# Architecture of the Five-Paper Yang–Mills Reconstruction Program

Finite Coercivity, Admissible Branch Selection,
Positive-Thickness Realization, Thick-to-Thin Reconstruction,
and Spectral Closure

**A structural overview and manuscript-dependency guide**

J. Scott Little
Independent Researcher

August 2026

Version 1.3.0

**Abstract**

This note is a source-audited public guide to the architecture of a five-paper reconstruction program for four-dimensional Yang–Mills theory. Its organizing spine is

$$\text{finite coercivity} \longrightarrow \text{coherent scale-branch selection} \longrightarrow \text{positive-thickness realization} \longrightarrow \text{thick-to-thin reconstruction} \longrightarrow \text{invariant spectral closure.}$$

The document is not a sixth technical paper and does not abridge the proofs. Its purpose is to identify which manuscript owns each obligation, what each stage imports and exports, which source-native theorems anchor the four handoffs, and where conditional or branch-specific qualifications remain visible. In particular, it separates Paper-II scale selection from Paper-III derived incision, Paper-IV thin/hyperbolic reconstruction from Paper-V Poincare and invariant-mass ownership, and finite witness calculations from the general continuation theorems.

This note is also distinct from the methodological manuscript *The Architecture of Long Mathematical Proofs*. That work develops a general framework for auditing distributed arguments and uses a five-stage spectral reconstruction only as a worked case study; it explicitly does not validate the Yang–Mills program. The present note instead follows this program's source-audited manuscript record.

The five technical manuscripts were submitted to the *Journal of Mathematical Physics* in August 2026. While they remain under editorial consideration, the author has elected not to place the complete submitted set on the author's website. This overview makes the manuscript architecture publicly intelligible during that interval and is intended to remain a permanent navigation document after the technical papers become publicly available.

## Status at a glance

| | |
|---|---|
| **Document** | Source-audited architecture and navigation note. It is not a sixth technical manuscript, a substitute for Papers I–V, or a peer-review judgment on their correctness. |
| **Submitted set** | Papers I–V were submitted to the *Journal of Mathematical Physics* in August 2026. The author is not independently posting the complete submitted set while editorial consideration remains active. |
| **Visible conditions** | Full multiparameter Paper-II selection remains conditional on its extension hypotheses; the explicit $SU(2)$ depth realization is fixed-stratum. Paper-III Theorems III.17–III.18 and Paper-IV H1–H3 retain their stated conditional status. |
| **Terminal scope** | Paper V owns the Poincare, joint-spectrum, no-leakage, and invariant-mass layer on its certified terminal branches. Sector-II and non-scalar correspondence-valued conclusions retain separately declared conditions. |
| **External boundary** | The paper-owned terminal results are not presented here as an unqualified declaration that the Clay Yang–Mills problem has been resolved; that stronger problem-level claim requires a separate end-to-end audit against the accepted external requirements. |

# Contents

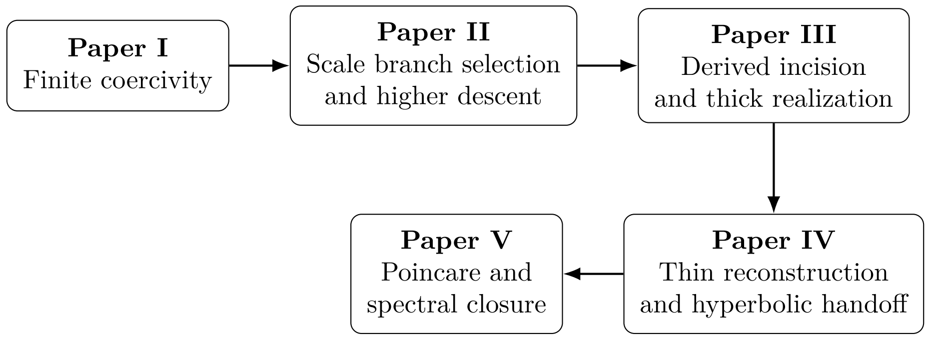


Figure 1: The five-paper dependency spine. Each arrow represents a theorem-level handoff rather than an informal change of description.

**How to use this document.** The five manuscripts contain the technical verification. The separate methodological work *The Architecture of Long Mathematical Proofs* develops a general theory of proof architecture and uses this program only as a worked case study. The present document identifies where each Yang–Mills obligation is established, what each manuscript imports, and what it hands forward.

# 1 Why the Long Proof Became Five Papers

The five-paper division reflects the internal dependency structure of the construction.

The endpoint is spectral, but the object whose spectrum is ultimately studied cannot be taken as a primitive. It must first be constructed through a sequence of mathematically distinct stages.

$$\boxed{\begin{array}{c}\text{finite coercivity} \longrightarrow \text{coherent branch selection} \longrightarrow \text{positive-thickness realization} \longrightarrow \\ \text{thick-to-thin reconstruction} \longrightarrow \text{spectral closure.}\end{array}}$$

Each arrow represents a theorem-level obligation.

Finite coercivity does not select a coherent refinement history. A coherent branch does not automatically possess a spacetime realization. A positive-thickness realization does not automatically determine a thin theory. A reconstructed thin theory does not automatically possess the spectral separation required by the final theorem.

The proof became long because these implications had to be established rather than assumed.

## 1.1 The construction changes mathematical type at every stage

The progression can be written schematically as

$$\mathsf{Finite} \xrightarrow{\mathrm{I}} \mathsf{Certified} \xrightarrow{\mathrm{II}} \mathsf{Selected} \xrightarrow{\mathrm{III}} \mathsf{Thick} \xrightarrow{\mathrm{IV}} \mathsf{Thin} \xrightarrow{\mathrm{V}} \mathsf{Spectral}.$$

These labels do not describe five discussions of the same object. They mark changes in mathematical status.

A raw finite system becomes *certified* only after the required coercive control has been proved. Certified objects become *selected* only after their refinement and descent relations have been organized into a coherent admissible branch. The selected branch becomes *thick* only after it has been realized as a positive-thickness spacetime system. The thick system becomes *thin* only after a controlled reconstruction has been carried out. Only then can the reconstructed system become spectrally certified through the final no-leakage and spectral-separation argument.

Thus the five-paper architecture is not merely organizational. It records what kind of object exists at each stage of the construction.

### 1.2 Why a direct spectral argument would hide the real obligations

One could attempt to begin much closer to the desired endpoint. For example, one might posit a continuum Hamiltonian, identify a vacuum vector, and seek a lower bound on the orthogonal spectrum.

Such an argument would be shorter. It would also leave unanswered the questions that organize the present construction:

- Where did the endpoint operator-theoretic object come from?
- Which finite structures survive refinement?
- How were incompatible local choices resolved?
- What carries nontrivial descent data?
- Which branch was actually realized?
- Why does the zero-thickness object exist?
- Why does the distinguished spectral sector survive reconstruction?

