Day 11 - Turaev Torsion, Signed Rays, and the Thurston Norm

Course date: Monday, November 2, 2026. **Study budget:** 4 h 55 min. **Level:** graduate / oral qualifying examination. **Coefficient field:** $\mathbb{F} = \mathbb{F}_2$ for Heegaard Floer theory; torsion and Euler characteristics are integer-valued.

> **Citation and scope policy** Theorem/lemma numbers below are those of the **published 2022 companion paper**. Theorems 13 and Proposition 14 are in its introduction; Lemma 61, Proposition 62, Corollary 64, Proposition 65, Lemma 66, Proposition 67, Lemma 68 and Theorem 69 are in Section 5. The paper often written "HRW23" here is the **three-author preprint arXiv:1604.03466**, not the different two-author *A calculus for bordered Floer homology* (Geom. Topol. 27 (2023), 823-924). We mark every independent algebraic calculation as **derived**. We mark the deeper knot-Floer detection and pegboard claims as **proof sketches** when their original diagrammatic or external theorem inputs are not reproved.

1. Reading map and today's theorem chain

The 2023 main work constructs the geometric invariant and establishes its grading cover. The 2022 companion then extracts classical topology from that geometry. Our main input from **HRW23, Section 6** is the Spin-c-refined immersion $\mathcal{C}(M,\mathfrak{s})$ in $\overline{T}_{M,\mathfrak{s}}$, oriented using the relative $\mathbb{Z}/2$ grading. From **HRW22, Section 4**, we use the filling-core pairing with a properly immersed test arc $L_{\mu,\bar{\mathfrak{s}}}$.

Today's three implications are

$$\boxed{\begin{array}{l}\text{oriented, Spin-c-refined immersed curves} \xrightarrow{\text{signed ray intersection}} \tau_M,\\ \text{pegboard extremes} + \widehat{HFK} \text{ detection} \longrightarrow x_M([\Sigma]),\\ \text{Floer simplicity} + \tau_M \in \{0,1\} \longrightarrow \text{L-space interval endpoints.}\end{array}}$$

> **Source theorem map HRW22, Theorem 13 (introduction) / Proposition 62:** Turaev torsion equals signed intersections of the oriented lifted curve with downward rays from lifted punctures. **HRW22, Proposition 14 (introduction) / Proposition 65:** the obstruction to reaching an end of the grading cylinder detects the complexity of a norm-minimizing surface. **HRW22, Theorem 69:** for Floer-simple, non-solid-torus-like M, the interior L-space filling interval is the maximal interval avoiding a set of slopes determined by torsion. Each statement has hypotheses; none follows from unsigned intersection numbers alone.

The exact passage from the three-dimensional Euler characteristic to a signed 1-cycle is the conceptual center of the lecture. There are *two* mathematical ingredients: an algebraic finite-difference equation from sutured torsion and a geometric finite-difference equation from relative intersections. Matching their initial conditions proves the result.

2. Spin-c torsors, homology orientations, and three kinds of grading

Let M be compact, connected, oriented, with ∂M a torus and chosen puncture z. Let $\mathcal{C}(M,\mathfrak{s})$ denote the **decorated immersed multicurve** representing the fixed $\mathfrak{s} \in \mathrm{Spin}^c(M)$ summand. It is not the closed-manifold vector space $\widehat{HF}(Y)$.

There are three related but distinct affine sets. The ordinary $\mathrm{Spin}^c(M)$ is an affine $H^2(M;\mathbb{Z})$-set. The relative $\mathrm{Spin}^c(M,\partial M)$, defined by outward boundary vector fields, is an affine $H^2(M,\partial M) \cong H_1(M)$-set. The sutured $\mathrm{Spin}^c(M,\gamma_\mu)$ uses a **different** boundary condition for sutures of slope μ. There is no canonical identification of the latter two torsors; HRW22 **Lemma 61** constructs two compatible relative lifts $\bar{\mathfrak{s}}_\pm$.

Let $j_* : H_1(\partial M) \to H_1(M)$ be inclusion, and let λ be a generator of its kernel (the **homological longitude**). It need not be primitive as a boundary class. Choose a primitive μ with $[\Sigma]\cdot\mu > 0$, where $[\Sigma] \in H_2(M,\partial M)$ and $\partial[\Sigma] = \lambda$. The homology orientation in HRW22, Section 5.1, is specified by this choice, or equivalently by a positive generator for the rational image of $H_1(\partial M) \to H_1(M)$.

The graded multicurve lives in the cover

$$\overline{T}_{M,\mathfrak{s}} = (\mathbb{R}^2 \setminus \mathbb{Z}^2)/\langle\lambda\rangle, \qquad \mathrm{Deck}(\overline{T}_{M,\mathfrak{s}}/T_M) \cong H_M = \mathrm{im}(j_*).$$

Filling in the lifted punctures gives a cylinder with a **positive** and a **negative** end, fixed by the choice of μ. Lifts of the puncture are labelled $z_{\hat{\mathfrak{s}}}$ by relative Spin-c structures mapping to $\mathfrak{s}$.

