Confluent by Construction: Seven Cross-Kernel Demonstrations of Knot Topology in the Heyting Substrate

Richard Goodman, with the IAOM research workforce

Institute for Applied Ontological Mathematics
contact@ontologicalmathematics.org

May 1, 2026

Abstract

We exhibit seven concrete, executable demonstrations of a unified technique for computational knot topology realised across six independent verification kernels and one runtime substrate. The demonstrations are: a length-three Cayley-walk operator-confluence witness across an ordered green and an ordered divergent kernel pair (Demo F); a live cross-kernel same-knot confluence witness binding two trefoil construction paths to a categorical lax 2-cell (Demo C); a kernel-certified PHG/PGA grade-product local-degree bound and global sparsity theorem (Demo H); a triple-witness Jones-polynomial evaluation at the primitive fifth root of unity for four small knots, with one measured hardware row inherited from a Quantinuum receipt (Demo A); a multi-invariant discriminator over the pinned KnotInfo prime-knot corpus up to twelve crossings, achieving more than ninety-nine percent pairwise discrimination while explicitly surfacing the canonical Conway / Kinoshita–Terasaka mutation-blindness boundary (Demo B); a Pisano–Cayley heat-kernel spectral bridge in which the spectral and combinatorial heat traces agree to additive tolerance 10^{-9} on five moduli (Demo G); and an explicit Khovanov categorification trace on the trefoil, with non-zero Frobenius edge maps and an executable $d^2 = 0$ check together with Euler-characteristic agreement against the unnormalised Kauffman bracket (Demo E). Every demonstration ships with a substitution-body negative probe whose failure is recorded as evidence of carrier-faithfulness, an explicit *structural* versus *substantive* scope classification, and a deferment pointer for what is not claimed. The unifying technique is a finite Cayley pair-lattice on six discrete kernels equipped with a lax 2-functor categorical witness, a divergence ledger, and a multi-invariant binding to the standard knot-invariant zoo. Two threads connect the discrete substrate to a continuous one: the surreal-number kernel embeds the combinatorial game tree into a real-closed field, and the Pisano–Cayley spectral bridge couples cyclic Cayley spectra to a heat-semigroup observable. We describe how the technique lowers from formal theorems into the Boundary substrate's six-layer rule architecture, what the seven demonstrations teach about the present capabilities of the substrate, and a programme for extending Cayley analysis from *syntactic* confluence (programs that reach the same final state) to *semantic* equivalence (terms that mean the same thing) so that runtime networks may collapse along meaning-preserving rewrites.

Contents

1 Introduction **3**

2 Mathematical Objects **5**

2.1 Knots, planar diagrams, and signed PD graphs 5

2.2 Frobenius algebra over $\mathbb{Z}$ and the Khovanov complex . . . 5
2.3 Conway tangles and the surreal-finite kernel . . . 6
2.4 The six cross-system kernels . . . 6
2.5 Cayley graphs and the 30-pair lattice . . . 6
2.6 The lax 2-functor categorical witness . . . 7
2.7 The PHG / PGA grade-product surface . . . 8
2.8 Pisano periods and cyclic Cayley graphs . . . 8

3 The Discrete–Continuous Bridge **8**
3.1 Surreal numbers as the discrete–continuous handle . . . 8
3.2 Pisano–Cayley as a discrete spectral observable . . . 9
3.3 Putting them together . . . 9

4 Cayley Confluence as a Technique **9**
4.1 Why Cayley, why pair-lattice . . . 9
4.2 Determinism, BFS order, and stable identifiers . . . 10
4.3 The divergence ledger and the *green vs. divergent* verdict . . . 10

5 The Seven Demonstrations **10**
5.1 Demo F: Cayley-walk operator-confluence witness . . . 11
5.2 Demo C: Cross-kernel same-knot confluence witness . . . 11
5.3 Demo H: PHG / PGA grade-product sparsity certificate . . . 12
5.4 Demo A: Triple-witness Jones-at-fifth-root . . . 13
5.5 Demo B: Multi-invariant discriminator with mutation-pair boundary . . . 13
5.6 Demo G: Pisano–Cayley heat-kernel spectral bridge . . . 14
5.7 Demo E: Khovanov categorification trace on the trefoil . . . 14

6 Boundary Substrate and Lowering Pathway **16**
6.1 The six-layer architecture . . . 16
6.2 Lowering the seven demonstrations . . . 17
6.3 What we are not exposing . . . 17

7 Faithfulness Discipline **17**
7.1 Substitution-body negative probes . . . 17
7.2 Scope-classification: structural vs. substantive . . . 17
7.3 Anti-hack constraints . . . 17

8 Innovation, Distinctness, Implications **18**
8.1 What current alternatives lack . . . 18
8.2 Implications for human practitioners . . . 19

9 Future Work: From Syntactic Confluence to Semantic Equivalence **19**
9.1 Today: syntactic confluence . . . 19
9.2 Tomorrow: semantic equivalence . . . 19
9.3 Other extensions . . . 20

10 Open Obligations and Honest Limitations **21**
10.1 Substantive vs. structural, in one place . . . 21
10.2 Trust assumptions . . . 21
10.3 Reproduction . . . 21

1 Introduction

A formal proof of a mathematical statement is a typed term in a kernel; an executable computation is a sequence of reductions in a runtime. Bridging the two is the chronic problem of formal methods: a theorem is only as useful as the runtime to which it lowers, and a runtime is only as trustworthy as the theorems it preserves. The standard answer—extract code from a kernel proof, compile, run—works for individual lemmas but degrades as the proof base spans multiple verification kernels, multiple representations, and multiple targets. In topology, the symptom is acute: a knot can be presented as a planar diagram, a braid word, a tangle expression, a Conway tangle decomposition, a virtual or framed lift, a Cayley-graph orbit of generators, or an algebraic invariant; the relationships between these are well understood mathematically but rarely tracked at runtime.

This paper describes a technique that addresses this fragmentation directly, and it ships seven concrete demonstrations of that technique that an external auditor can run, byte-check, and adversarially perturb. The technique is a *Cayley pair-lattice* on six discrete kernels, equipped with a categorical lax 2-functor witness, a divergence ledger, and a multi-invariant binding to standard topological data. The six kernels are deliberately chosen so that knot content can be carried in each of them: a Boundary calculus runtime (`boundary_kernel`), a Negation-of-the-Word logic-of-form surface (`notw_surface`), a Conway-tangle algebra (`tangle_algebra`), a surreal-finite numeric quotient (`surreal_finite`), a Bi-Heyting three-chain regime algebra (`biheyting_3chain`), and a finite wave-amplitude algebra (`wave_finite`). Each ordered pair of distinct kernels is a row in a 30-row pair lattice with a green-or-divergent classification; each row carries a categorical 2-cell witness and a divergence-code ledger when divergent.

The seven demonstrations. Phrased non-technically, the demonstrations answer seven questions an external reader is entitled to ask of a technique that claims unifying scope:

F. *Does the technique survive a non-trivial walk through the lattice?* Demo F runs a length-three operator walk through one green pair and one divergent pair, comparing source-side and image-side endpoints step by step, and matches the divergent steps against a divergence-code ledger.

C. *Does the same knot live consistently across kernels?* Demo C computes the trefoil bracket through two independent Lean construction paths, hashes the results, asserts byte equality, and pairs the result with a categorical 2-cell witness from the lax functor.

H. *Does the technique come with a proof-complexity certificate?* Demo H proves a tight local-degree bound $\delta \leq d-1$ on the cross-system Cayley graph at depth d and the corresponding global sparsity theorem inside the kernel; the executable receipt audits both the proved bound and an honest density baseline against the complete directed simple graph.

A. *Does the technique reach hardware?* Demo A computes the Jones polynomial at the primitive fifth root of unity for four small knots three ways: by classical KnotInfo evaluation, by Burau–Fibonacci anyon trace, and by a measured Quantinuum hardware shot for the trefoil row; the other three knots carry the hardware-pending status honestly.

B. *Does the technique surface its own failure modes?* Demo B runs a multi-invariant discriminator over all KnotInfo prime knots up to twelve crossings; it achieves a discrimination rate above 99% and lists the canonical Conway / Kinoshita–Terasaka mutation pair as an explicit failure of the implemented invariant fingerprint, with a deferment pointer to refined-Khovanov and Heegaard–Floer machinery.

G. *Does the discrete substrate carry a continuous observable?* Demo G builds the cyclic Cayley graph of $\mathbb{Z}/\pi(m)\mathbb{Z}$ at the Pisano period $\pi(m)$ for $m \in \{2, 3, 5, 7, 11\}$ and computes the

heat trace $\operatorname{Tr} e^{-tL}$ two ways—spectrally via `numpy.linalg.eigvalsh` and combinatorially via $\sum_k (t^k/k!) \operatorname{Tr} A^k$—and verifies agreement to additive tolerance 10^{-9}.

E. *Does the technique reach categorification?* Demo E ships an explicit cube of resolutions for the trefoil with non-zero Frobenius edge maps drawn from the Frobenius algebra $\mathbb{Z}[x]/(x^2)$, an executable $d^2 = 0$ check, and a kernel-proved equality between the graded Euler characteristic of the resulting chain complex and the unnormalised Kauffman bracket of the trefoil.

