking protocol. The citations below distinguish the three-author main immersed-curve paper from the different *two-author* papers on bordered Floer calculus and cabling, which should not be substituted for it. The theorem numbers and PDF page locations of HRW23 were checked directly against arXiv:1604.03466v3. The companion citation's journal identity was checked against its publisher record; its arXiv numbering can differ from the 2022 published version.

1. **[HRW23]** J. Hanselman, J. Rasmussen, L. Watson, *Bordered Floer homology for manifolds*

with torus boundary via immersed curves, arXiv:1604.03466v3, April 25, 2023; electronically published August 23, 2023, *Journal of the American Mathematical Society* **37 (2024)**, 391–498, DOI:10.1090/jams/1029. Source statements used: Theorems **1.1, 1.2, 1.4, 1.5, 2.2, 4.11**; Definition **3.1**; Propositions **3.8, 3.13, 4.16, 5.1**; Corollary **4.10**; §§2–5.

2. **[HRW22]** J. Hanselman, J. Rasmussen, L. Watson, *Heegaard Floer homology for manifolds with torus boundary: properties and examples*, *Proceedings of the London Mathematical Society* **125 (2022)**, 879–967, DOI:10.1112/plms.12473, arXiv:1810.10355. Source statements used: **Theorems 5, 7, 11** and **Propositions 10, 12**, §§2–4. The companion's gradings/elliptic theorem numbers refer to the published exposition where applicable.
3. **[LOT18, background only]** R. Lipshitz, P. Ozsváth, D. Thurston, *Bordered Heegaard Floer Homology*, *Memoirs AMS* **254 (2018)**, no. 1216, DOI:10.1090/memo/1216. It supplies the bordered pairing and the torus algebra; the HRW papers remain today's primary subjects.

Next: Day 18. We will interrogate the classification's *injectivity* and *rigidity* more closely: how Floer test curves detect free homotopy class, rank, and monodromy; how admissibility works when curves are parallel; and why abstract curve objects must be distinguished from manifold-realizable ones.