The five-paper program treats these as parts of the proof rather than as preliminary intuition.

Its organizing rule is simple: when an implication is not automatic, it must appear as a mathematical handoff.

### 1.3 The no-backward-borrowing rule

**No later conclusion may be imported backward as an earlier assumption.**

Paper I does not assume the eventual invariant spectral conclusion in order to obtain finite coercivity. Paper II does not select a branch by first deciding which branch produces the desired endpoint. Paper III does not assume the existence of the thin theory it is eventually intended to support. Paper IV does not identify its reconstructed object by appealing to the final spectrum. Paper V must analyze the system exported by Paper IV rather than replace it with a more convenient one.

The direction of dependence remains

$$\mathrm{I} \longrightarrow \mathrm{II} \longrightarrow \mathrm{III} \longrightarrow \mathrm{IV} \longrightarrow \mathrm{V}.$$

This directional discipline is one of the principal reasons the proof is distributed across five manuscripts.

### 1.4 The proof resides partly in the interfaces

Each manuscript contains substantial mathematics of its own, but the complete theorem depends equally on the validity of the handoffs.

Table 1: The four principal interfaces of the five-paper construction.

| **Interface** | Required handoff |
|---|---|
| **I–II** | Certified finite data become admissible source data for selection. |
| **II–III** | The selected branch becomes the branch actually realized. |
| **III–IV** | The realized thick system becomes the source of thin reconstruction. |
| **IV–V** | The reconstructed thin system becomes the object of spectral analysis. |

A failure at any one of these boundaries would interrupt the proof even if the neighboring manuscripts were individually correct. The arrows are therefore part of the theorem.

### 1.5 Five principal obligations

The program can be reduced to five successive questions, with the short labels in this document understood as architectural summaries rather than replacements for the source-native theorem statements.

**Paper I – finite-scale coercivity.** Can a marked finite Wilson-tube Hamiltonian package be physically certified, with exact complete-ground identification and a positive retained floor on its complement?

**Paper II – coherent scale-branch selection and higher descent.** Can the finite-good packages be organized, under the stated amalgamation and extension hypotheses, into a coherent higher pseudosection carrying one quantitative profile and one branch-appropriate coefficient-line descent type?

**Paper III – derived incision and positive-thickness spacetime realization.** Can the selected marked time-zero branch determine its obstruction carrier and controlled incised habitat, and can that branch then be realized as a regional, Cauchy-compatible positive-thickness system while preserving its complete-ground and coercive provenance?

**Paper IV – thick-to-thin reconstruction and selected hyperbolic handoff.** Can the positive-thickness diagram be reconstructed as a sector-faithful thin observable and represented system with a closed comparison–Mosco form, exact complete-ground kernel, positive off-ground thin stiffness, and a selected physical hyperbolic energy certificate?

**Paper V – Poincare, no-leakage, and invariant spectral closure.** Can the exact Paper-IV handoff be promoted on the certified terminal branches to strong translation/Poincare structure, a joint four-momentum spectral measure, an invariant mass operator, sharp no-leakage, and a positive invariant-mass threshold?

The long proof therefore retains the compact five-stage description

$$\text{certify} \longrightarrow \text{select} \longrightarrow \text{realize} \longrightarrow \text{reconstruct} \longrightarrow \text{close},$$

but the technical ownership inside those verbs is deliberately typed. In particular, Paper II selects the scale branch but does not construct the spacetime incision; Paper III derives the incision and realizes the thick regional system; Paper IV reconstructs the thin system and produces the selected hyperbolic handoff; and Paper V alone owns the invariant four-momentum/mass and physical no-leakage conclusions.

## 2 The Five-Paper Program at a Glance

The five manuscripts divide the long proof according to mathematical responsibility. The table below has been aligned to the source-audited stage contracts. The short role names are architectural summaries; they do not collapse source-native objects, constants, or theorem labels into one notation.

Table 2: Principal responsibility and export of the five manuscripts.

| **Paper** | Owns | Principal export |
|---|---|---|
| **I** | Finite Wilson-tube certificate | Marked late-tail finite packages with exact complete-ground provenance, refinement data, and retained coercive floor $\kappa_{\mathrm{cof}} > 0$. |
| **II** | Coherent scale branch and branch-appropriate higher descent | A selected higher pseudosection when the full-scale extension hypotheses are discharged, or a precisely typed fixed-stratum depth package otherwise; in either case the coefficient-line descent type and quantitative provenance remain explicit. |
| **III** | Dynamically derived incision and positive-thickness regional realization | A typed positive-thickness regional diagram carrying observable, represented, state, Gauss–Ward, entropy/seam, form, projection, sector, comparison, and provenance data, together with thin-readiness witnesses. |
| **IV** | Sector-faithful thick-to-thin reconstruction and physical hyperbolic handoff | Thin observable/represented systems, a closed thin form with exact complete-ground kernel and positive off-ground stiffness, plus the selected physical hyperbolic certificate and energy threshold handed to Paper V. |
| **V** | Poincare, invariant spectral transfer, interaction, and no-leakage | Branch-appropriate Poincare/translation implementation, joint four-momentum spectral structure, invariant mass operator, sharp no-leakage, positive invariant-mass threshold, and the separately proved interaction/non-free witness on the declared terminal branches. |

Table 3: Principal import and scope boundary at each stage.

| **Paper** | Imports | Does not yet claim |
|---|---|---|
| **I** | A marked finite Wilson-tube Hamiltonian package. | A coherent full-scale branch, spacetime realization, thin system, Poincare representation, or invariant-mass threshold. |
| **II** | A type-identified finite-good package, with the Paper-I complete-ground and coercive provenance left visible. | An incised spacetime base, regional local-normal net, thin operator, relativistic spectrum, or mass conclusion. |
| **III** | The branch actually certified by Paper II, with its exact cofinality status and coefficient-line provenance. | The completed zero-thickness system, physical hyperbolic generator, Poincare translation representation, or invariant spectral conclusion. |
| **IV** | The Paper-III positive-thickness diagram and its thin-readiness packages. | A full Poincare representation, joint four-momentum PVM, invariant-mass operator theorem, physical no-leakage theorem, or final mass threshold. |
| **V** | The exact Paper-IV thin/hyperbolic handoff certificate, including the positive selected-frame energy threshold. | Confinement, scattering completeness, hadronization, or other stronger infrared conclusions not contained in the terminal spectral and interaction theorems. |

**Canonical source labels.** The frozen authority-set labels currently associated with the sequence are:

**Paper I.** *Coarse Transport, Central Twist, and Finite-Scale Coercivity for Time-Zero Yang–Mills Wilson-Tube Systems in 3+1 Dimensions.*

**Paper II.** *Geometric Exhaustion, Entropy Surgery, and Time-Zero Branch Selection for Yang–Mills Wilson-Tube Systems.*

**Paper III.** *Spacetime Realization, Vacuum Gerbe Descent, and Spectral Threshold Formation for Twisted Yang–Mills $W^*$-Observable Systems.*

**Paper IV.** *Scale Bordism and Gerbe-Compatible Thick-to-Thin Reconstruction for Wilson-Tube Observable Systems.*

**Paper V.** *Coercive Spectral Closure and Mass-Threshold Formation for Thin Yang–Mills Wilson Observable Systems.*

### 2.1 Source-native theorem concordance

The following labels are source-native anchors from the frozen authority sets and the source-audited supplementary concordance. Page references are to the canonical or normalized PDFs identified by that audit. They are included for navigation; they do not create new theorem statements or strengthen conditional source results.