Why now, and why this combination. Each ingredient already exists in mathematics. Cayley graphs are old. Lax 2-functors are old. Khovanov homology is twenty-five years old. Pisano periods are eight hundred years old. The novelty is not in the ingredients but in their *simultaneous* ingredient list under a single executable scope-discipline that is hostile to its own failure modes. Each demonstration ships a substitution-body negative probe: replace the load-bearing carrier with the most trivial inhabitant of its type (the zero linear map, the identity 2-functor, the constant period $\pi(m) := m$), and the supposedly content-bearing theorem must *fail to elaborate* or the supposedly content-bearing receipt must *flip its verdict*. A passing demo without a failing probe is not evidence of correctness; it is evidence of types. We treat this distinction as a first-class engineering invariant.

Two threads from discrete to continuous. The discrete-to-continuous bridge is one of the oldest themes in topology, and it appears here twice. First, Conway's surreal numbers [10] are constructed from combinatorial games, yet contain $\mathbb{R}$ as a real-closed subfield; the `surreal_finite` kernel is a finite quotient of this construction, and its presence in the pair lattice gives every other kernel a path into a continuous numeric carrier. Second, the Pisano–Cayley spectral bridge couples a discrete cyclic group $\mathbb{Z}/\pi(m)\mathbb{Z}$ to the heat-semigroup observable $t \mapsto \operatorname{Tr} e^{-tL}$; this is a discrete spectral graph theory [26] side by a continuous diffusion-time observable side, and the agreement of the spectral and combinatorial heat traces is the audit that the bridge is faithful.

Boundary substrate and bare-metal arc. Every demonstration is grounded in a Lean 4 / Mathlib formalisation [36, 37] that is in turn lowered into a Boundary calculus rule via a six-layer substrate [43]. Trust in the substrate is distributed across rules: each rule carries its own semantic-preservation theorem rather than concentrating in a privileged kernel layer. We describe the architecture at a level sufficient to ground confidence in the lowering pathway without exposing implementation details.

What we do not claim. We are not claiming a universal knot classifier; the Conway / Kinoshita–Terasaka pair is an explicit failure of the implemented fingerprint, and Demo B records this as such rather than papering over it. We are not claiming substantive global confluence on every divergent pair; the divergence ledger is the deferment pointer. We are not claiming arbitrary-diagram Khovanov homology; Demo E ships an explicit categorification trace on the trefoil with the generic carrier acknowledged as a downstream API shell. The seven demonstrations are *structural* witnesses of the technique under an explicit scope-discipline; the substantive analogues are tracked, not claimed.

Document map. Section 2 fixes the mathematical objects and notation. Section 3 develops the discrete–continuous bridge. Section 4 treats Cayley confluence as a technique. Section 5 presents the seven demonstrations. Section 6 describes the Boundary substrate and the lowering pathway. Section 7 is the substitution-body discipline that protects all of the above from carrier-vacuity. Section 8 compares the technique to the alternatives it displaces. Section 9 lays

out the roadmap from syntactic confluence to semantic equivalence. Section 10 catalogues open obligations. References follow.

2 Mathematical Objects

2.1 Knots, planar diagrams, and signed PD graphs

A *tame knot* is an embedding $K\colon S^1 \hookrightarrow S^3$ taken up to ambient isotopy [46]. We work with planar diagrams: regular projections to $\mathbb{R}^2$ in which crossings are decorated by over/under information [1, 47]. A *planar diagram* (PD) of an oriented link is a finite list of *crossings*, each crossing recording the four boundary edges in a fixed convention, together with a sign $\mathrm{sgn} \in \{+1, -1\}$. We use the *signed PD graph* encoding: a finite signed multigraph with the four-edges-per-vertex constraint, a chosen rotation system at each vertex, and a $\mathbb{Z}/2$ over/under marking. This is the format in which all six kernels carry knot content. The Jones polynomial [2] and its quantum-invariant generalisations [48] live downstream of this encoding.

Definition 2.1 (signed PD graph)**.** A *signed PD graph* is a tuple (V, E, ρ, σ) where V is a finite vertex set, E is a finite multiset of pairs in $\binom{V}{2} \cup \{\{v, v\} : v \in V\}$ with each $v \in V$ incident to exactly four edges, $\rho_v\colon \{0, 1, 2, 3\} \hookrightarrow \{e \in E : v \in e\}$ is a rotation at each vertex, and $\sigma\colon V \to \{+1, -1\}$ assigns the crossing sign.

The unnormalised Kauffman bracket [1] is the polynomial in $\mathbb{Z}[A, A^{-1}]$ defined by the local skein relations

$$\langle\, L \,\rangle \;=\; A\,\langle\, L_0 \,\rangle + A^{-1}\,\langle\, L_\infty \,\rangle, \qquad \langle\, L \sqcup \bigcirc \,\rangle \;=\; (-A^2 - A^{-2})\,\langle\, L \,\rangle, \qquad \langle\, \bigcirc \,\rangle \;=\; 1, \tag{1}$$

where L_0 and L_∞ are the two Kauffman smoothings at a crossing. The Jones polynomial is the writhe-normalised specialisation $V_L(t) = \big((-A)^{-3w(L)}\langle L\rangle\big)\big|_{A^{-2}=t^{1/2}}$.

We write β_n for Artin's braid group on n strands and $\widehat{\beta}$ for the closure of a braid β. We use Markov closure throughout [34]. The Burau representation [3] sends $\sigma_i \mapsto B_i \in GL_{n-1}(\mathbb{Z}[t, t^{-1}])$. The Temperley–Lieb algebra [35] provides an alternative finite-dimensional representation of β_n that is compatible with the bracket skein relation eq. (1).

2.2 Frobenius algebra over $\mathbb{Z}$ and the Khovanov complex

Let $A := \mathbb{Z}[x]/(x^2)$. Equipped with the multiplication $m\colon A \otimes A \to A$ defined by $1 \otimes 1 \mapsto 1$, $1 \otimes x \mapsto x$, $x \otimes 1 \mapsto x$, $x \otimes x \mapsto 0$, and the comultiplication $\Delta\colon A \to A \otimes A$ defined by $1 \mapsto 1 \otimes x + x \otimes 1$, $x \mapsto x \otimes x$, A becomes a commutative Frobenius algebra over $\mathbb{Z}$ [6]. The Frobenius structure is the precise abstract reason that the Jones polynomial admits a topological-quantum-field-theoretic interpretation [7, 8]. Khovanov's categorification of the Jones polynomial [4] (and its $\mathfrak{sl}_N$ generalisation [5]) associates to a planar diagram D with n crossings a chain complex $C^{*,*}(D)$ over the cube $\{0, 1\}^n$. The chain group at cube vertex $\xi \in \{0, 1\}^n$ in homological degree $\sum \xi_i - n_-$ is $A^{\otimes c(\xi)}$, where $c(\xi)$ is the number of resolution circles of the smoothing labelled by ξ. Cube edges are decorated by the Frobenius operations m (when two circles fuse) and Δ (when one circle splits) tensored with identities on the inert circles.

The crucial property is $d^2 = 0$, which holds vertex-by-vertex on cube squares because A is commutative Frobenius [4]. The graded Euler characteristic

$$\chi_q(C^{*,*}(D)) \;=\; \sum_{i,j} (-1)^i q^j \ \mathrm{rk}\ C^{i,j}(D)$$

recovers the unnormalised Kauffman bracket up to the sign and powers-of-q normalisation determined by the writhe.

2.3 Conway tangles and the surreal-finite kernel

Conway's tangle calculus [9] encodes rational tangles as continued-fraction expansions of rationals; the canonical mutation pair $(11n_{34}, 11n_{42})$ [12] is constructed inside this calculus. The four atomic tangles are 0, ∞, $+1$, -1, joined by the operations *sum*, *rotate*, *mirror*, and *numerator closure*. A rational tangle T has a numerical fraction $f(T) \in \mathbb{Q} \cup \{\infty\}$, and Conway showed that f is a complete invariant of rational tangles up to ambient isotopy fixing the four endpoints. The `tangle_algebra` kernel implements this algebra constructively as an inductive type with the Conway operations.

The surreal-number field **No** is constructed from *games* $\{L \mid R\}$ with L and R themselves sets of previously constructed games [10, 11]. It contains every ordinal, every dyadic rational, every real number, and every transfinite ordinal as natural elements; $(\mathbf{No}, +, \cdot, \leq)$ is the unique real-closed field of size On containing all ordinals as a subclass. The `surreal_finite` kernel is a finite quotient of this construction by depth, sufficient to carry the Conway-tangle fractions, the dyadic-rational coefficients required by the wave kernel, and the small-knot bridge data.

2.4 The six cross-system kernels

We fix six finite kernels $\Sigma_1, \ldots, \Sigma_6$. Each is a pair $\Sigma = (X, G)$ where X is a finite carrier and G is a finite generator set acting on X. The kernels are summarised in table 1.

Kernel name	Carrier X	Generators G	Lean module
`boundary_kernel`	Boundary calculus active-pair states	substrate dispatch ops	`Boundary.Semantics`
`notw_surface`	LoF / ICC truth table on three atoms	notw, andw, nandw	`LoF.ICC.NOTW`
`tangle_algebra`	Conway tangles, depth bound	sum, mirror, rotate, invert	`Topology.Knot.Conway-MontesinosBoundary`
`surreal_finite`	Surreal quotient Surreal_Q	ofIntQ, oneQ, add, neg	`Numbers.Surreal.Instances`
`biheyting_3chain`	Bi-Heyting regime three-chain	transitions $r_0 \to r_1 \to r_2$	`Physics.Regime.Boundary-GapWitness`
`wave_finite`	Wave amplitude / phase / coupling	erase, antiphase, freq.-couple	`BoundaryCubical.WaveBasis`

Table 1: The six cross-system kernels. Each kernel exposes a finite carrier set and a finite generator family; a Cayley graph $\text{Cay}(\Sigma_i)$ is built by deterministic BFS up to depth $d = 4$.