> **The sign and boundary conventions are essential** HRW22 uses $\lambda\cdot\mu > 0$ **on the boundary of the complement**, whereas common knot-theory conventions often use $\mu\cdot\lambda > 0$ on the tubular-neighborhood boundary. Reversing the orientation of the cylinder or of the ray changes signed intersections. Also $\mathrm{Spin}^c(M,\partial M)$ and $\mathrm{Spin}^c(M,\gamma_\mu)$ are not literally the same affine set.

3. Relative Turaev torsion and the sutured difference identity

Once the homology orientation is fixed, Turaev torsion is an integer-valued function

$$\tau_M : \mathrm{Spin}^c(M, \partial M) \longrightarrow \mathbb{Z}.$$

For $b_1(M) > 1$, it is finitely supported. For $b_1(M) = 1$, the normalization of HRW22 is $\tau_M(\hat{\mathfrak{s}}) = 0$ sufficiently far toward the negative end and $\tau_M(\hat{\mathfrak{s}}) = 1$ sufficiently far toward the positive end (within the relevant torsor component). These are asymptotic statements, not a claim that τ_M only takes values zero or one for all manifolds.

The Euler characteristic of sutured Floer homology decategorifies to the torsion $\tau_{(M,\gamma_\mu)}$. The underlying foundational input is Friedl–Juhasz–Rasmussen, *The decategorification of sutured Floer homology*. Relative to a compatible choice of origins, the group-ring relation is

$$\sum_{\bar{\mathfrak{s}}} \tau_{(M,\gamma_\mu)}(\bar{\mathfrak{s}})[\bar{\mathfrak{s}}] \sim (1 - [\mu]) \sum_{\hat{\mathfrak{s}}} \tau_M(\hat{\mathfrak{s}})[\hat{\mathfrak{s}}]. \tag{3.1}$$

Here $\sim$ includes an **overall translation** by $H_1(M)$. Lemma 61 fixes it compatibly with the two ends of each test arc.

> **HRW22, Lemma 61 (source statement)** Every $\bar{\mathfrak{s}} \in \mathrm{Spin}^c(M, \gamma_\mu)$ determines $\bar{\mathfrak{s}}_\pm \in \mathrm{Spin}^c(M, \partial M)$ such that
>
> $$\bar{\mathfrak{s}}_+ - \bar{\mathfrak{s}}_- = [\mu],$$
>
> these assignments commute with the affine $H_1(M)$-action, and
>
> $$\boxed{\tau_{(M,\gamma_\mu)}(\bar{\mathfrak{s}}) = \tau_M(\bar{\mathfrak{s}}_+) - \tau_M(\bar{\mathfrak{s}}_-).}$$

A rigorous algebraic check of (3.1) uses the coefficient of a class h: multiplication by $(1-[\mu])$ yields the difference $\tau(h) - \tau(h - [\mu])$. The orientation of the plus/minus labels then fixes which ordered difference is used.

This is the first finite-difference equation. Our next task is to show that oriented intersections obey the *same* equation.

4. The signed downward-ray intersection function

Choose the absolute orientation of $\mathcal{C}(M, \mathfrak{s})$ by matching the $\mathbb{Z}/2$-graded Euler characteristic of the filling-core pairing with sutured torsion, as in HRW22, Section 5.1. Do **not** orient each component arbitrarily. A rank-r local system contributes with multiplicity r to signed Euler characteristic, but the full monodromy conjugacy class is invisible to the resulting homology class.

For each relative $\hat{\mathfrak{s}}$ choose a path $\gamma_{\hat{\mathfrak{s}}}$ starting at its lifted puncture $z_{\hat{\mathfrak{s}}}$ and tending to the **negative** end, transverse to $\mathcal{C}(M, \mathfrak{s})$. Define

$$\boxed{n_{\hat{\mathfrak{s}}} = [\gamma_{\hat{\mathfrak{s}}}] \cdot [\mathcal{C}(M, \mathfrak{s})] \in \mathbb{Z}.} \tag{4.1}$$

Since $\mathcal{C}$ is compactly supported in the cylinder, the signed intersection is well defined as a relative/locally finite intersection pairing. Deforming a ray without crossing its specified puncture leaves the intersection number unchanged; two newly created transverse intersections have opposite signs.

Derived local cancellation lemma. Suppose a transverse isotopy of a test ray crosses a tangency with an oriented curve and produces two transverse intersections. Their local intersection signs are opposite, so the algebraic count is unchanged. Proof: choose oriented coordinates where the curve is the x-axis and the moving ray is a parabola $y = x^2 - \varepsilon$; its two crossings at $x = \pm\sqrt{\varepsilon}$ have opposite derivative signs. This is why unsigned intersection counts are the wrong torsion invariant.

The relative homology classes of the rays satisfy (HRW22, proof of Proposition 62)

$$[\gamma_{\hat{\mathfrak{s}}+[\mu]}] - [\gamma_{\hat{\mathfrak{s}}}] = -[L_{\mu,\bar{\mathfrak{s}}}]. \tag{4.2}$$

With the intersection convention of HRW22, equations (4.1)–(4.2) give

$$\boxed{n_{\hat{\mathfrak{s}}+[\mu]} - n_{\hat{\mathfrak{s}}} = [\mathcal{C}(M, \mathfrak{s})] \cdot [L_{\mu,\bar{\mathfrak{s}}}].} \tag{4.3}$$

The orientation on the test arc $L_{\mu,\bar{\mathfrak{s}}}$ runs from its negative to its positive end. This is the second finite-difference equation.