**Paper I.** Theorems XI.3 (canonical p. 18), XIII.1 (p. 21), XVI.5 (p. 31), and XVII.4 (pp. 35–36), together with Corollary XVII.5 and Scope Note XVII.6 (p. 36). Theorem XVII.4 is the controlling selected-source identity and interface-readiness boundary.

**Paper II.** Theorems 4.5 (canonical p. 25), 5.1 and Corollary 5.2 (p. 27), 6.1 (p. 30), 6.4 (p. 31), 7.1 (p. 34), and 7.2–7.3 (p. 36). Full multiparameter cofinal selection remains conditional on the extension hypotheses stated in that paper; the explicit $SU(2)$ realization is fixed-stratum depth-cofinal.

**Paper III.** Theorems II.132 (normalized p. 106), II.233 (p. 166), II.261 (p. 181), II.286 (p. 194), III.17 (p. 198), III.18 (p. 200), and G.1 (p. 257). Theorem II.286 is the selected front-end closure theorem; III.17–III.18 remain conditional at the source-audited authority level.

**Paper IV.** Theorems IV.1/A (canonical p. 21), V.7/B (p. 27), VI.1/C (p. 32), VII.4/D (p. 37), VIII.5/E (p. 40), XII.1/H1 (p. 57), XII.2/H2 (p. 58), XIII.2/H3 (p. 63), and XIV.1 (p. 64). Theorem XIV.1 is the direct physical hyperbolic handoff; it is not an invariant-mass theorem.

**Paper V.** Theorems 2.8.1 (p. 21), 3.9.1 (p. 30), 4.2.1 (p. 33), 8.4.2 (p. 72), 9.3.1 (p. 76), 10.7.1 (p. 89), 12.7.3 (p. 110), and 13.10.1 (p. 119). The independent canonical interaction result is Theorem 14.12.1 (p. 129).

### 2.2 Source-native interface concordance

The four paper interfaces also have controlling source anchors. The symbols $\Phi_{I\to II}$, $\Phi_{II\to III}$, $\Phi_{III\to IV}$, and similar maps used in architectural summaries are *not* source-native map names; they are bookkeeping labels for the typed identifications certified by the cited results.

**I $\to$ II.** Paper I Theorem XVII.4, Corollary XVII.5, and Scope Note XVII.6; Paper II Definition 2.3, Theorem 4.5, Theorem 5.1 and Corollary 5.2, and Theorem 6.4. The Paper-I readiness theorem does not itself establish membership in Paper II's finite-good locus.

**II $\to$ III.** Paper II Theorems 4.5, 5.1, 6.1, 6.4, 7.1, and 7.2–7.3; receiving control in Paper III through CPR1 and CPR8 in the coercive-provenance ledger together with Theorems II.132, II.233, II.261, and II.286. The incoming branch's fixed-stratum or full-multiparameter cofinality status must remain visible.

**III $\to$ IV.** Paper III Theorems II.132, II.233, II.261, II.286, conditional III.17–III.18, and G.1; Paper IV primitive packages $OJ$, $RJ$, $SJ$, $GJ$, and $M$, followed by Theorems IV.1/A, V.7/B, VI.1/C, VII.4/D, VIII.5/E, XII.1/H1, XII.2/H2, XIII.2/H3, and XIV.1. This interface is simultaneous reconstruction across observable, represented, state, graph-domain, form, ground, and provenance data, not a pointwise thin limit followed by gluing.

**IV $\to$ V.** Paper IV Theorems XII.1/H1, XII.2/H2, XIII.2/H3, and XIV.1; Paper V Definition 8.1.1, Theorems 8.1.2 and 8.2.1, Lemma 8.2.2, Theorems 8.3.4, 8.4.2, 9.3.1, 10.7.1, 12.7.3, and 13.10.1. The mobile-incision PDIRT theorem is a Paper-V discharge and must not be read backward into the Paper-IV handoff.

The complete dependency remains usefully summarized by

$$\mathcal{F}_{\mathrm{cert}} \longrightarrow \mathcal{B}_{\mathrm{adm}} \longrightarrow \mathcal{T}_{>0} \longrightarrow \mathcal{T}_0 \longrightarrow \mathcal{S}_{\mathrm{inv}},$$

but these symbols are intentionally schematic. In the source record, complete-ground projections, selected invariant states, GNS vectors, zero-momentum projections, zero-mass projections, and the several coercive constants are different typed objects until an owning theorem identifies or transports them.

# 3 Paper I – Finite Coercivity

## 3.1 Responsibility

Paper I owns the finite certificate. It begins from a marked finite Wilson-tube Hamiltonian package containing, in source-native form, a gauge-invariant Wilson quotient, a complete-ground projection, a selected center character, an ultraviolet/channel/residual decomposition, and an exact mixed block.

Its central question is not merely whether a finite Hamiltonian has separated eigenvalues. It is whether the finite physical package can be certified strongly enough to carry a rigid source identity into the later scale-selection problem.

Schematically,

$$\mathcal{F}_{\mathrm{marked}} \longrightarrow \mathcal{F}_{\mathrm{cert}}.$$

## 3.2 Why this stage is needed

The later papers cannot begin from an arbitrary finite approximation. They need a finite source whose ground structure, physical Wilson channel, center marking, refinement maps, mixed-block information, and quantitative floor have already been fixed.

The source-audited finite theorem package proves finite self-adjointness and exact complete-ground identification, physicalized Wilson-channel control, sectorwise ultraviolet–infrared–mixing estimates, and a positive lower bound on the complete-ground complement.

The important architectural fact is that the positive quantity entering the next paper has a specific provenance. The retained floor is not a generic symbol that later stages may silently rename.

## 3.3 Principal mechanism

The coercive mechanism is time-zero and finite-scale. It acts on the physical Wilson-tube package before a coherent full-scale branch, regional spacetime system, thin observable theory, or relativistic spectrum has been constructed.