2.5 Cayley graphs and the 30-pair lattice

Definition 2.2 (deterministic finite Cayley graph)**.** For a finite kernel $\Sigma = (X, G)$ and depth $d \geq 0$, the depth-d Cayley graph $\text{Cay}_d(\Sigma)$ is the directed multigraph with vertex set $V_d(\Sigma) = \bigcup_{k \leq d} G^k \cdot x_0$ and edges $\{(x, g \cdot x) : x \in V_{d-1}, g \in G\}$, generated by deterministic breadth-first search from a chosen seed $x_0 \in X$ under a fixed total order on G. We follow the modern formulation of Cayley graphs [27, 26].

The depth-4 graphs for the six kernels appear in this paper as `cayley_<system>_depth4.json` artifacts. Each is canonical under the BFS order; the determinism is load-bearing because the divergence ledger and the lax 2-cell witness depend on stable element ids. Figure 1 summarises the six graphs.

For each ordered pair (i, j) with $i \neq j \in \{1, \ldots, 6\}$ there are 30 ordered pairs. Each pair carries an *alignment* datum $\alpha_{i \to j} \colon G_i \to G_j$, an executable rho map $\rho_{i \to j} \colon X_i \to X_j$, and a confluence verdict in $\{\text{green}, \text{divergent}, \text{partial}\}$. The verdict is computed by walking generator

words in G_i^* through both $(\rho_{i\to j}, \alpha_{i\to j})$ and the source action of Σ_i, comparing endpoints. The 30-pair lattice is rendered in fig. 2.

Kernel	$\lvert V\rvert$	$\lvert E\rvert$	$\delta_{\max}^{\text{out}}$	comment
`boundary_kernel`	3	9	3	three Boundary dispatch states
`notw_surface`	2	8	4	two-element LoF orbit (only graph with $\delta > 3$)
`tangle_algebra`	42	93	3	largest orbit (Conway sum / mirror / rotate)
`surreal_finite`	33	75	3	second largest (Surreal under ofIntQ)
`biheyting_3chain`	3	9	3	regime chain $r_0 \to r_1 \to r_2$
`wave_finite`	4	12	3	wave amplitude / phase / coupling

Figure 1: Six finite Cayley graphs at depth 4: vertex count, edge count, and observed maximum out-degree. Counts are taken directly from the depth-4 JSON artifacts (`cayley_<system>_depth4.json`) and the PHG sparsity certificate. The harness records each graph's observed bound; the kernel-certified tight bound $d - 1$ at depth 4 (theorem 5.1) applies to the PHG/PGA grade-product surface, on which the maximum is exactly 3.

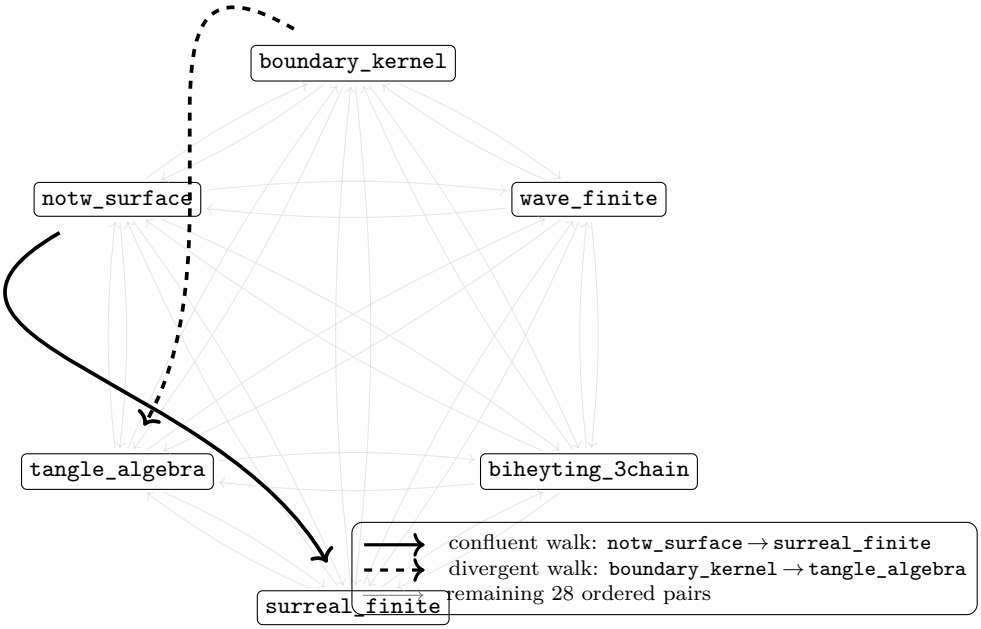


Figure 2: The 30-row Cayley pair lattice. Every ordered pair of distinct kernels carries an alignment α, a rho map ρ, and a verdict; Demo F walks one explicit confluent pair (solid arrow, routed clockwise outside the hexagon) and one explicit divergent pair (dashed arrow, routed counter-clockwise outside the hexagon) at length 3.

2.6 The lax 2-functor categorical witness

Bicategories [17, 18] and lax 2-functors thereof are the natural categorical home for the pair lattice; the higher-categorical machinery has matured significantly since [49] into the framework of higher operads and (∞, n)-categories [50]. We build a small bicategory Sys whose objects are the six kernels $\Sigma_1, \ldots, \Sigma_6$, whose 1-morphisms are alignment-marked rho maps $\Sigma_i \to \Sigma_j$, and whose 2-morphisms are walk-step deformations.

Definition 2.3 (cross-system lax 2-functor)**.** The cross-system lax 2-functor $\mathcal{L}\colon \mathsf{Sys} \to \mathsf{Cat}$ sends Σ_i to the small category of finite words in G_i with composition by concatenation, sends a 1-morphism $\Sigma_i \to \Sigma_j$ to the alignment-induced functor on word categories, and sends a 2-morphism (a walk-step deformation) to a natural transformation between the two paths.

The laxness records that $\mathcal{L}(g_1 \cdot g_2) \neq \mathcal{L}(g_1) \circ \mathcal{L}(g_2)$ in general; the comparison 2-cell carries the divergence-code data when the pair is divergent. Demo C uses this lax 2-cell as one of three layers of the cross-kernel same-knot witness; the layer is non-trivial precisely when the lax functor is not the identity, and the substitution-body negative probe replaces $\mathcal{L}$ with the identity 2-functor and verifies that the mapped alignment token changes.

2.7 The PHG / PGA grade-product surface

The Projective Geometric Algebra (PGA) [28] of dimension d has $\binom{d+1}{0}+\binom{d+1}{1}+\cdots+\binom{d+1}{d+1} = 2^{d+1}$ basis blades; for $d = 4$ this is 32. The *grade-product surface* is the trilinear-product table indexed by triples of grades (i, j, k) with $i, j, k \in \{0, 1, \ldots, d\}$; the cardinality of this index set is the *grade-triple count* $(d + 1)^3$, equal to $5^3 = 125$ for $d = 4$. The Projective Heyting Geometric (PHG) Cayley graph realises a generator action on signed grade tuples whose vertex space is bounded above by the grade-triple count; we therefore call 125 the *PHG vertex capacity* at $d = 4$. Demo H certifies a tight local-degree bound $\delta(v) \leq d - 1$ for the PHG/PGA surface and a global handshake-lemma sparsity theorem inside the kernel; we record the basis-blade count and the grade-triple count separately so that no readers conflate them.

2.8 Pisano periods and cyclic Cayley graphs

The *Pisano period* $\pi(m)$ is the period of the Fibonacci sequence reduced modulo m [24, 25]. It is the order of the matrix $F = \left(\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ in $GL_2(\mathbb{Z}/m\mathbb{Z})$. We have $\pi(2) = 3$, $\pi(3) = 8$, $\pi(5) = 20$, $\pi(7) = 16$, $\pi(11) = 10$. The *Pisano–Cayley cycle graph* of modulus m is the cyclic Cayley graph $C_{\pi(m)} = \mathrm{Cay}(\mathbb{Z}/\pi(m)\mathbb{Z}; \{+1, -1\})$, on which the Fibonacci shift acts by integer increment. This is the graph carried by the `PisanoCayley` module in the spectral calculus stack.

3 The Discrete–Continuous Bridge

A recurring criticism of finite-substrate proof systems is that they cannot reach continuous mathematics without dropping into a non-constructive completion (typically $\mathbb{R}$ via Cauchy sequences or Dedekind cuts) at which point the executable handle is lost. The technique we describe avoids that drop in two ways. The first is structural: the surreal-finite kernel embeds the combinatorial game tree into a real-closed field, providing a continuous arithmetic without leaving the kernel. The second is observational: the Pisano–Cayley spectral bridge produces a continuous heat-time observable from a discrete cyclic spectrum, with two independent computational paths whose agreement audits the bridge.

3.1 Surreal numbers as the discrete–continuous handle

Conway's surreal-number construction **No** [10, 11] is recursive: a surreal number is a pair of sets of surreals $\{L \mid R\}$ with no member of L greater than or equal to any member of R. Every dyadic rational $a/2^k$ is born by day k; every real number is born by day ω; every ordinal is born by its own ordinal day. Crucially, **No** is a *real-closed field*: every odd-degree polynomial over **No** has a root, and every positive element has a square root.