5. Proof of HRW22 Proposition 62

HRW22, Proposition 62; Theorem 13 in the introduction For every $\widehat{\mathfrak{s}} \in \mathrm{Spin}^c(M, \partial M)$,

$$\boxed{\tau_M(\widehat{\mathfrak{s}}) = \gamma_{\widehat{\mathfrak{s}}} \cdot \mathcal{C}(M, \mathfrak{s}).}$$

Here $\mathfrak{s}$ is the restriction of $\widehat{\mathfrak{s}}$ to M, and the curves, rays, cylinder ends and homology orientation carry the conventions just fixed.

Proof of Proposition 62, with the foundational inputs identified The filling-core pairing from HRW22, Proposition 12 (and the refined knot-Floer construction in Section 4), identifies the signed Euler characteristic of the Floer pairing with the sutured Floer Euler characteristic. Consequently

$$[\mathcal{C}(M, \mathfrak{s})] \cdot [L_{\mu, \bar{\mathfrak{s}}}] = \chi(SFH(M, \gamma_\mu, \bar{\mathfrak{s}})) = \tau_{(M, \gamma_\mu)}(\bar{\mathfrak{s}}).$$

The second equality is the sutured Floer decategorification theorem, an **external theorem**.

By (4.3) and HRW22, Lemma 61,

$$\begin{aligned} n_{\widehat{\mathfrak{s}}+[\mu]} - n_{\widehat{\mathfrak{s}}} &= \tau_{(M, \gamma_\mu)}(\bar{\mathfrak{s}}) \\ &= \tau_M(\widehat{\mathfrak{s}} + [\mu]) - \tau_M(\widehat{\mathfrak{s}}). \end{aligned}$$

Thus the difference $n - \tau_M$ is constant along each $[\mu]$ orbit. Move sufficiently far into the negative end along this orbit. Compactness of $\mathcal{C}$ means the downward ray no longer meets it, so $n = 0$ there. The normalized torsion also vanishes at that end, so $\tau_M = 0$. Therefore the constant is zero. Since the orbit was arbitrary, $n = \tau_M$ everywhere. □

The key logical point: the theorem is not a claim that signed intersections with a closed test slope *directly* equal torsion. One must use noncompact rays, their choice of end, and the normalization. The proof is a uniqueness argument for solutions to a discrete derivative equation with an initial condition.

6. Alexander series: three fully computed examples

Assume for this section that $M = S^3 \setminus \nu K$, $H_1(M) = \mathbb{Z}$, and choose the usual symmetric normalization $\Delta_K(t) = \Delta_K(t^{-1})$, $\Delta_K(1) = 1$. Label relative Spin-c structures by $j \in \mathbb{Z} + \frac{1}{2}$ with the homology-orientation normalization of HRW22. Write $\tau_j = \tau_M(\widehat{\mathfrak{s}}_j)$ and

$$T_K(t) = \sum_{j \in \mathbb{Z}+1/2} \tau_j t^j.$$

HRW22, introduction, equation preceding Theorem 13: as a formal Laurent series expanded in increasing powers,

$$\boxed{T_K(t) = \frac{\Delta_K(t)}{t^{-1/2} - t^{1/2}} = \frac{t^{1/2}\Delta_K(t)}{1 - t}.} \tag{6.1}$$

Derived discrete derivative lemma If $\Delta_K(t) = \sum_{n \in \mathbb{Z}} a_n t^n$, then

$$\boxed{a_n = \tau_{n+1/2} - \tau_{n-1/2}.}$$

Proof: multiply (6.1) by $t^{-1/2} - t^{1/2}$ and compare coefficients of t^n. Conversely, the difference equation together with $\tau_j = 0$ for $j \ll 0$ uniquely determines T_K.

Example A: unknot. With $\Delta_U(t) = 1$,

$$T_U(t) = \frac{t^{1/2}}{1 - t} = t^{1/2} + t^{3/2} + t^{5/2} + \cdots.$$

Hence $\tau_j = 0$ for $j \le -1/2$ and $\tau_j = 1$ for $j \ge 1/2$. The single jump is the Euler characteristic of $\widehat{HFK}(U)$ in Alexander grading zero.

Example B: right-handed trefoil. From Day 10, $\Delta_{T(2,3)} = t^{-1} - 1 + t$. Then

$$\begin{aligned} T_{T(2,3)}(t) &= \frac{t^{-1/2} - t^{1/2} + t^{3/2}}{1 - t} \\ &= t^{-1/2} + \frac{t^{3/2}}{1 - t} \\ &= t^{-1/2} + t^{3/2} + t^{5/2} + \cdots. \end{aligned}$$

Thus

$$\tau_j = \begin{cases} 0, & j \le -3/2, \\ 1, & j = -1/2, \\ 0, & j = 1/2, \\ 1, & j \ge 3/2. \end{cases}$$

The successive jumps at integer indices $-1, 0, 1$ are $+1, -1, +1$, recovering $t^{-1} - 1 + t$. These are exactly the **signed** ray-intersection values reported in HRW22, introduction, Figure 3.