For each selected grade and finite-demand datum, the boundary theorem exports a late-tail family with a fixed source fingerprint, refinement data, nontrivial marked source holonomy, and an exact retained floor

$$\kappa_{\mathrm{cof}} > 0.$$

The corresponding no-branch-switching statement preserves the source identity that Paper II must subsequently organize.

This is stronger than a naked finite eigenvalue computation but deliberately weaker than any continuum or invariant-mass conclusion.

**Source control.** The finite export is controlled by Paper I Theorems XI.3, XIII.1, XVI.5, and XVII.4, with Corollary XVII.5 and Scope Note XVII.6. Theorem XVII.4 and its adjacent corollary/scope note are the controlling source-identity and interface-readiness boundary.

### 3.4 Imports

Paper I imports no preferred continuum branch and no endpoint spectrum. Its inputs are finite and marked. In particular, it does not use a later spacetime realization to decide which finite packet should count as the source.

### 3.5 Exports

At architectural resolution, Paper I exports

$$\begin{aligned}\mathfrak{P}_{\mathrm{I}} = \big(&\text{marked finite package, complete-ground data, } \kappa_{\mathrm{cof}},\\ &\text{refinement maps, center/holonomy provenance}\big).\end{aligned}$$

Paper II must perform an explicit type-identification into its finite-good analytic package. Equality of notation is not the handoff theorem.

### 3.6 Firewall and handoff

The principal firewall is

$$\text{finite/cofinal coercive certification} \not\equiv \text{a coherent full-scale branch or invariant mass.}$$

Paper I does not produce a full-scale branch, a spacetime realization, a thin object, a Poincare representation, or an invariant-mass threshold.

The first handoff is therefore

$$\underbrace{\mathfrak{P}_{\mathrm{I}}^{\mathrm{cert}}}_{\text{marked finite certificate, } \kappa_{\mathrm{cof}}>0} \quad\longrightarrow\quad \underbrace{\mathfrak{P}_{\mathrm{II}}^{\mathrm{good}}}_{\text{finite-good selection input}}.$$

**Interface anchors.** Paper I Theorem XVII.4, Corollary XVII.5, and Scope Note XVII.6 control the export boundary; Paper II Definition 2.3 and Theorems 4.5, 5.1/Corollary 5.2, and 6.4 control the receiving finite-good identification.

# 4 Paper II – Admissible Branch Selection and Higher Descent

## 4.1 Responsibility

Paper II owns the scale direction: coherent branch selection, scale-bordism comparison, and branch-appropriate coefficient-line higher descent.

Its central problem is the quantifier change from certified endpointwise or depthwise data to one coherent cofinal history. Under its stated compactness, amalgamation, restriction, and finite-extension hypotheses, the target is a full multiparameter physical higher pseudosection carrying one common quantitative profile and one fixed branch-appropriate descent type.

Thus the architectural transition is more accurately written as

$$\mathcal{F}_{\text{good}} \longrightarrow \mathcal{B}_{\text{scale}},$$

where the exact status of the resulting branch—fixed-stratum depth or full multiparameter cofinal—must remain visible.

## 4.2 Why this stage is needed

Paper I supplies a rigid finite source identity. It does not prove that independently acceptable finite packages assemble into one coherent full-scale branch.

The missing burden includes directed finite-shape amalgamation, compatibility of restriction maps, common quantitative profiling, and control of the distinction between relative refinement data and coefficient-band associator data.

Pointwise existence therefore remains weaker than coherent selection:

$$\forall\lambda\ \mathcal{F}_{\text{good}}(\lambda) \neq \varnothing \qquad \not\Rightarrow \qquad \text{one coherent cofinal pseudosection.}$$

## 4.3 Principal mechanism

Paper II consumes a finite-good package; its principal upstream instance is the triality-sensitive $SU(3)$ package exported from Paper I with retained floor $\kappa_{\text{cof}}$.

At full scale, the selection theorem is conditional on the stated extension hypotheses. When those hypotheses are discharged, the resulting higher pseudosection carries one common positive profilewise bound. Separately, Paper II constructs an explicit end-to-end $SU(2)$ depth exhaustion at one fixed finite stratum, including strict and nontrivially center-banded coefficient-line globalizations with the same finite analytic shadow.

That explicit $SU(2)$ construction is a depth-cofinal witness at a fixed stratum; it is not silently promoted to a full multiparameter continuum branch.

The output is a typed obstruction-and-coercivity packet containing the physicalized packages, line-assisted scale bordisms, coefficient-line descent data, distinct relative and band classes, one quantitative profile, and one named coercive certificate.

Crucially, Paper II does *not* yet construct the spacetime obstruction locus or incision. It organizes the scale direction and exports the higher-descent data from which Paper III must derive the regional obstruction carrier.

**Source control.** The branch-selection and coefficient-line export is controlled by Paper II Theorems 4.5, 5.1, 6.1, 6.4, and 7.1–7.3, with Corollary 5.2. The source audit keeps the conditional full-scale selection theorem distinct from the explicit fixed-stratum depth realization.

## 4.4 Imports

Paper II imports the actual Paper-I certificate after an explicit type-identification. It must preserve the complete-ground projection, the selected center/background grade, marked holonomy, refinement maps, and the numerical provenance of $\kappa_{\mathrm{cof}}$.

It must not conflate that source-native floor with its separate profilewise feasibility constant.

## 4.5 Exports

Paper II exports, at the strength actually achieved on the relevant branch,

$$\begin{aligned}\mathfrak{P}_{\mathrm{II}} = \big(&\text{marked scale branch},\ \text{higher-pseudosection status},\\ &\text{scale-bordism and coefficient-line descent data},\\ &\text{obstruction/coercivity packet}\big).\end{aligned}$$

Where full multiparameter cofinality depends on the abstract extension theorem, that hypothesis remains visible rather than being replaced by the fixed-stratum witness.

## 4.6 Firewall and handoff

The principal firewall is

$$\text{coherent scale selection} \not\equiv \text{an incised regional spacetime realization.}$$

Paper II produces no incised spacetime base, no regional local-normal net, no thin operator, and no relativistic spectrum.

The second handoff is therefore

$$\underbrace{\mathfrak{P}_{\mathrm{II}}^{\mathrm{scale}}}_{\text{selected branch + descent data}} \longrightarrow \underbrace{\mathfrak{P}_{\mathrm{III}}^{>0}}_{\text{derived habitat + thick realization}}.$$

**Interface anchors.** Paper II Theorems 4.5, 5.1, 6.1, 6.4, and 7.1–7.3 control the outgoing scale package; Paper III CPR1/CPR8 and Theorems II.132, II.233, II.261, and II.286 control the receiving provenance and selected front-end realization.

# 5 Paper III – Positive-Thickness Realization

## 5.1 Responsibility

Paper III is where the regional spacetime direction first becomes structural. It consumes the marked time-zero branch exported by Paper II and owns both the dynamically derived incision and the positive-thickness regional realization.