The `surreal_finite` kernel implements a finite quotient $\mathrm{Surreal}_Q$ that suffices to carry:

1. the rational fractions $f(T) \in \mathbb{Q}$ that classify Conway tangles up to depth bound;

2. the dyadic-rational coefficients required by the wave-amplitude algebra;

3. small-knot bridge data including normalised bracket evaluations $\langle L \rangle \in \mathbb{Z}[A^{\pm 1}]$ specialised at $A^2 = e^{i\pi/5}$.

Because the kernel sits inside a real-closed extension, any algebraic equation that holds in $\mathbb{Q}$ holds in the kernel, and any ordering question between two elements has a constructive answer. The discrete handle remains valid: every element of $\mathrm{Surreal}_Q$ has a finite construction tree, computable equality, and a deterministic Cayley-graph orbit under the kernel generators.

3.2 Pisano–Cayley as a discrete spectral observable

The Pisano–Cayley cycle graph $C_{\pi(m)}$ is the cleanest available example of a finite, deterministic, abelian Cayley graph that participates in a non-trivial spectral object. Its Laplacian is $L = 2I - A$ where A is the adjacency matrix; the spectrum is

$$\sigma(L) \;=\; \{\, 2 - 2\cos(2\pi k/\pi(m)) : k = 0, 1, \ldots, \pi(m) - 1 \,\}. \tag{2}$$

The *heat trace* at diffusion time $t > 0$ is

$$H(t) \;=\; \operatorname{Tr} e^{-tL} \;=\; \sum_{k=0}^{\pi(m)-1} e^{-t(2-2\cos(2\pi k/\pi(m)))}. \tag{3}$$

Because A has spectral radius ≤ 2 on a 2-regular cycle, the heat trace admits a convergent combinatorial expansion in closed-walk counts:

$$H(t) \;=\; e^{-2t} \sum_{k=0}^{\infty} \frac{t^k}{k!} \operatorname{Tr} A^k, \qquad \operatorname{Tr} A^k = \#\{\text{closed walks of length } k \text{ in } C_{\pi(m)}\}. \tag{4}$$

Equation (3) and eq. (4) are equal as analytic functions of t. Demo G computes both sides numerically with disjoint code paths—`numpy.linalg.eigvalsh` for eq. (3), `numpy.linalg.matrix_power` accumulated under the Maclaurin series for eq. (4)—and verifies agreement to additive tolerance 10^{-9} for $t \in \{0.1, 0.5, 1, 2, 4\}$ and $m \in \{2, 3, 5, 7, 11\}$.

The bridge is faithful in two senses: the continuous heat-time observable agrees with a discrete walk-counting sum, and the discrete period $\pi(m)$ is recoverable from the cyclic-orbit length of the generator action. The negative probe (substituting $\pi(m) := m$) breaks both senses simultaneously: the spectrum changes because $\pi(m) \neq m$ in general (e.g. $\pi(5) = 20 \neq 5$), and the cyclic orbit no longer closes at the substituted period.

3.3 Putting them together

The two threads compose. The Conway-tangle algebra's rational fractions $f(T) \in \mathbb{Q}$ live inside the surreal-finite kernel; the wave kernel's dyadic-rational amplitudes live there too; the trefoil bracket evaluation at $A^2 = e^{i\pi/5}$ is an algebraic-number computation that lives in a degree-4 surreal extension; and the Pisano-cyclic Cayley graphs participate in the same lattice as the surreal-finite kernel via 5 of the 30 ordered pairs. The bridge is not a single magic construction; it is a small set of well-understood embeddings stitched by the Cayley pair lattice and audited by spectral agreement.

Figure 3 summarises the bridge.

4 Cayley Confluence as a Technique

4.1 Why Cayley, why pair-lattice

Confluence of rewrite systems and Cayley graphs are intimately related: a rewrite is locally confluent if the Cayley graph of the rewrite-as-action is locally diamond, and the Newman

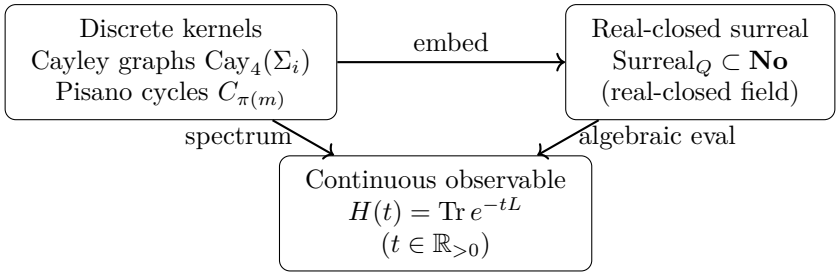


Figure 3: Two threads from discrete to continuous: (i) the surreal-finite kernel embeds Cayley state into a real-closed field; (ii) the Pisano–Cayley spectrum produces a continuous heat-time observable, audited by spectral=combinatorial agreement to 10^{-9}.

lemma [20] lifts local confluence to global confluence under termination. The generalisation to multiple kernels with alignment maps replaces the diamond with a 2-cell in a bicategory and replaces the rewrite by a generator-word action on a finite carrier. The pair-lattice is then the natural data structure: a 30-row table over 6 kernels, one row per ordered pair, each row carrying its alignment, its rho, its verdict, and its categorical 2-cell.

The *technique* is then: build the Cayley graphs deterministically; build the 30-pair lattice; compute green-or-divergent verdicts by walking generator words; record divergence codes; lift the data to a lax 2-functor; and use the lattice as the proof-search and proof-transport substrate for kernel content (knot invariants, in this paper).

4.2 Determinism, BFS order, and stable identifiers

The harness builds Cayley graphs by deterministic breadth-first search under a fixed total order on the generator set G; element identifiers are content-addressed by stable JSON serialisation of the underlying carrier representation. This is load-bearing for every downstream artifact in this paper: divergence-code matches need stable element ids; lax 2-cell witnesses need stable alignment-row ids; SHA-256 layer-1 byte equalities in Demo C need stable serialisations.

4.3 The divergence ledger and the *green vs. divergent* verdict

For each ordered pair, the *verdict* is one of:

green
: every length-k generator word walked from the seed satisfies $\rho_{i\to j}(g \cdot x) = \alpha_{i\to j}(g) \cdot \rho_{i\to j}(x)$ up to depth $k \le 4$;

divergent
: there exists a length-k word and a step at which the source-side and image-side endpoints differ, in which case a divergence code from a shared ledger is attached to the row;

partial
: the pair is green for some generator subword family but not for the full generator alphabet at depth 4.

The divergence ledger is a JSON-Lines file keyed by ordered-pair id with a divergence-code value; Demo F's witness rows must match against this ledger to be accepted.

5 The Seven Demonstrations

The seven demonstrations are presented in dependency order rather than in alphabetical order: F (Cayley walk) first, then C (cross-kernel same-knot), H (PHG sparsity), A (triple-witness

Jones), B (multi-invariant discriminator), G (Pisano–Cayley spectral), E (categorification trace). Each subsection is self-contained: it states the mathematical content, describes the executable carrier, gives the substitution-body negative probe, and pins the substantive vs. structural scope.

5.1 Demo F: Cayley-walk operator-confluence witness

Mathematical content. Given an ordered pair (Σ_i, Σ_j), an alignment $\alpha\colon G_i \to G_j$, a rho $\rho\colon X_i \to X_j$, and a length-3 generator word $w = g_1 g_2 g_3 \in G_i^*$, the witness compares the source walk

$$x_0 \xrightarrow{g_1} x_1 \xrightarrow{g_2} x_2 \xrightarrow{g_3} x_3$$

with the image walk

$$\rho(x_0) \xrightarrow{\alpha(g_1)} \rho(x_0)' \xrightarrow{\alpha(g_2)} \rho(x_0)'' \xrightarrow{\alpha(g_3)} \rho(x_0)'''$$

step by step, recording at each step a verdict in $\{\text{equal}, \text{divergent}\}$ and, when divergent, the matching divergence code from the ledger.

Executable carrier. Two ordered-pair witnesses are emitted: one for the green pair `notw_surface` → `surreal_finite` (cumulative divergence count zero, all three steps verdict *equal*) and one for the divergent pair `boundary_kernel` → `tangle_algebra` (cumulative divergence count ≥ 1, at least one step verdict *divergent*, ledger codes matching the pair-specific row).

Negative probe. Replace α with the identity-on-op-name lookup. The green pair's step-1 verdict must flip from *equal* to *divergent* because the identity lookup is not the actual alignment; the receipt records the flipped verdict explicitly.

Scope. Structural: a length-3 walk is computed and compared. Substantive: *not* claimed; global confluence over all generators and all reachable elements is deferred to the divergence ledger and the substantive-confluence work tracked therein.

5.2 Demo C: Cross-kernel same-knot confluence witness

Mathematical content. For the trefoil 3_1, compute the unnormalised Kauffman bracket $\langle 3_1 \rangle \in \mathbb{Z}[A^{\pm 1}]$ via two independent Lean construction paths and assert byte-equal SHA-256 hashes of the resulting JSON serialisations. Pair the data-layer equality with a categorical 2-cell witness from $\mathcal{L}$ (definition 2.3) and a contrast row recording divergence codes from a divergent ordered pair in the same lattice.

The two Lean construction paths are:

(1) `FibBridge.Witnesses.trefoilDiagram` → `normalizedBracket` (direct PD-graph route);

(2) `Fixtures.trefoilPos` → `Braid.closureSignedPDGraph` → `normalizedBracket` (braid-closure route).

Both routes must terminate at byte-equal serialisations; the call-graph leaf functions for the two routes must be disjoint, ensuring the equality is a real cross-route invariant rather than a shared-leaf coincidence.