Example C: figure-eight knot. Its symmetrized Alexander polynomial is $\Delta_{4_1}(t) = -t^{-1} + 3 - t$. Therefore

$$\begin{aligned} T_{4_1}(t) &= \frac{-t^{-1/2} + 3t^{1/2} - t^{3/2}}{1 - t} \\ &= -t^{-1/2} + 2t^{1/2} + t^{3/2} + t^{5/2} + \cdots. \end{aligned}$$

In particular $\tau_{-1/2} = -1$, $\tau_{1/2} = 2$, $\tau_j = 1$ for $j \ge 3/2$, and $\tau_j = 0$ for $j \le -3/2$. **This refutes the tempting assertion that Turaev torsion is always zero or one.** The $\{0, 1\}$ feature in Section 5.3 relies on Floer simplicity, not on the general definition of torsion.

Example D (extra): $T(2, 5)$. The symmetric polynomial is $t^{-2} - t^{-1} + 1 - t + t^2$. The derivative rule gives successively the torsion values $1, 0, 1, 0, 1$ at $j = -3/2, -1/2, 1/2, 3/2, 5/2$; beyond the last jump the value remains 1. This is a genus-two example, in contrast with the trefoil's genus one. It will be used to cross-check the pegboard norm theorem.

7. Homological projection and information lost by decategorification

HRW22, Corollary 64 Let $p : \overline{T}_{M,\mathfrak{s}} \to T_M$ denote projection. With the source's orientation convention,

$$p_*[\mathcal{C}(M, \mathfrak{s})] = \begin{cases} \lambda, & b_1(M) = 1, \\ 0, & b_1(M) > 1. \end{cases}$$

Proof (reconstructed). In the puncture-filled cylinder, a cross-section joining the ends intersects the homology class of the core longitude with algebraic count one. For $b_1 = 1$, the lower and upper tails of torsion are respectively zero and one; the ray formula therefore gives total algebraic crossing number one. For $b_1 > 1$, torsion is finitely supported in each affine orbit, so its two tails agree (both zero); the total intersection vanishes. The projection class is fixed by its intersection with a transverse section. This is the argument of HRW22 immediately following Proposition 62. □

The full geometric invariant remembers the regular-homotopy classes, components and local-system similarity classes, whereas τ_M only remembers **their signed homology class in the punctured grading cover** (HRW22, Remark 63). In particular, two nonisomorphic local systems of the same rank on the same oriented curve have the same torsion contribution. Local systems are invisible to signed decategorification except through their fiber ranks.

> **Do not conflate three quantities** (i) $\#(L \cap \mathcal{C})$ is an **unsigned chain-generator count**; (ii) $HF(L, \mathcal{C})$ may have a differential; (iii) $L \cdot \mathcal{C}$ is the **signed intersection/Euler characteristic**. Turaev torsion is recovered from (iii), not from (i). For the figure-eight example, negative and greater-than-one values already show that naive geometric counts are inadequate.

8. Pegboard geometry and the Thurston norm

For a properly embedded compact oriented surface $\Sigma \subset M$, define

$$\chi_-(\Sigma) = \sum_{\Sigma_i} \max\{0, -\chi(\Sigma_i)\}, \qquad x_M(a) = \min_{[\Sigma]=a} \chi_-(\Sigma).$$

For a **connected** genus-g surface with one boundary component, $\chi(\Sigma) = 1 - 2g$. Thus

$$-\chi(\Sigma) = 2g - 1, \qquad x_M([\Sigma]) = \max\{0, 2g - 1\}. \tag{8.1}$$

The difference matters for disks: if $g = 0$, then $-\chi = -1$ but the Thurston norm is **zero**. HRW22, Proposition 65 is stated using $-\chi(\Sigma)$ for the norm-minimizing surface, and the boundary-compressible/solid-torus-like case is handled separately. Do not indiscriminately substitute the Thurston norm in all formulas.

Pull the lifted curves tight around a grid of punctures (pegs). For each lifted puncture $z_{\hat{\mathfrak{s}}}$, ask whether there exists a path in the **complement** of the curves from that puncture toward the positive end of the grading cylinder. Define

$$\boxed{k^+(\mathcal{C}(M,\mathfrak{s})) = \max\{\langle c_1(\hat{\mathfrak{s}}), [\Sigma]\rangle : z_{\hat{\mathfrak{s}}} \text{ cannot reach } +\infty \text{ without crossing } \mathcal{C}\}.} \tag{8.2}$$

This is the geometric obstruction height of HRW22, Section 5.2. The associated k^- is defined analogously using the negative end, after normalizing the lift and its height coordinates. A pegboard **corner** is a puncture around which a tightened curve changes direction.

HRW22, Proposition 65; Proposition 14 in the introduction Let $[\Sigma] \in H_2(M, \partial M)$ with $\partial[\Sigma] = \lambda$, and let Σ be norm-minimizing, in the setting of the proposition. Then

$$\boxed{-\chi(\Sigma) = \max_{\mathfrak{s}\in \mathrm{Spin}^c(M)} k^+(\mathcal{C}(M,\mathfrak{s})).}$$

In the introductory rank-one formulation, if k_+ is the upper obstruction height and k_- the lower one, then $k_+ - k_- = 2g(\Sigma) - 1$.