Its central question is therefore twofold:

$$\text{selected scale branch} \longrightarrow \text{derived controlled habitat} \longrightarrow \text{positive-thickness regional system.}$$

## 5.2 Why this stage is needed

A selected higher pseudosection in the scale direction is not yet a spacetime object. Nor does Paper II supply a pre-existing incision that Paper III may simply adopt.

Paper III must derive the obstruction carrier from the compatible cofinal lifting problem, establish the support regularity needed to make that carrier geometrically usable, pass to the controlled incised habitat, and only then build the regional local-normal system.

This ordering prevents the incision or later state-surgery data from being chosen circularly to force the desired realization.

## 5.3 Principal mechanism

At the architectural level, the front end is

$$\text{marked scale branch} \longrightarrow \text{cofinal lifting problem} \longrightarrow \Sigma_{\mathrm{obs}}^{\Gamma_{\mathrm{cof}}} \longrightarrow M_{\mathrm{inc}} = \mathbb{R}^{1,3} \setminus \Sigma_{\mathrm{obs}}^{\Gamma_{\mathrm{cof}}}.$$

Under the stated support-regularity hypotheses, the derived carrier acquires the stratified, tubular, and collared structure required by the positive-thickness reconstruction.

Paper III then proves the regional time-zero extension, admissible Cauchy extraction, restriction- and Cauchy-compatible local-normal positive-thickness nets, represented descent and obstruction-carrier data, transported complete-ground and coercive provenance, and the thin-readiness ledger. Where positive-thickness entropy/seam surgery is part of a branch realization, it belongs here rather than being silently treated as a Paper-II spacetime operation.

One additional firewall remains important: if the incoming Paper-II object is certified only as a fixed-stratum depth-cofinal witness, the missing multiparameter cofinal extension remains a separate load-bearing obligation. Paper III does not promote it by notation.

**Source control.** The positive-thickness realization is controlled by Paper III Theorems II.132, II.233, II.261, II.286, III.17, III.18, and G.1. Theorem II.286 is source-audited as selected front-end closure; III.17–III.18 retain their conditional status.

### 5.4 Imports

Paper III imports the branch actually certified by Paper II, with its exact status left visible. The package includes complete-ground data, scale-correspondence and coefficient-line provenance, center/descent marking, and a named coercive certificate.

It does not import a ready-made spacetime incision.

### 5.5 Exports

Paper III exports

$$\begin{aligned}\mathfrak{P}_{\mathrm{III}} = \big(&\text{derived incision/habitat},\ \text{positive-thickness regional diagram},\\ &\text{represented descent and Gauss–Ward/state/seam data},\\ &\text{coercive provenance},\ \text{thin-readiness witnesses}\big).\end{aligned}$$

The readiness ledger records what Paper IV is entitled to consume; it is not itself the thin theory.

### 5.6 Firewall and handoff

The principal firewall is

$$\text{positive-thickness regional realization} \not\equiv \text{completed zero-thickness reconstruction.}$$

Paper III does not construct the thin algebra or representation, the closed thin form, the Paper-IV physical hyperbolic generator, a Poincare translation representation, or the invariant spectral conclusion.

The third handoff is therefore

$$\underbrace{\mathfrak{P}_{\mathrm{III}}^{>0}}_{\text{thick diagram + readiness}} \longrightarrow \underbrace{\mathfrak{P}_{\mathrm{IV}}^{0}}_{\text{sector-faithful thin reconstruction}}.$$

**Interface anchors.** Paper III Theorems II.132, II.233, II.261, II.286, conditional III.17–III.18, and G.1 control the thick export; Paper IV's primitive packages and Theorems IV.1/A through XIV.1 control the simultaneous reconstruction intake and selected hyperbolic handoff.

## 6 Paper IV – Thick-to-Thin Reconstruction

### 6.1 Responsibility

Paper IV owns sector-faithful thick-to-thin reconstruction and the selected physical hyperbolic handoff to Paper V.

Its task is not merely to produce a zero-thickness object. It must reconstruct the Paper-III thick datum as an ordinary and represented thin observable system while preserving branch type, complete-ground multiplicity, Gauss–Ward and surgery domains, comparison provenance, and the ancestry of the positive threshold.

Schematically,

$$\mathcal{T}_{>0} \longrightarrow \mathcal{T}_0 \longrightarrow \mathcal{H}^{\mathrm{sel}}_{\mathrm{hyp}}.$$

## 6.2 Why this stage is needed

A formal limit can exist while enlarging the zero-energy kernel, erasing sector provenance, or losing the comparison data needed downstream. Paper IV therefore treats zero thickness as a reconstruction problem rather than parameter substitution.

The source-audited burden includes ordinary incised thin observable descent, branch-appropriate represented descent, restart and lifting closure, Gauss–Ward/collar reconstruction, and a closed comparison–Mosco limit on the correctly identified thin carrier.

## 6.3 Principal mechanism

Paper IV consumes the positive-thickness diagram and thin-readiness packages from Paper III. Its theorem chain establishes, in source-native form:

- ordinary normal incised thin observable descent;
- branch-appropriate represented descent;
- corrected state/lifting/restart closure and the inherited expectation graph;
- Gauss–Ward and collar reconstruction;
- a closed comparison–Mosco limit with exact complete-ground kernel and positive off-ground thin stiffness;
- compact-semisimple/global-form coefficient realization and explicit analytic/physical non-vacuity witnesses; and
- the exact branchwise certification shell needed for the Paper-IV-to-Paper-V handoff.

The physical handoff identifies the selected hyperbolic generator on the thin carrier as

$$H_{\mathrm{hyp}} = K_{\mathrm{thin}}^{1/2},$$

with the source-native normalization $c_E = 1$ and a selected-frame energy-squared threshold

$$\Delta^2_{E,\mathrm{IV}} = c_E\, \kappa_{\mathrm{IV,thin}}.$$

This is a strong endpoint for Paper IV. It is still not an invariant-mass theorem.

**Source control.** The thin reconstruction and hyperbolic handoff are controlled by Paper IV Theorems IV.1/A, V.7/B, VI.1/C, VII.4/D, VIII.5/E, XII.1/H1, XII.2/H2, XIII.2/H3, and XIV.1. The last of these is a selected-frame hyperbolic/energy theorem, not the invariant-mass theorem.

### 6.4 Imports

Paper IV imports the Paper-III positive-thickness diagram and its thin-readiness data, together with the branchwise observable, represented, state/surgery, Gauss–Ward, and analytic comparison packages specified by the reconstruction category.