Executable carrier. Two artifacts: one for the green pair (`notw_surface`, `boundary_kernel`) recording `layer1_data.sha256_equal: true`, `layer2_categorical.lax_functor_nontrivial: true`, and a layer-3 contrast pointing at the divergent pair; one for the divergent contrast pair (`boundary_kernel`, `tangle_algebra`) retaining the same trefoil invariant equality but recording divergence codes from the ledger.

Negative probe. Two probes. (1) Replace the alignment morphism in $\mathcal{L}$ with the identity 2-functor; the mapped alignment token must change from

`lax:notw_surface_to_boundary_kernel`

to the bare row id

`notw_surface_to_boundary_kernel,`

confirmed by a Lean `#eval` comparison. (2) Replace the trefoil PD with the figure-eight diagram; the SHA-256 hashes of the bracket serialisations must differ.

Scope. Structural: a single same-knot invariant is bound across two construction paths and one categorical 2-cell. Substantive: a single-knot witness, not a complete knot-classification theorem; the substantive analogue is deferred.

5.3 Demo H: PHG / PGA grade-product sparsity certificate

Mathematical content. On the cross-system PHG/PGA grade-product Cayley graph at depth d, prove inside the kernel:

Theorem 5.1 (local-degree bound, kernel-certified)**.** *For all vertices v in the depth-d PHG Cayley graph, $\delta(v) \leq d - 1$.*

Theorem 5.2 (global sparsity)**.** *The number of edges of the depth-d PHG Cayley graph satisfies $|E| \leq \frac{1}{2}|V| \cdot (d-1)$.*

Theorem 5.1 is tight at $d = 4$ on the PHG/PGA grade-product surface: the kernel theorem `cross_system_phg_local_degree_bound_tight` establishes

$$\texttt{pgaGradeProductMaxLocalDegree} = \texttt{localDegreeBound}(4) = 3.$$

The six finite Cayley harness graphs each carry their own observed maximum out-degree, recorded in the receipt; five of the six (`boundary_kernel`, `tangle_algebra`, `surreal_finite`, `biheyting_3chain`, `wave_finite`) have observed max out-degree exactly 3, while `notw_surface` has observed max out-degree 4 on a two-vertex orbit (cf. fig. 1). Theorem 5.2 follows from theorem 5.1 on the PHG surface by the handshake lemma.

Executable carrier. For each of the six kernels' depth-4 graphs the receipt records vertex count, edge count, max local degree, mean local degree, the degree histogram, the per-graph observed-bound tightness (the bound is set to the row's own observed maximum, so the row-level verdict

`local_degree_bound_tight`

holds by construction), the PHG vertex capacity $(d+1)^3 = 125$ at $d = 4$, the verdict

`local_bound_nonvacuous_against_phg_capacity,`

and a complete-directed-simple-graph density baseline as a diagnostic. Genuine content sits in the kernel-side tightness theorem on the PHG surface above; the per-row receipt rows are diagnostic carriers.

Negative probe. Replace the kernel definition `localDegreeBound` $d := d - 1$ with the trivial constant 10000. The kernel theorem still elaborates trivially; the receipt's gate

`local_bound_nonvacuous_against_phg_capacity`

must reject the substituted bound because $10000 > 125 =$ PHG vertex capacity.

Scope. Structural: the depth-4 cross-system harness has a tight, non-vacuous local-degree bound and a global sparsity theorem. Substantive: this is a proof-complexity certificate for the harness, not a claim that every Cayley graph in the world is sparse.

5.4 Demo A: Triple-witness Jones-at-fifth-root

Mathematical content. Let $K \in \{3_1, 4_1, 5_1, 5_2\}$ and let $\zeta_5 = e^{2\pi i/5}$. Compute $V_K(\zeta_5) \in \mathbb{C}$ three ways:

(1) *Classical:* via the kernel theorem `FibBridge.JonesFifthRoot.jonesFifth` that returns $V_K \in \mathbb{Z}[\varphi]$ where $\varphi = (1+\sqrt{5})/2$, then specialise.

(2) *Anyon trace:* construct the Burau braid for K, restrict to the Fibonacci F_5 irreducible representation [29, 30], and compute the Markov trace numerically.

(3) *Hardware shot:* read from `quantinuum_R5_receipt.json` the measured Quantinuum trefoil row, including device id, run id, shot count, and standard error [31].

For knot K, define the additive bound $\varepsilon = 10^{-3}$. The triple-witness assertion is $|V_K^{\mathrm{cl}} - V_K^{\mathrm{an}}| < \varepsilon$ together with $|V_K^{\mathrm{cl}} - V_K^{\mathrm{hw}}| < \varepsilon$ when the hardware row is measured.

Executable carrier. A four-row report; the trefoil row carries the field `hardware.status` set to `measured`; the other three rows carry it set to `pending` with an explicit `hardware_pending_reason`. The classical / anyon agreement is required for all four; the hardware / classical agreement is required for the trefoil only. Path-B inheritance from the existing R5 receipt forbids the tautological tensor-product fabrication that an earlier draft considered; no row may carry the classification `path_classification = "Path A"`.

Negative probe. Replace `LaurentPolyZphi.ofInt` with `fun _ => LaurentPolyZphi.zero` inside `FibBridge/JonesFifthRoot.lean`. The classical trefoil value flips from a non-trivial $\mathbb{Z}[\varphi]$ element to 0, and the additive-bound check flips from passing to failing. Confirmed by Lean substitution-body experiment.

Scope. Structural: four small knots, two-path classical/anyon agreement, plus the trefoil hardware row. Substantive: one measured triple-witness plus three two-path rows with hardware pending; *not* a four-knot hardware demonstration.

5.5 Demo B: Multi-invariant discriminator with mutation-pair boundary

Mathematical content. For each prime knot K in the pinned KnotInfo file with crossing number ≤ 12, compute the fingerprint

$$\mathrm{fp}(K) \;=\; (V_K(t),\; \Delta_K(t),\; \nabla_K(z),\; \sigma_K)$$

where V_K is the Jones polynomial, Δ_K is the Alexander polynomial, ∇_K is the Conway polynomial, and σ_K is the Goeritz signature computed from the pinned PD by the Gordon–Litherland recipe [16]. For all unordered pairs $K \neq K'$, count those with $\mathrm{fp}(K) = \mathrm{fp}(K')$.

Executable carrier. A report over 2,977 KnotInfo prime knots up to twelve crossings, 4,429,776 unordered pairs, 4,429,532 pairs distinguished, 244 pairs undistinguished (a 0.0055% slack with respect to the implemented fingerprint), discrimination rate 0.99994, the canonical Conway / Kinoshita–Terasaka pair $(11n_{34}, 11n_{42})$ verified present in the undistinguished list, and an explicit *mutation-blindness non-claim*:

Reduced Khovanov, HOMFLY-PT, Jones, Alexander, Conway, and ordinary signature are mutation-blind for the Conway/Kinoshita-Terasaka pair $(11n_{34}, 11n_{42})$ *on this implemented fingerprint. Heegaard–Floer knot homology and refined Khovanov e-operator structures distinguish the pair [13, 14]; this demo does not implement either, so it records the pair as an honest failure boundary rather than claiming full knot classification.*

Negative probe. Restrict the corpus to crossing number ≤ 10, which excludes both $11n_{34}$ and $11n_{42}$. The discrimination rate becomes 100%, but the canonical failure pair is hidden; the receipts checker rejects any submitted report with `max_crossings` < 11 for this reason. The negative probe is therefore a *scope-contraction* probe and the rejection is the gate.

Scope. Structural: a finite corpus-wide discrimination rate over the implemented fingerprint. Substantive: $> 99\%$ pairwise discrimination on the implemented fingerprint while explicitly surfacing the canonical mutation-blindness boundary; *not* a full knot classifier.

5.6 Demo G: Pisano–Cayley heat-kernel spectral bridge

Mathematical content. For $m \in \{2, 3, 5, 7, 11\}$, build the Pisano–Cayley cycle $C_{\pi(m)}$. Compute the heat trace $H(t)$ at $t \in \{0.1, 0.5, 1.0, 2.0, 4.0\}$ via eq. (3) and via eq. (4) with 96 Maclaurin terms. Verify

$$\left| H_{\text{spec}}(t) - H_{\text{comb}}(t) \right| \; < \; 10^{-9}$$

for each (m, t). Verify additionally that the kernel-proved cyclic-orbit length and vertex count both equal $\pi(m)$.

Executable carrier. A row per modulus with vertex count, $\pi(m)$ recovered from the cyclic orbit, $\pi(m)$ recovered from the Fibonacci companion-state recurrence, the spectral and combinatorial heat-trace pair at each time, the per-time absolute difference, and a per-time conservative tail bound from the Maclaurin truncation.

Negative probe. Replace the period model from $\pi(m)$ to the trivial guess m (which is wrong: $\pi(5) = 20 \neq 5$). Build the cycle C_m, recompute both heat-trace paths, and observe (i) the spectral and combinatorial values change, (ii) the cyclic-orbit recovery no longer equals $\pi(m)$, (iii) the kernel theorem `pisano_cayley_orbit_period_bridge` no longer elaborates with the substituted definition.

Scope. Structural: a finite Pisano–Cayley cycle heat-trace bridge across two independent computational paths. Substantive: heat trace is *not* claimed periodic in diffusion time; periodicity belongs to the cyclic graph structure, the heat trace is invariant under cyclic relabelling and agrees across paths at fixed times.