Caution about normalizations: The k^+ in the general proposition is an evaluation of a relative first Chern class; the schematic half-integer *vertical coordinates* used for knot-complement plots are not identical to these integer evaluations without the source's chosen affine normalization. In particular, $2g-1$ should not be equated naively with the highest ordinary Alexander grading g.

9. Proof architecture of HRW22 Proposition 65

The proof has a geometric part and a Floer-detection part. The geometric part can be described precisely; the final detection step invokes the external genus/Thurston-norm theorem for knot Floer homology.

Step 1: highest pegs and blocked ends. Suppose a tightened curve has at least one pegboard corner, and let z_* be its highest corner peg. Immediately above the level of z_*, every peg can be connected to the positive end by a path avoiding the curve. The peg z_* itself is separated from that end by an arc of the curve. Hence the maximal blocked evaluation is the evaluation of a highest corner peg. This uses minimal-position geometry of the tightened pegboard.

Step 2: a longitudinal test arc. Let l be the *primitive rational longitude* and let $L_{l,k}$ be the noncompact test arc of its slope passing through the puncture at relative first-Chern evaluation k. HRW22, Theorem 52 from Day 10 identifies

$$HF(\mathcal{C}(M,\mathfrak{s}), L_{l,k}) \cong \widehat{HFK}(M(l), K_l, \hat{\mathfrak{s}}_k). \tag{9.1}$$

Here K_l is the core of rational-longitude filling. The test arc is not the closed slope curve calculating $\widehat{HF}(M(l))$.

Step 3: support at the highest corner. Above the highest peg, the tightened curve has no intersections with the test arc, hence the pairing is zero. At the highest corner, one sees a nontrivial intersection class; in the local pegboard analysis it cannot be killed by a competing differential. Thus the largest k with nonvanishing pairing is the largest corner-peg evaluation. This is the figure-dependent argument in HRW22, proof of Proposition 65, not a formal consequence of signed torsion alone.

Step 4: invoke norm detection. The knot Floer theorem for the filled core identifies the largest relative Spin-c evaluation with $-\chi(\Sigma)$ for the norm-minimizing surface. This uses results of Ozsvath–Szabo and Yi Ni; HRW's proof does not reprove the general norm-detection theorem.

Step 5: control solid-torus-like components. A component with no corners can lie between adjacent rows of pegs and can have zero longitudinal test-arc Floer homology, yet still block access to an end. HRW addresses this special case separately: if every component is solid-torus-like, use the boundary-compressible result cited there; if a solid-torus-like component sits above the highest genuine corner, a different filling-core pairing and norm detection give a contradiction. This exclusion is necessary; dropping it would create an unjustified proof gap.

Proof-status disclosure The five steps accurately reproduce the logical organization of HRW22, Proposition 65. The *corner pairing is nontrivial* statement is a **diagrammatic proof sketch** here, and the final detection of the Thurston norm is **an external theorem**. The algebraic computations of Sections 4–6 are complete independently of those deeper inputs.

Worked genus checks. The trefoil has Seifert genus $g = 1$: $-\chi(\Sigma) = 1$, so Proposition 65 predicts maximal obstruction evaluation 1. The torus knot $T(2,5)$ has $g = 2$, so $-\chi(\Sigma) = 3$, predicting maximal obstruction evaluation 3. These are evaluations of c_1 in the source's normalization, not the heights 1 and 2 of the top hat

knot-Floer generators. For the unknot, the norm-minimizing surface is a disk, $-\chi = -1$ while $x_M = 0$; this is the crucial boundary-compressible exception to naive norm-language.

Why torsion alone need not detect genus. τ_M is an *Euler characteristic* and can have cancellations. The support of $\widehat{HFK}$ is finer. For example the symmetrized Alexander polynomial of a knot may be 1 even when the knot has positive genus. The pegboard theorem accesses the full knot-Floer support, not merely the signed intersection class.

10. Floer-simple manifolds and L-space intervals

A rational-homology-sphere Y is an L-space if $\dim_{\mathbb{F}} \widehat{HF}(Y) = |H_1(Y)|$. Let

$$\mathcal{L}(M) = \{\alpha \in \mathrm{Sl}(\partial M) : M(\alpha) \text{ is an L-space}\}$$

and let $\mathcal{L}^\circ(M)$ denote its interior in the slope circle. A torus-boundary manifold is **Floer simple** if $\mathcal{L}^\circ(M) \neq \emptyset$.

HRW22, Lemma 66. If $\alpha \in \mathcal{L}^\circ(M)$ and coordinates are chosen so that α is vertical, each tightened lift $\widetilde{\mathcal{C}}(M, \mathfrak{s})$ is a graph $y = f(x)$, with no vertical segments. Sketch: a generic vertical line meets the lifted curve once because the corresponding filling is an L-space; minimal position prevents extra intersections. A vertical pegboard segment would put α among the tangent or singular slopes, contrary to being interior.