### 6.5 Exports

Paper IV exports

$$\begin{aligned}\mathfrak{P}_{\mathrm{IV}} = \big(&\text{thin observable/represented systems},\ \text{closed thin form/operator},\\ &\text{exact complete-ground kernel},\ \text{comparison/projection data},\\ &\text{selected hyperbolic certificate},\ \Delta^2_{E,\mathrm{IV}}\big).\end{aligned}$$

The export includes the unresolved Paper-V obligation ledger. That ledger is part of the handoff rather than an admission that Paper IV should have proved those later theorems.

### 6.6 Firewall and handoff

The principal firewall is

$$\text{selected-frame thin stiffness/energy separation} \not\equiv \text{invariant mass or physical no-leakage}.$$

Paper IV does not produce the full Poincare translation representation, joint four-momentum PVM, forward-cone theorem, non-ground null-orbit exclusion, invariant-mass operator theorem, physical no-leakage theorem, or final mass threshold.

The fourth handoff is therefore

$$\underbrace{\mathfrak{P}_{\mathrm{IV}}^{\mathrm{thin/hyp}}}_{\text{thin + hyperbolic handoff}} \longrightarrow \underbrace{\mathfrak{P}_{\mathrm{V}}^{\mathrm{inv}}}_{\text{Poincare + invariant closure}}.$$

**Interface anchors.** Paper IV Theorems XII.1/H1, XII.2/H2, XIII.2/H3, and XIV.1 control the exact handoff; Paper V Definition 8.1.1 and the sequence through Theorem 13.10.1 own the Poincare, joint-spectrum, null-exclusion, no-leakage, and invariant-mass discharge.

## 7 Paper V – Spectral Closure

### 7.1 Responsibility

Paper V owns invariant spectral and interaction closure. It is the first stage permitted to construct the final four-momentum spectral object and to identify the selected-frame hyperbolic threshold with an invariant mass threshold under the stated terminal hypotheses.

Its architectural transition is

$$\mathcal{T}_0 + \mathcal{H}^{\mathrm{sel}}_{\mathrm{hyp}} \longrightarrow \mathcal{S}_{\mathrm{inv}}.$$

## 7.2 Why this stage is needed

Paper IV supplies a thin system with exact complete-ground kernel and a positive selected-frame energy threshold. That is not yet a relativistic invariant statement.

Paper V must construct the branch-appropriate represented Poincare implementation, commuting translation generators, the joint four-momentum projection-valued measure, and the invariant mass operator. It must then exclude precisely the spectral leakage mechanisms that could destroy the threshold when passing from the selected time direction to the full Lorentz-invariant spectrum.

This is the last major change of mathematical type in the five-paper chain.

## 7.3 Principal mechanism

The source-audited terminal mechanism contains several distinct steps.

First, the thin form/operator package is closed on the correct carrier and the ground projections are stabilized with the required Mosco/Riesz controls.

Second, Paper V constructs the translation generators $P^\mu$ and their joint spectral measure $E$, and defines

$$M^2 = P_\mu P^\mu.$$

The construction includes the branch-appropriate Poincare realization, forward-cone support, and the same-carrier time-axis identification relating the selected hyperbolic evolution to the time-translation generator.

Third, the proof distinguishes exact complete-ground identification, zero selected-frame energy, zero four-momentum, and zero invariant mass until the owning theorems identify them. The non-ground null-orbit exclusion and the sharp no-leakage theorem are what prevent the positive threshold from being lost along lightlike or otherwise prohibited spectral directions.

The threshold transfer then yields, on the certified terminal branch in the theorem's scope, an inequality of the schematic form

$$M^2 \geq m_{\mathrm{inv}}^2(1 - P_\Omega), \qquad m_{\mathrm{inv}} > 0.$$

Paper V also carries a separately proved interaction package, including connected-cumulant/Wick-defect witnesses and a Haag-type non-equivalence statement on the declared canonical branches. This interaction conclusion is kept distinct from the spectral-gap statement rather than being inferred from it.

**Source control.** The terminal invariant spectral chain is controlled by Paper V Theorems 2.8.1, 3.9.1, 4.2.1, 8.4.2, 9.3.1, 10.7.1, 12.7.3, and 13.10.1; Theorem 14.12.1 is the independent canonical interaction result.

## 7.4 Imports

Paper V imports the exact Paper-IV handoff certificate: the incised thin observable and represented systems, invariant-state preparation data, complete-ground face and projections, Gauss–Ward/collar package, closed thin form and represented operator, selected hyperbolic group and generator, branch typing, certification shell, and positive thin/energy threshold.

### 7.5 Exports

At terminal resolution, Paper V exports

$$\begin{aligned}\mathfrak{P}_{\mathrm{V}} = \big(&\text{Poincare/translation representation},\ E,\ P^{\mu},\ M^2,\\ &\text{sharp no-leakage},\ \text{positive invariant-mass threshold},\\ &\text{interaction witness}\big).\end{aligned}$$

The terminal scope remains branch-sensitive. Fully certified Sector-I and scalar Sector-III branches receive the stated invariant spectral conclusions; Sector-II and non-scalar correspondence-valued conclusions retain their separately declared conditions.

### 7.6 Firewall and closure of the chain

The final firewall is scope:

$$\begin{array}{c}\text{positive invariant-mass separation}\\ +\text{ non-free interaction witness}\end{array} \not\equiv \begin{array}{c}\text{confinement}\\ \text{or scattering completeness}\end{array}.$$

Confinement, scattering completeness, hadronization, and broader transport-groupoid interpretations remain separate questions unless independently proved.

The corrected terminal spine is therefore

$$\begin{aligned}\text{finite certificate} &\longrightarrow \text{coherent scale branch}\\ &\longrightarrow \text{derived incision and thick regional realization}\\ &\longrightarrow \text{thin reconstruction and selected hyperbolic handoff}\\ &\longrightarrow \text{Poincare, no-leakage, and invariant spectral closure.}\end{aligned}$$

## 8 Higher Descent, Obstruction Carrier, Incision, and Realization

Sections 3–7 follow manuscript ownership vertically. This section follows the early cross-paper mechanism horizontally, with the source-audited ownership boundaries left explicit:

$$\begin{aligned}\text{Paper I finite certificate} &\longrightarrow \text{Paper II scale branch/higher descent}\\ &\longrightarrow \text{Paper III obstruction carrier and incision}\\ &\longrightarrow \text{Paper III positive-thickness regional realization.}\end{aligned}$$

The important correction is that higher-descent information and the spacetime carrier are not owned by the same paper.

## 8.1 Paper I supplies marked source information

Paper I exports a finite/cofinal source packet carrying complete-ground, center-grade, marked-holonomy, refinement, and coercive provenance. Those markings are later inputs to the scale-selection problem, but Paper I does not construct a regional obstruction locus.

## 8.2 Paper II organizes the scale direction

Paper II asks whether the finite-good packages admit one coherent higher pseudosection with a common quantitative profile and a fixed branch-appropriate coefficient-line descent type. It records relative and band classes and produces the typed obstruction-and-coercivity packet needed downstream.