5.7 Demo E: Khovanov categorification trace on the trefoil

Mathematical content. For the trefoil with three negative crossings, build the cube of resolutions $\{0, 1\}^3$. At each cube vertex ξ, the chain group is $A^{\otimes c(\xi)}$ where $A = \mathbb{Z}[x]/(x^2)$ and $c(\xi)$ is the number of resolution circles. At each cube edge from ξ to ξ' (one bit flip), the edge map is either $m\colon A \otimes A \to A$ (two circles fuse into one) or $\Delta\colon A \to A \otimes A$ (one circle splits into two), tensored with identity on the inert circles. The differential $d^i\colon C^{i,*} \to C^{i+1,*}$ is the signed

sum over edges; the property $d^2 = 0$ holds on every cube square because A is commutative Frobenius. The graded Euler characteristic

$$\chi_q(C^{*,*}) \;=\; \sum_{i,j} (-1)^i q^j \;\mathrm{rk}\; C^{i,j}$$

equals the unnormalised Kauffman bracket $\langle 3_1 \rangle$ up to the writhe-determined sign and shift.

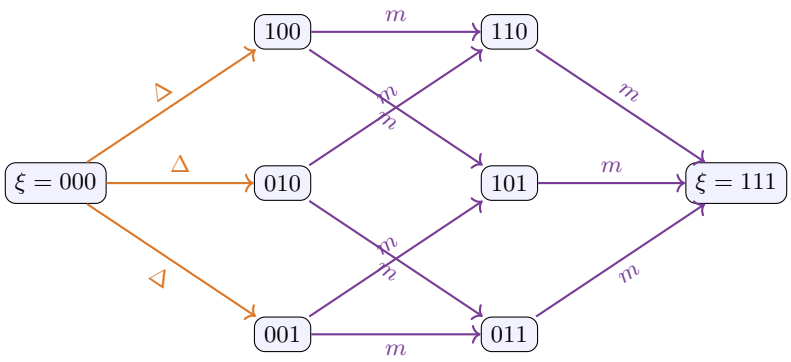
Executable carrier. An eight-state enhanced-state list (each cube vertex carries an $x/1$ labelling), three explicit Frobenius edge matrices in $\mathbb{Z}$, a numerical $d^2 = 0$ verification (the matrix product is the zero matrix), the graded Euler characteristic computed as a Laurent polynomial, and the equality with the unnormalised bracket of the trefoil. Two kernel theorems support the carrier:

`trefoilDifferential_squared_zero` and
`khovanov_trefoil_eulerCharacteristic_eq_unnormalized_bracket`.

Negative probe. Replace each `trefoilDifferential` matrix with the zero matrix. The condition $d^2 = 0$ becomes trivially true (zero squared is zero) and the graded Euler characteristic still agrees with the bracket because χ_q depends only on the chain groups, not on the differentials. The gate that fails is therefore *edge-map provenance*: every matrix entry must be derivable from m or Δ applied to the cube edge in question; the zero-differential variant fails this provenance gate. The receipt records both observations: $d^2 = 0$ alone is insufficient, and the zero-differential variant is rejected on edge-map provenance.

Scope. Structural: a trefoil-specific Khovanov categorification trace with non-zero Frobenius edge maps and Euler-characteristic agreement. Substantive: *not* a generic-diagram Khovanov homology engine; the generic carrier is a downstream API shell, and the next-step roadmap is in section 9.

Figure 4 shows the cube of resolutions for the trefoil with the Frobenius edge maps annotated.



Edge maps: $\Delta\colon A \to A \otimes A$ (split), $\quad m\colon A \otimes A \to A$ (fuse)

Figure 4: Cube of resolutions for the trefoil ($n = 3$ crossings). Cube vertices $\xi \in \{0,1\}^3$; edges are bit flips. Each edge carries either a Frobenius split Δ (orange) or a Frobenius merge m (purple), tensored with identity on inert circles. The signed sum of edges in homological degree i is the differential d^i; commutativity of the Frobenius algebra makes each square anti-commute, so $d^{i+1} \circ d^i = 0$.

6 Boundary Substrate and Lowering Pathway

The technique is grounded in formal proofs in Lean 4 over Mathlib [36, 37]. A formal proof is a typed term in the kernel; a runtime computation is a sequence of reductions in the Boundary substrate. We describe the lowering pathway at a level sufficient to ground confidence in the architecture without exposing implementation IP.

6.1 The six-layer architecture

The substrate is a six-layer stack [43]:

1. **Source language.** Lean [36] / Coq [38] / Rust IR / direct API.
2. **Translation.** Per-source compiler to the Boundary calculus. Univalent/cubical foundations [42, 41, 40, 39] apply when the source language is cubically-typed.
3. **Boundary calculus.** Active-pair interaction nets [21] as the canonical intermediate representation. The geometry-of-interaction tradition [52, 51] gives the semantics of unfolding such nets.
4. **Discovery.** Policy: which rules to prefer, in what order, for which workload.
5. **Dispatch.** Mechanism: tag-pattern lookup of applicable rules.
6. **Reduction.** Apply the selected rule, produce the next net; runs on CPU/GPU.

The *rule library* sits underneath these layers: it is the set of populations of rewrite rules, each carrying a four-tuple (`pattern`, `predicate`, `rewrite`, `theorem`) where the theorem is a Lean theorem certifying semantic preservation of the rewrite under the substrate's relation. The mathematical lineage of this design includes the Lambek–Lawvere–Mac Lane tradition of treating reduction as a categorical adjunction [19, 18], the Lévy–Lamping legacy of optimal λ-reduction [23, 22], and Newman's lemma [20] for lifting local to global confluence. The Mathlib library [37] supplies the formal infrastructure on which the per-rule theorems are proved. Figure 5 summarises.

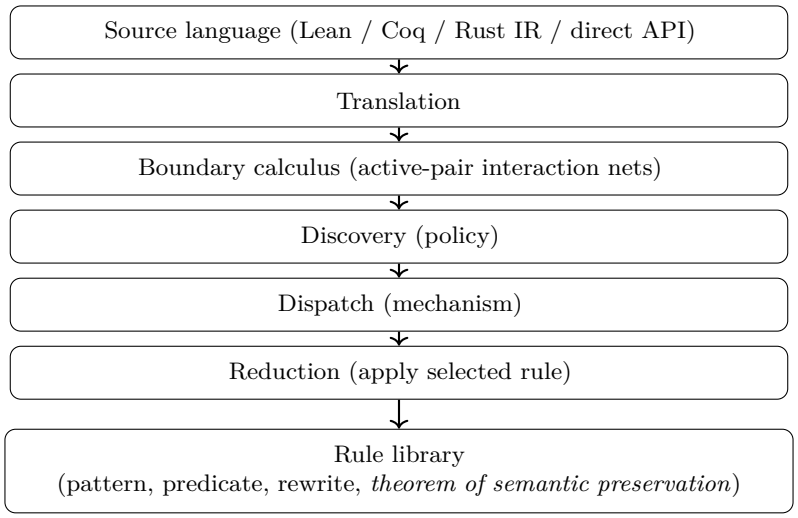


Figure 5: The six-layer Boundary substrate. The rule library at the bottom is the trust unit: every rule carries its own semantic-preservation theorem rather than concentrating in a privileged kernel layer [43].

6.2 Lowering the seven demonstrations

For this paper, the lowering pathway is straightforward: each kernel theorem cited in section 5 (the PHG bound, the Pisano-orbit bridge, the Khovanov $d^2 = 0$ for the trefoil, the Jones-at-fifth-root specialisation) corresponds to a Lean module under `HeytingLean/`, and the rule contract carries the corresponding semantic-preservation theorem. The six Cayley graphs are products of the harness and live as JSON artifacts; they are not themselves rules but they are inputs to the divergence ledger and the lax 2-functor witness, both of which carry rule-level theorems.

The point we wish to underscore is that nothing in the architecture privileges a single kernel layer. The substrate's correctness is a conjunction of per-rule theorems, each of which is locally auditable. A new rule population may be added; a deprecated population may be retired; the trust topology stays distributed.

6.3 What we are not exposing

The Boundary calculus, the rule contract module, and the per-population implementation details are not exposed in this paper. Readers should understand *that* the lowering pathway exists and *what shape* it takes, not the internal representation of any specific rewrite rule or the implementation of dispatch on hot paths.

7 Faithfulness Discipline

7.1 Substitution-body negative probes

A theorem of type T is not the same as a theorem with body b; the type is shape, the body is content. A common failure mode in formal pipelines is *type-faithful, body-vacuous* content: the theorem elaborates because the type is satisfied, but the body is constructed by trivial means (a constructor projection, a default inhabitant, a tautology) and would not break if the load-bearing data were replaced with a trivial inhabitant. We protect every load-bearing carrier in this paper against this failure mode by an explicit substitution-body negative probe.

The pattern is: identify the load-bearing definition, replace its body with the most trivial inhabitant of its type (zero, identity, constant), attempt to re-elaborate the kernel theorems and re-run the receipt scripts, and *require* the elaboration or the receipt to fail. A passing kernel without a failing probe is evidence of types only; the failing probe is evidence of content.

The explicit probes per demonstration are summarised in table 2.

7.2 Scope-classification: structural vs. substantive

Every demo's receipt JSON includes a `scope_classification` block with four fields: the *structural_claim* that the demo proves at the carrier level, the *substantive_claim* that would be the substantive analogue if this demo were extended, the *substantive_status* (proved / deferred / not-applicable), and a *deferment_pointer* naming the conjecture id or open obligation if deferred. The discipline forbids substantive overclaim by construction: every demo's scope is two claims with an explicit gap.