The graph separates the plane into two regions. HRW colors pegs **black** on the negative side and **white** on the positive side. By Proposition 62, black pegs have torsion 0, white pegs torsion 1. This is the geometric reason a Floer-simple torsion function behaves like a two-coloring. It is not true for the figure-eight knot, which is not Floer simple.

Let l be the primitive rational longitude and j_* the map $H_1(\partial M) \to H_1(M)$. Define

$$D^\tau(M) = \{\hat{\mathfrak{s}}_0 - \hat{\mathfrak{s}}_1 : \tau_M(\hat{\mathfrak{s}}_i) = i,\ i = 0, 1,\ l \cdot (\hat{\mathfrak{s}}_0 - \hat{\mathfrak{s}}_1) \geq 0\} \subset H_1(M). \tag{10.1}$$

Set $\mathbb{P}(j_*^{-1} D^\tau(M)) \subset \mathrm{Sl}(\partial M)$ to be the slopes of primitive representatives in the indicated preimage. Here $\mathbb{P}$ means **projectivization of nonzero classes**; we do not include the zero vector as a slope.

> **HRW22, Theorem 69 (torsion characterization)** If M is Floer simple and **not solid-torus-like**, and $\alpha \in \mathcal{L}^\circ(M)$, then $\mathcal{L}^\circ(M)$ is the largest interval in the slope circle containing α and avoiding $\mathbb{P}(j_*^{-1} D^\tau(M))$. This is HRW's geometric new proof of the corresponding Rasmussen–Rasmussen torsion theorem.

Proof architecture. HRW22, Proposition 67 identifies the singular/tangent slope interval with the smallest interval spanned by slopes joining a black peg to a white peg in the downward ordering. The inclusion uses the intermediate value theorem: a ray from a black peg toward a white peg must cross the graph and, farther along the ray, return to the black region, so somewhere the tangent slope equals that ray's slope. For the reverse inclusion, an extremal graph slope occurs on a pegboard edge whose endpoints are controlled by opposite-colored pegs. Lemma 68 identifies these geometric peg differences with $j_*^{-1} D^\tau(M)$. Finally the interior L-space slopes are the complement of the singular slopes. This last chain is a **proof sketch** because the extremal pegboard configurations and the L-space/singular-slope theorem are cited from HRW, not reproduced in full.

> **Do not drop the hypotheses** Theorem 69 assumes Floer simplicity and excludes the solid-torus-like case. It identifies the **interior interval**; boundary slopes require separate analysis. The set D^τ is built from relative torsion values, not by selecting coefficients of the Alexander polynomial without the Spin-c/lattice information.

11. Pitfalls and dependence on the two HRW papers

Source distinction. HRW23 (three authors, arXiv:1604.03466, Sections 4–6) gives the immersed-curve invariant, grading cover, and pairing machinery. HRW22 (three authors, *Proc. Lond. Math. Soc.* 125 (2022)) derives Proposition 62, Proposition 65, and Theorem 69. The latter are **not** presented as theorems of the separate two-author *A calculus for bordered Floer homology.*

Signs. Turaev torsion is signed; orientation of curves is fixed by a homology orientation, not chosen component by component. A ray goes to the *negative* end. Comparing puncture lifts requires a fixed affine origin.

The genus distinction. For a knot exterior $-\chi(\Sigma) = 2g - 1$, whereas the top Alexander grading is g; their equality would be false without the appropriate first Chern class normalization. For the unknot, $x_M = 0$ but $-\chi(\text{disk}) = -1$.

Torsion versus complete curves. The homology class of a weighted decorated multicurve is a decategorification. It cannot reconstruct local-system monodromy, immersed self-intersections, or differential information; nor does the degree of Δ_K always equal genus.

The L-space interval. Floer simple is a hypothesis; $\tau \in \{0, 1\}$ and the black/white peg picture do not extend to arbitrary M.

12. Serious exercises

Exercise 1 - Spin-c torsors and torsion differences Starting from $\mathrm{Spin}^c(M, \partial M)$ and $\mathrm{Spin}^c(M, \gamma_\mu)$, explain why they cannot simply be identified. Derive the coefficient formula for $(1-[\mu])\sum_s \tau_M(s)[s]$, then reconstruct all three parts of HRW22, Lemma 61. **Hint:** an affine origin must be chosen to interpret the group-ring identity as an equality of functions.

Exercise 2 - Ray intersection invariance Prove directly that the signed intersection number of a compact oriented multicurve with a proper downward ray is invariant under compactly supported isotopies of the ray fixing its initial puncture. Give a local coordinate calculation for a birth/death of two intersections. Why does the same statement fail for the unsigned count? **Hint:** use $y = x^2 - \varepsilon$.

Exercise 3 - Proof of the torsion theorem Write a complete proof of HRW22, Proposition 62 assuming Lemma 61 and the filling-core pairing/Euler-characteristic identification. State exactly where compactness and the negative-end normalization enter. Explain how the proof adapts when $b_1(M) > 1$. **Hint:** compare two finite-difference equations along a fixed $[\mu]$ orbit.