The relevant logical distinction is

$$\text{higher compatibility/obstruction data in scale} \not\equiv \text{a spacetime incision.}$$

A nontrivial higher class may constrain what kind of carrier can realize the branch, but Paper II does not choose or construct the regional carrier on that basis alone.

## 8.3 Paper III derives the obstruction carrier

Paper III formulates the simultaneous cofinal regional lifting problem and derives the common obstruction carrier from it. Before support regularity is proved, this carrier is a lifting obstruction locus rather than a fully geometric defect surface. Under the stated support-regularity hypotheses, it acquires the structure needed to define the controlled incised habitat.

Thus the direction is

$$\text{selected scale branch and lifting data} \longrightarrow \Sigma_{\mathrm{obs}}^{\Gamma_{\mathrm{cof}}} \longrightarrow M_{\mathrm{inc}}.$$

The incision is upstream of the final spectral theorem. It is not chosen because a later branch produces a desirable mass threshold.

**Obstruction to strict globalization does not imply obstruction to globalization.**
The obstruction data must be analyzed at their own type; where nontrivial higher data survive, they constrain the carrier in which the represented branch can be realized.

## 8.4 The incision is a derived regional object

The source audit distinguishes several objects that should not be conflated: a marked source carrier, a dynamically derived incision, a chosen incised coordinate, the mobile class of avoiding incisions, and the later incision-relative descended net.

The present overview therefore uses *derived incision* only for the Paper-III construction that follows the obstruction-carrier and support-regularity analysis.

### 8.5 Positive-thickness realization follows the derived habitat

Once the controlled habitat has been obtained, Paper III constructs the regional time-zero extension, admissible Cauchy extraction, local-normal positive-thickness nets, represented descent, and the remaining branchwise state/seam/form/projection data.

The early cross-cutting mechanism is therefore

$$\boxed{\begin{array}{c}\text{finite marked certificate} \to \text{coherent scale branch and higher descent} \to \text{dynamically}\\ \text{derived obstruction carrier} \to \text{controlled incision} \to \text{positive-thickness regional realization}\end{array}}$$

### 8.6 Location in the proof record

| Structure | Primary ownership |
|---|---|
| Finite complete-ground/coercive certificate | Paper I |
| Cofinal scale selection / coefficient-line descent | Paper II |
| Fixed-stratum $SU(2)$ depth witness | Paper II |
| Regional lifting obstruction carrier | Paper III |
| Dynamically derived incision / controlled habitat | Paper III |
| Regional/Cauchy positive-thickness realization | Paper III |

This is exactly the kind of cross-interface mechanism that the general framework in *The Architecture of Long Mathematical Proofs* is designed to audit, while the present document preserves the source-native manuscript ownership specific to this Yang–Mills program.

## 9 Why Thick-to-Thin Is Reconstruction, Not a Limit

The late half of the construction begins from the Paper-III positive-thickness regional diagram and ends with the Paper-V invariant spectral theorem. The source-audited chain is more precise than a single symbol $\varepsilon \to 0$:

$$\begin{aligned}\text{thick regional system} &\longrightarrow \text{thin reconstruction}\\ &\longrightarrow \text{selected hyperbolic handoff}\\ &\longrightarrow \text{Poincare/joint spectrum}\\ &\longrightarrow \text{no-leakage}\\ &\longrightarrow \text{invariant mass threshold.}\end{aligned}$$

### 9.1 Why ordinary convergence is too weak

The Paper-III system carries branch identity, complete-ground data, fixed-thickness descent type, incision provenance, regional/Cauchy compatibility, and an inherited coercive certificate. A limit that forgets any load-bearing part of that packet would not be the object required by Paper V.

Thus

$$\text{existence of a zero-thickness limit} \not\equiv \text{sector-faithful reconstruction.}$$

### 9.2 The thin object is an output

Paper IV does not start from a pre-existing thin Yang–Mills theory. It constructs the thin observable and represented systems from the admissible positive-thickness data supplied by Paper III.

The reconstruction is branchwise and incision-aware. It retains the distinction between ordinary and represented targets and does not identify different carrier types merely because the same notation can be reused.

### 9.3 Comparison–Mosco closure belongs to Paper IV

The analytic limiting step in Paper IV produces a closed comparison–Mosco limit, exact complete-ground kernel, and positive off-ground thin stiffness threshold on the thin carrier. Together with the descent, lifting, restart, Gauss–Ward, and collar theorems, this makes the thin system an actual output of the proof rather than a formal endpoint symbol.

Strong control of the ground projections and comparison maps is part of the handoff preparation. It is a necessary input to the final spectral argument, but it should not itself be called the physical no-leakage theorem.

### 9.4 Paper IV stops at a selected physical hyperbolic certificate

The final Paper-IV step identifies the selected hyperbolic generator and its positive selected-frame energy-squared threshold. This is the precise boundary at which the reconstruction manuscript stops.

The distinction is

$$\text{selected-frame energy threshold} \not\equiv \text{invariant mass threshold}.$$

The latter requires the full four-momentum spectral construction of Paper V.

### 9.5 Poincare and joint spectral structure belong to Paper V

Paper V constructs commuting translation generators and the joint PVM $E$, establishes the branch-appropriate Poincare implementation, and defines $M^2 = P_\mu P^\mu$. The source audit explicitly forbids treating a complete-ground projection, selected invariant state, GNS vector, zero-momentum projection, and zero-mass projection as the same object until the relevant theorem identifies them.

This typed discipline is why the terminal stage cannot be reduced to "the gap survives the limit."

### 9.6 No-leakage is a terminal spectral theorem

Physical no-leakage belongs to Paper V. It uses the reconstructed thin carrier, the Paper-IV threshold provenance, projection control, light-cone/forward-cone structure, and the exclusion of non-ground null orbits to prevent the selected threshold from disappearing in the invariant spectral assembly.

Only after those steps can the selected-frame energy separation be transferred to the invariant mass operator.