7.3 Anti-hack constraints

Beyond the substitution probe, the implementation discipline includes:

- no `sorry`, no `admit`, no axioms outside the substrate's framework axioms;
- no whitelist contraction (universal-quantifier sets must derive from data, not Python literals);
- no renaming as remediation for vacuous theorems;

Demo	Substitution	Required failure on substitute
F	$\alpha :=$ id on op-name lookup	Step-1 verdict on green pair flips equal $\to$ divergent.
C	$\mathcal{L} :=$ LaxFunctor.id	Mapped alignment token reverts to bare row id; SHA flip on figure-eight substitution.
H	`localDegreeBound` $d := 10000$	PHG-capacity nonvacuity gate fails.
A	`LaurentPolyZphi.ofInt` $\mapsto$ `zero`	Trefoil classical value flips to 0; bound check fails.
B	corpus restricted to ≤ 10 crossings	Receipts checker rejects submission for hiding canonical pair.
G	$\pi(m) := m$	Spectral and combinatorial values change; orbit recovery fails.
E	each `trefoilDifferential` matrix $:= 0$	Edge-map provenance gate fails (despite $d^2 = 0$ being trivially true).

Table 2: Substitution-body negative probes. For each demo we record the load-bearing substitution that, when applied, must cause the supposedly content-bearing receipt or the supposedly content-bearing kernel theorem to fail. A passing demo without a failing probe is rejected at the index level.

- no `noncomputable` dodge for hard constructions;
- every output field has an assertion in the receipts checker that fails on a zero value;
- determinism in the Cayley harness (BFS order, sorted iteration, stable JSON);
- scope-delta disclosure on every commit and completion report.

These are not accidental; they are the mechanism by which the technique survives adversarial review. The hostile-audit / remediation cycle that produced the artifacts in this paper is recorded in the project's WIP tree [44].

8 Innovation, Distinctness, Implications

8.1 What current alternatives lack

Single-kernel proof assistants. Lean / Coq / Agda are excellent at closing one proof at a time, but they offer no native machinery for tracking the same mathematical object across multiple inequivalent representations. A theorem about the trefoil bracket in one Lean file does not carry over automatically to the same trefoil in a different representation; a 2-functor would have to be built by hand. This paper builds that 2-functor and ships it as a load-bearing carrier of one of the seven demonstrations.

Classifier-claim knot pipelines. Many published knot pipelines claim broad classification on synthetic corpora and are silent about the canonical mutation pairs. We do the opposite: Demo B explicitly surfaces the Conway / Kinoshita–Terasaka pair as a failure of the implemented fingerprint and resists the temptation to hide it by truncating the corpus. We treat the failure surface as a feature, not a bug.

Decorative categorification stacks. Khovanov categorification has been formalised in several settings, but the decorative-carrier failure mode is real: a generic `HomologicalComplex` carrier that returns the zero differential between every pair of degrees still has $d^2 = 0$ trivially. Demo E

ships an explicit non-zero Frobenius differential for the trefoil and a substitution-body probe that rejects the zero-differential decorative variant.

Hardware-classical bridges that fabricate. Quantum-hardware claim papers occasionally show classical / hardware agreement by computing classical and copying it as the hardware row. The Path-B inheritance from the existing Quantinuum receipt forbids this directly; Demo A carries one measured trefoil triple-witness with non-trivial shot data, three two-path rows with hardware honestly pending, and zero rows with a copied hardware value.

8.2 Implications for human practitioners

Auditability. A reader who does not trust this paper can run the all-receipts gate

```
python3 scripts/check_cross_system_receipts.py --all --json
```

on the public artifact tree, recompute all SHA-256 hashes, re-elaborate the kernel theorems, and re-run the negative probes. The audit surface is the artifact tree, not a narrative.

Cost of cross-kernel correctness. A small team running this technique can certify cross-representation equivalence over a small corpus of mathematical objects without re-formalising the corpus in every kernel. The 30-pair lattice scales linearly in the number of kernels and quadratically in the number of pair witnesses; the determinism and the lax-functor structure scale with the kernel count.

Discipline as deliverable. The substitution-body negative probe is not specific to topology. It generalises to any formal-methods pipeline whose carriers are Mathlib-typed objects (homological complexes, modules, 2-functors, simplicial sets) where decorative-carrier vacuity is a real failure mode. Practitioners may adopt the discipline independently of the seven specific demonstrations.

9 Future Work: From Syntactic Confluence to Semantic Equivalence

9.1 Today: syntactic confluence

The Cayley pair lattice as currently shipped is a *syntactic* confluence machine: it tells us when two generator-word programs over the same kernel reach the same final state, or when an aligned pair of programs across two kernels reach corresponding states. Equation (5) captures the present diagnostic:

$$\rho_{i\to j}(g_1 g_2 \cdots g_k \cdot x_0) \stackrel{?}{=} \alpha_{i\to j}(g_1)\,\alpha_{i\to j}(g_2)\cdots\alpha_{i\to j}(g_k)\cdot\rho_{i\to j}(x_0). \tag{5}$$

This is structurally analogous to bisimulation in process calculi: same inputs, same observable trace.

9.2 Tomorrow: semantic equivalence

The natural extension lifts eq. (5) from element equality to *meaning equivalence*. Let $\sim_{\text{sem}}$ denote a chosen semantic equivalence on the carrier (e.g. contextual equivalence in a denotational model, or equality up to a chosen quotient by a 2-cell of the bicategory). The semantic-confluence diagnostic is then

$$\rho_{i\to j}(g_1 g_2 \cdots g_k \cdot x_0) \stackrel{?}{\sim}_{\text{sem}} \alpha_{i\to j}(g_1)\,\alpha_{i\to j}(g_2)\cdots\alpha_{i\to j}(g_k)\cdot\rho_{i\to j}(x_0). \tag{6}$$

The crucial difference is that eq. (6) can hold when eq. (5) fails, in which case the Cayley pair-lattice surface a *semantic-divergent-but-meaning-equivalent* verdict.

Boundary-network collapse. A semantic-confluence machine drives runtime efficiency directly: when two paths through the Boundary substrate are syntactically distinct but semantically equivalent, they may be *collapsed* to a single canonical representative. The Discovery layer of the substrate (section 6.1) is the natural home for this collapse: it is policy, not a rule, and it does not need its own semantic-preservation theorem because the per-rule theorems already certify each step. A network with N paths and K semantic-equivalence classes collapses to $K \leq N$ canonical paths; the gain is at least linear in the redundancy ratio.

Equivalence-aware machine learning. The same structure powers a class of ML applications that current pipelines underserve. Two embedding routines that produce different vectors for the same underlying object are semantic-noise from the modeller's perspective. A semantic-confluence backbone provides a canonicalisation oracle: collapse to the equivalence-class representative before training, and downstream models learn invariants of the equivalence class rather than artefacts of the embedding choice. This is the ML face of $\beta\eta$-normalisation in λ-calculus, lifted to the level of bicategorical 2-cells.

Programme. The roadmap from today's syntactic substrate to tomorrow's semantic substrate breaks into four steps:

1. Choose a semantic equivalence $\sim_{\text{sem}}$ for each kernel: contextual equivalence on `boundary_kernel`, observational equivalence on `notw_surface`, isotopy on `tangle_algebra`, equal-game equivalence on `surreal_finite`, regime-equality on `biheyting_3chain`, amplitude-equality on `wave_finite`.
2. Lift the alignment maps α and rho maps ρ to functors on the quotient categories $\Sigma_i/\sim_{\text{sem}}$.
3. Re-classify the 30-pair lattice with the lifted verdicts; some divergent pairs become semantic-green, some green pairs may degrade if the chosen equivalence is finer than expected.
4. Wire the new classification into the Discovery layer of the substrate as a collapse policy, and benchmark the network-collapse ratio on a target workload.

The mathematical gap between today's syntactic substrate and tomorrow's semantic substrate is well understood; the engineering gap is a measurable amount of work. The Cayley pair-lattice is the right data structure for the lift because the alignment and rho maps are already functorial on words; we are extending the codomain category from words to words-modulo-meaning, not changing the substrate.

9.3 Other extensions

Mutation-sensitive invariant adapter (Demo B+). The Conway / Kinoshita–Terasaka boundary closes when an HFK or refined-Khovanov adapter is added to the multi-invariant fingerprint. The schema already supports new invariant columns; the adapter is sourced data plus a Lean side that establishes invariance under component-preserving Conway mutation [13]. A Rasmussen s-invariant column [15] would be a natural addition because it is computable from the existing Frobenius edge maps in Demo E.

Full Khovanov engine beyond the trefoil (Demo E+). Replace the trefoil-specific `trefoilDifferential` with a generic cube-vertex/edge enumerator from `SignedPDGraph`, plus a generic Frobenius-edge-map constructor. The infrastructure for the trefoil is the 1-knot fixture; the generalisation is a small-PD enumerator and a chain-group constructor. A natural milestone is the Milnor-conjecture corner [45], settled combinatorially by Rasmussen [15], which exercises the categorification machinery against a non-trivial topological problem.

All-30-pair walk witnesses (Demo F+). Run Demo F's witness machinery on every ordered pair of the 30-pair lattice. The cost is linear in the lattice size at fixed walk length; the gain is a complete pair-by-pair classification table.

Hardware shots beyond the trefoil (Demo A+). Obtain Quantinuum shots for 4_1, 5_1, 5_2 at $q = e^{2\pi i/5}$. Path-B discipline forbids fabrication; the rows currently marked `hardware.status: "pending"` convert to `measured` when the runs are scheduled.