Exercise 4 - Four Laurent-series calculations Compute the torsion series and the discrete derivative for the unknot, right-handed trefoil, figure-eight knot, and $T(2,5)$. Check $\sum a_n = 1$ in each case. Why does the figure-eight calculation contradict the claim $\tau \in \{0, 1\}$ for arbitrary knot exteriors? **Hint:** $T(t) = t^{1/2}\Delta(t)/(1-t)$, expanded toward increasing exponents.

Exercise 5 - Homology class of the lifted multicurve Deduce HRW22, Corollary 64 from Proposition 62. Explain why the stable torsion values at the two cylinder ends determine the coefficient of the longitudinal homology class. Discuss the $b_1 = 1$ and $b_1 > 1$ cases separately. **Hint:** intersect with a positive-to-negative proper cross-section.

Exercise 6 - Surface genus and normalized heights For connected one-boundary surfaces of genus $g = 0, 1, 2, 3$, compute $-\chi$, χ_-, and the value predicted by Proposition 65 when its hypotheses apply. Explain why the usual top Alexander grading g is not the same numerical convention as $\langle c_1, [\Sigma]\rangle$. Identify the extra proof step needed to turn a highest peg into a nonzero Floer class. **Hint:** use HRW22, Section 5.2 and the knot-Floer norm detection theorem.

Exercise 7 - Floer-simple slope geometry Reconstruct the proof of HRW22, Lemma 66 and explain why an interior L-space slope forbids vertical pegboard segments. State $D^\tau(M)$ precisely, including the sign inequality. Explain how a black-white peg pair yields a potential singular slope, and why the zero vector must be excluded in projectivization. **Hint:** use Proposition 67 and Lemma 68.

Exercise 8 - Detect information lost by torsion Exhibit two nonisomorphic rank-two local systems on the same oriented essential loop with the same signed torsion contribution. Give explicit matrices over $\mathbb{F}_2$, prove they are not similar, and explain why torsion nevertheless agrees. **Hint:** compare I_2 with $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 1 \end{smallmatrix}\right)$.

13. Hints and selected solutions

Solution 2. In a local oriented chart, take the curve as $y = 0$ and the moving arc as $y = x^2 - \varepsilon$. For $\varepsilon > 0$ the two crossings are at $\pm\sqrt{\varepsilon}$. The oriented normal derivative has opposite signs at the two roots, so their local intersection indices sum to zero. For $\varepsilon < 0$ there are no crossings. The unsigned count jumps from zero to two. Other generic isotopy events do not change the signed count.

Solution 3. Define $d_s = n_s - \tau(s)$. The paired finite-difference identities give $d_{s+[\mu]} = d_s$. Thus d is constant on every $[\mu]$ orbit. Translate sufficiently far toward the negative end: the compact curve is disjoint from the relevant downward rays, so $n = 0$. The normalized torsion has $\tau = 0$ there. Hence $d = 0$ throughout that orbit. This handles all orbits, including the multiple affine components when $b_1 > 1$.

Solution 4. For the trefoil, $T = t^{-1/2} + t^{3/2}/(1-t)$. Thus the jump sequence at $n = -1, 0, 1$ is $1, -1, 1$. For the figure-eight knot,

$$T = -t^{-1/2} + 2t^{1/2} + t^{3/2}/(1-t),$$

so the jumps are $-1, 3, -1$, and the intermediate torsion value 2 is unavoidable. For $T(2,5)$, the nonzero jumps $+1, -1, +1, -1, +1$ occur at degrees $-2, -1, 0, 1, 2$; the torsion is $1, 0, 1, 0, 1$ through the five successive half-integer levels. In every case $\Delta(1) = 1$.

Solution 6. If Σ has one boundary component, $\chi = 1 - 2g$. The table of $(-\chi, \chi_-)$ is $(-1, 0)$, $(1, 1)$, $(3, 3)$, $(5, 5)$ for genera $0, 1, 2, 3$. Proposition 65 applies to the signed $-\chi$ complexity, with the disk handled as a special boundary-compressible case. A highest peg contributes to Floer **homology**, not merely to the chain complex, only after showing the local intersection generator survives the differential; that is the pegboard argument in HRW22, Proposition 65.

Solution 8. Take $A = I_2$ and $B = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 1 \end{smallmatrix}\right)$. They are not similar: A has characteristic polynomial $(t+1)^2$, while B has irreducible characteristic polynomial $t^2 + t + 1$. Both have rank two, so both contribute twice the signed

homology class of the underlying oriented curve. Thus their torsion contributions agree even though their local systems are not isomorphic.

14. Oral examination with model answers

Q1. Explain torsion from curves in 90 seconds. The relative torsion is a function of outward-boundary Spin-c structures. Its finite difference under translation by $[\mu]$ equals the sutured Floer Euler characteristic. The signed intersection of a downward puncture ray with the oriented lifted immersed curve has exactly the same difference, by the filling-core pairing and the relation between adjacent rays. Both vanish at the negative end. Therefore they agree everywhere: HRW22, Proposition 62.

Q2. Why is an *absolute* orientation of the curve needed? Type D mod-two grading initially determines relative orientations. To get the integer torsion rather than its negative one fixes the global sign using a homology orientation and matches the curve/test-arc intersection Euler characteristic with sutured torsion. Arbitrary componentwise orientation reversals would corrupt the invariant.

Q3. Why not simply count intersections? Torsion is a signed Euler characteristic. Positive and negative transverse intersections can cancel, and homology also involves a differential. The figure-eight computation has torsion values -1 and 2, demonstrating that the invariant is not an unsigned count.

Q4. How can the same curve data detect genus when the Alexander polynomial sometimes cannot? The Alexander polynomial is an Euler characteristic and loses cancellation information. The curve invariant records the filtered knot-Floer support. HRW22, Proposition 65 uses the highest obstructing peg together with the filling-core knot Floer pairing and the external knot-Floer detection of the Thurston norm; the extra data is geometric and homological, not merely torsion.

Q5. What is the disk/unknot pitfall? The Thurston norm uses $\max(0, -\chi)$, so a disk has norm zero. But $-\chi(D^2) = -1$. The statement of Proposition 65 is phrased using $-\chi$ of a norm-minimizing surface and handles solid-torus-like boundary-compressible cases separately. Replacing $-\chi$ by x_M without qualification is wrong.

Q6. Why is torsion two-valued in the Floer-simple argument but not for every knot? A non-singular interior L-space slope makes each component of the lifted tightened multicurve a graph intersecting a generic vertical line exactly once. The graph separates punctures into two sides; Proposition 62 assigns torsion zero and one to the black and white sides. A generic knot exterior need not have an open interval of L-space fillings; the figure-eight torsion is a counterexample to unrestricted two-valuedness.

Q7. State the exact L-space interval theorem. HRW22, Theorem 69: for a Floer-simple, non-solid-torus-like torus-boundary manifold and an interior L-space slope α, the interior L-space interval is the largest slope interval containing α and avoiding the projectivized nonzero preimage under j_* of the torsion difference set $D^\tau(M)$. The theorem does not automatically decide endpoints and requires all of its stated hypotheses.

15. Reading assignment, proof checkpoints, and next lecture

Study schedule (total: 4 h 55 min):

Time	Paper and task	What to reproduce without notes
35 min	HRW23, Section 6; HRW22, Section 5.1 opening	Grading cover, relative Spin-c torsors and homology orientation
55 min	HRW22, Lemma 61 and Proposition 62	Both difference equations and the negative-end normalization
45 min	HRW22, Theorem 13 and Remark 63/Corollary 64	Trefoil, figure-eight series and projected homology class
60 min	HRW22, Section 5.2, Proposition 65	Pegboard corners, exception, and the external norm-detection input
45 min	HRW22, Section 5.3, Lemma 66–Theorem 69	Graph property; black/white pegs; D^τ and slope interval
35 min	Exercises 2, 3, 4, 8 and oral questions	Write full proofs and computations

Proof checkpoints:

- ☐ Explain why the two relative Spin-c torsors have different boundary conditions.
- ☐ Specify the homology orientation and the negative-end normalization.
- ☐ Derive the group-ring finite-difference formula from $(1 - [\mu])\tau$.
- ☐ Prove invariance of signed ray intersections under local isotopy.
- ☐ Reproduce the proof of HRW22, Proposition 62.
- ☐ Calculate the trefoil and figure-eight torsion coefficient by coefficient.
- ☐ Reconstruct HRW22, Corollary 64 from end behavior.
- ☐ State the precise hypothesis and the disk exception in Proposition 65.
- ☐ Explain the pegboard corner-to-HFK support argument and identify its unproved input.
- ☐ State Theorem 69 without omitting Floer simplicity or the solid-torus-like exception.

Next lecture, Day 12 (November 3): Torsion-computed L-space intervals and their behavior under slope changes, with explicit Floer-simple examples and a return to the geometric pairing from HRW23.

Exact references

[**HRW22**] J. Hanselman, J. Rasmussen and L. Watson, *Heegaard Floer homology for manifolds with torus boundary: properties and examples*, *Proceedings of the London Mathematical Society* **125** (2022), 879–967. Introduction: Theorem 13, Proposition 14. **Section 5.1:** Lemma 61, Proposition 62, Remark 63, Corollary 64. **Section 5.2:** Proposition 65. **Section 5.3:** Lemma 66, Proposition 67, Lemma 68, Theorem 69. Published theorem numbering is used throughout.

[**HRW23**] J. Hanselman, J. Rasmussen and L. Watson, *Bordered Floer homology for manifolds with torus boundary via immersed curves*, arXiv:1604.03466, **Section 4** (classification/pairing), **Section 5** (invariance), **Section 6** (Spin-c-refined lifts and orientations). This is the **three-author immersed-curve paper**, supporting the geometric construction; it is not the source of the new torsion/norm theorem numbers above.

[**FJR**] S. Friedl, A. Juhasz and J. Rasmussen, *The decategorification of sutured Floer homology*, source for the identification of the Euler characteristic of sutured Floer homology with its torsion invariant.

[**Ni**] Y. Ni, *Link Floer homology detects the Thurston norm*, *Geometry & Topology* **13** (2009), 2991–3019, supporting the norm-detection input used in the general setting.

[**MathOverflow**] Does knot Floer homology detect knot genus in rational homology spheres? includes a pointer to Ni's link-Floer theorem. It is supplementary discussion, **not** a source for the numbered HRW statements.