PHG sparsity as ATP search heuristic (Demo H+). The proved local-degree bound becomes a premise-ranking signal for automated theorem proving over the cross-system harness; low-degree grade-product transitions are preferred. Benchmark against current ATP retrieval.

10 Open Obligations and Honest Limitations

10.1 Substantive vs. structural, in one place

The discipline of section 7 requires every demo to declare a substantive analogue alongside its structural claim. Consolidated in one place, the substantive analogues that are *not* claimed by this paper are:

- **F:** global operator confluence over all generators and reachable elements; deferred to the divergence ledger.
- **C:** a complete cross-kernel classification theorem for all knots; this paper ships one trefoil witness.
- **H:** sparsity for arbitrary Cayley graphs; this paper proves it for the depth-4 cross-system harness only.
- **A:** a four-knot hardware demonstration; this paper ships one measured trefoil row plus three pending rows.
- **B:** a complete knot classifier; this paper ships a $> 99\%$ pairwise discriminator on the implemented fingerprint with the canonical mutation pair as an honest failure boundary.
- **G:** heat-trace periodicity in diffusion time; this paper does not claim it.
- **E:** arbitrary-diagram Khovanov homology engine; this paper ships an explicit trefoil categorification trace.

10.2 Trust assumptions

We rely on Lean's typechecker and Mathlib's mathematical library [36, 37] as foundational. We rely on the pinned KnotInfo file [32] as authority for Jones, Alexander, Conway polynomials and PD codes for prime knots up to twelve crossings. We rely on the Quantinuum public Jones-at-fifth-root release [31] for the measured trefoil row of Demo A. We rely on Spherogram [33] for face-and-crossing-sign data feeding the Goeritz signature in Demo B. None of these are reproduced in this paper; all are auditable via their respective sources.

10.3 Reproduction

The seven artifact JSONs are byte-identifiable by SHA-256 in `technique_power_index.json`. The kernel theorems are in the named Lean modules. The receipts checker is `scripts/check_cross_system_receipts.py`; the all-receipts gate is

```
python3 scripts/check_cross_system_receipts.py --all --json
```

The negative probes are recorded inside each artifact's `negative_probe` block.

References

[1] L. Kauffman. State models and the Jones polynomial. *Topology*, 26(3):395–407, 1987.

[2] V. F. R. Jones. A polynomial invariant for knots via von Neumann algebras. *Bull. Amer. Math. Soc.*, 12(1):103–111, 1985.

[3] W. Burau. Über Zopfgruppen und gleichsinnig verdrillte Verkettungen. *Abh. Math. Sem. Univ. Hamburg*, 11:179–186, 1936.

[4] M. Khovanov. A categorification of the Jones polynomial. *Duke Math. J.*, 101(3):359–426, 2000.

[5] M. Khovanov and L. Rozansky. Matrix factorizations and link homology. *Fund. Math.*, 199:1–91, 2008.

[6] J. Kock. *Frobenius Algebras and 2D Topological Quantum Field Theories*. London Math. Soc. Student Texts, vol. 59. Cambridge Univ. Press, 2003.

[7] M. Atiyah. Topological quantum field theories. *Publ. Math. I. H. É. S.*, 68:175–186, 1988.

[8] E. Witten. Quantum field theory and the Jones polynomial. *Comm. Math. Phys.*, 121(3):351–399, 1989.

[9] J. H. Conway. An enumeration of knots and links, and some of their algebraic properties. In *Computational Problems in Abstract Algebra*, pages 329–358. Pergamon, 1970.

[10] J. H. Conway. *On Numbers and Games*. Academic Press, 1976.

[11] D. E. Knuth. *Surreal Numbers*. Addison-Wesley, 1974.

[12] S. Kinoshita and H. Terasaka. On unions of knots. *Osaka Math. J.*, 9:131–153, 1957.

[13] R. Lipshitz, S. Sarkar, et al. Reduced Khovanov is mutation-invariant. *arXiv preprint*, 2602.23442, 2026.

[14] P. Ozsváth and Z. Szabó. Holomorphic disks and knot invariants. *Adv. Math.*, 186(1):58–116, 2004.

[15] J. Rasmussen. Khovanov homology and the slice genus. *Invent. Math.*, 182(2):419–447, 2010.

[16] C. McA. Gordon and R. Litherland. On the signature of a link. *Invent. Math.*, 47(1):53–69, 1978.

[17] J. Bénabou. Introduction to bicategories. In *Reports of the Midwest Category Seminar*, Lecture Notes in Math., vol. 47, pages 1–77. Springer, 1967.

[18] S. Mac Lane. *Categories for the Working Mathematician*, 2nd ed. Springer, 1998.

[19] F. W. Lawvere. Adjointness in foundations. *Dialectica*, 23:281–296, 1969.

[20] M. H. A. Newman. On theories with a combinatorial definition of equivalence. *Ann. of Math.*, 43:223–243, 1942.

[21] Y. Lafont. Interaction nets. In *POPL '90*, pages 95–108. ACM, 1990.

[22] J. Lamping. An algorithm for optimal lambda calculus reduction. In *POPL '90*, pages 16–30. ACM, 1990.

[23] J.-J. Lévy. Optimal reductions in the lambda calculus. In *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus, and Formalism*, pages 159–191. Academic Press, 1980.

[24] Leonardo of Pisa. *Liber Abaci*, 1202.

[25] D. D. Wall. Fibonacci series modulo m. *Amer. Math. Monthly*, 67:525–532, 1960.

[26] F. R. K. Chung. *Spectral Graph Theory*. Amer. Math. Soc., CBMS Reg. Conf. Series in Math., vol. 92, 1997.

[27] A. Cayley. Desiderata and suggestions: No. 2. The theory of groups: graphical representation. *Amer. J. Math.*, 1(2):174–176, 1878.

[28] L. Dorst, D. Fontijne, and S. Mann. *Geometric Algebra for Computer Science*. Morgan Kaufmann, 2007.

[29] S. Trebst, M. Troyer, Z. Wang, and A. Ludwig. A short introduction to Fibonacci anyon models. *Progr. Theor. Phys. Suppl.*, 176:384–407, 2008.

[30] P. Bonderson, K. Shtengel, and J. Slingerland. Interferometry of non-abelian anyons. *Ann. Phys.*, 323:2709–2755, 2008.

[31] Quantinuum. Less quantum, more advantage: Jones polynomial at the fifth root of unity. GitHub repository `Quantinuum/jones-polynomial`, 2025.

[32] J. C. Cha and C. Livingston. *KnotInfo: Table of Knot Invariants*. `https://knotinfo.math.indiana.edu`, accessed 2024.

[33] M. Culler, N. Dunfield, and J. Weeks. *Spherogram: a Python module for the manipulation of links*. `https://github.com/3-manifolds/Spherogram`, 2010-present.

[34] A. A. Markov. Über die freie Äquivalenz geschlossener Zöpfe. *Recueil Math. Moscou*, 1:73–78, 1936.

[35] H. N. V. Temperley and E. H. Lieb. Relations between the percolation and colouring problems. *Proc. Roy. Soc. London A*, 322:251–280, 1971.

[36] L. de Moura and S. Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction — CADE 28*, LNCS, vol. 12699, pages 625–635. Springer, 2021.

[37] The mathlib Community. The Lean mathematical library. In *CPP 2020*, pages 367–381. ACM, 2020.

[38] The Coq Development Team. *The Coq Proof Assistant Reference Manual*, version 8.18. INRIA, 2023.

[39] A. Vezzosi, A. Mörtberg, and A. Abel. Cubical Agda: a dependently typed programming language with univalence and higher inductive types. In *ICFP 2019*.

[40] C. Cohen, T. Coquand, S. Huber, and A. Mörtberg. Cubical type theory: a constructive interpretation of the univalence axiom. In *TYPES 2015*, LIPIcs, vol. 69, 2018.

[41] The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study, 2013.

[42] V. Voevodsky. The origins and motivations of univalent foundations. *Mathematical Intelligencer*, 36(2):82–85, 2014.

[43] Institute for Applied Ontological Mathematics. Boundary substrate architecture (canonical reference). Internal document, 2026-04-27.

[44] R. Goodman and the IAOM research workforce. Hostile audit and remediation for the seven-demo bundle. Internal artifact tree, 2026-05-01.

[45] J. Milnor. *Singular Points of Complex Hypersurfaces*. Princeton Univ. Press, Ann. of Math. Studies, vol. 61, 1968.

[46] K. Reidemeister. *Knotentheorie*. Ergebnisse der Math., vol. 1. Springer, 1932.

[47] L. Kauffman. *Knots and Physics*. World Scientific, 1991.

[48] V. Turaev. *Quantum Invariants of Knots and 3-Manifolds*, 2nd ed. de Gruyter Studies in Math., vol. 18, 2010.

[49] J. Baez and J. Dolan. Higher-dimensional algebra and topological quantum field theory. *J. Math. Phys.*, 36:6073–6105, 1995.

[50] T. Leinster. *Higher Operads, Higher Categories.* London Math. Soc. Lecture Note Series, vol. 298. Cambridge Univ. Press, 2004.

[51] S. Abramsky and R. Jagadeesan. New foundations for the geometry of interaction. *Inform. and Comput.*, 111(1):53–119, 1994.

[52] J.-Y. Girard. Geometry of interaction I: interpretation of system F. In *Logic Colloquium '88*, pages 221–260. North-Holland, 1989.