### 9.7 The corrected late-stage diagram

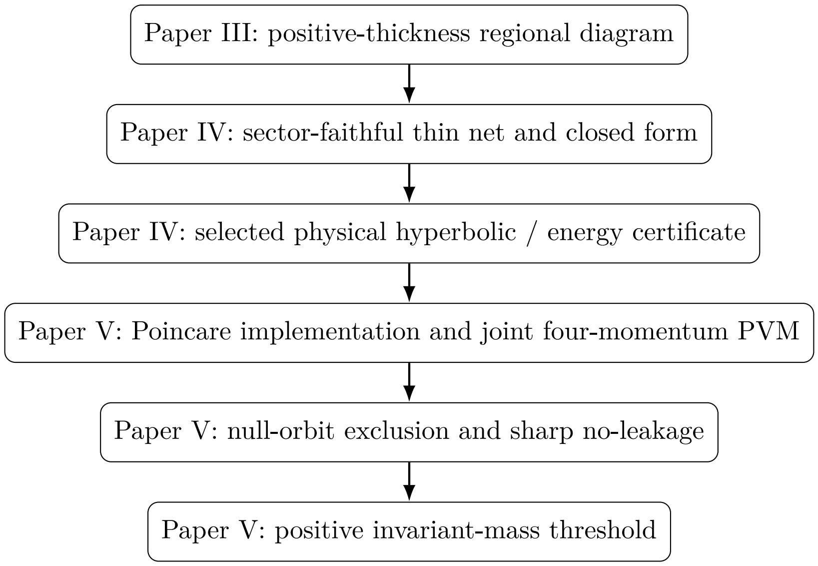


Figure 2: The source-audited late-stage mechanism. Paper IV reconstructs the thin system and selected hyperbolic certificate; Paper V owns the Poincare, no-leakage, and invariant-mass steps.

### 9.8 What reconstruction is required to remember

The governing rule remains

**Forget auxiliary presentation; preserve admissible provenance.**

In source-native terms, the provenance includes branch/sector type, complete-ground multiplicity, incision and surgery data where active, Gauss–Ward domains, comparison maps, and the ancestry of the positive threshold. The proof may forget redundant presentation choices; it may not forget the data needed to authorize the Paper-IV-to-Paper-V handoff.

## 10 Explicit Witnesses and Reproducibility

The five-paper construction contains explicit witness material at more than one stage. A fidelity audit therefore should not compress all $SU(2)$ and $SU(3)$ information into "finite examples in Paper I." Different witnesses certify different obligations.

### 10.1 Witnesses are theorem-local

The general rule is

$$\text{explicit witness} \neq \text{general continuation theorem.}$$

A witness can demonstrate non-vacuity, distinguish two organizational behaviors, or verify a source-native quantitative packet. It does not automatically discharge the later branch, realization, reconstruction, or invariant-spectral interfaces.

### 10.2 Paper-I source witness

Paper I carries the marked finite Wilson-tube source package and the exact retained floor $\kappa_{\text{cof}} > 0$ on the complete-ground complement. The source-audited interface to Paper II tracks an actual triality-sensitive $SU(3)$ packet, including its center character, marked holonomy, UV–IR–mixing profile, refinement data, and coercive provenance.

The role of this witness is finite and source-identifying. It is not a continuum branch or invariant-mass calculation.

### 10.3 Paper-II fixed-stratum control witness

Paper II constructs an explicit end-to-end $SU(2)$ depth exhaustion at one fixed finite stratum. In that witness, strict and nontrivially center-banded coefficient-line globalizations have the same finite analytic shadow.

This example is especially useful because it separates two questions:

$$\text{finite analytic shadow} \qquad \text{and} \qquad \text{higher coefficient-line globalization.}$$

Its scope must remain exact: the construction is depth-cofinal at the fixed stratum; it is not, by itself, a proof of a full multiparameter continuum branch.

### 10.4 Paper-III readiness witnesses

Paper III exports a thin-readiness ledger, including the named TR1–TR12 readiness witnesses, attached to the positive-thickness regional diagram. Their purpose is to certify that the data required by Paper IV have actually survived regionalization: branch identity, complete-ground family, descent type, incision provenance, Cauchy compatibility, form/projection data, and coercive ancestry.

These are continuation witnesses, not a thin system.

### 10.5 Paper-IV non-vacuity and threshold witnesses

Paper IV contains explicit analytic and physical non-vacuity witnesses for the thin reconstruction and proves the exact complete-ground kernel and positive off-ground thin stiffness needed for the selected hyperbolic handoff.

Again, this does not yet certify an invariant mass operator or physical no-leakage. It certifies that the reconstructed thin/hyperbolic package delivered to Paper V is nonempty and quantitatively separated in the required selected-frame sense.

### 10.6 Paper-V interaction witnesses

Paper V contains the terminal $SU(2)/SU(3)$ interaction witness package, including connected-cumulant and Wick-defect information and its survival in the thin reconstructed setting, together with a separately proved Haag-type non-equivalence result on the declared canonical branches.

The architecture should therefore keep two terminal statements separate:

$$\text{positive invariant-mass threshold} \qquad \text{and} \qquad \text{non-quasi-free / non-free interaction witness.}$$

Neither is inferred from the other merely by physical expectation.

### 10.7 Compact witness ledger

| **Paper** | Witness type | What it certifies |
|---|---|---|
| **I** | Triality-sensitive marked finite source packet | Exact finite source identity, complete-ground/coercive provenance, center/holonomy marking, and retained floor $\kappa_{\mathrm{cof}}$. |
| **II** | Fixed-stratum $SU(2)$ depth exhaustion | Non-vacuity of strict and center-banded coefficient-line scale behavior at the stated depth scope; not full multiparameter cofinality. |
| **III** | TR1–TR12 thin-readiness witnesses | Survival of the data needed to enter the thick-to-thin reconstruction. |
| **IV** | Analytic/physical thin non-vacuity witnesses | Nonempty sector-faithful thin reconstruction, exact complete-ground kernel, and selected-frame positive threshold. |
| **V** | $SU(2)/SU(3)$ interaction witnesses | Non-quasi-free interaction structure and its distinction from the invariant spectral conclusion. |

### 10.8 What should be independently reproducible

A public reproducibility record should identify, for each witness:

- the owning manuscript and theorem or certificate;
- the carrier and branch/sector on which the witness lives;
- the finite or regional input data and normalization conventions;
- the exact quantity or cohomological/coefficent-line datum being checked;
- whether the result is exact algebra, numerical evaluation, analytic theorem, or conditional full-scale extension; and

- the downstream theorem that consumes the witness.

This is more informative than a single undifferentiated "$SU(2)/SU(3)$ reproducibility" label.

### 10.9 The reproducibility firewall

Reproducing a finite or branch-local witness is not the same as reproducing the long proof. Conversely, a general continuation theorem does not make the explicit source witnesses dispensable.

The correct division is

$$\begin{aligned} \text{witness} &\longrightarrow \text{typed source or non-vacuity certificate,} \\ \text{theorem chain} &\longrightarrow \text{transport of that certificate through the program.} \end{aligned}$$

## 11 Scope, Conditional Dependencies, and Explicit Nonclaims

The final theorem of a long construction is easiest to misread when the distinction between proved statements, stated hypotheses, concrete witnesses, and physical interpretations becomes blurred.

The five-paper program therefore maintains an explicit scope discipline.

At the highest level, the construction establishes a directed chain

$$\mathcal{F}_{\mathrm{cert}} \longrightarrow \mathcal{B}_{\mathrm{adm}} \longrightarrow \mathcal{T}_{>0} \longrightarrow \mathcal{T}_0 \longrightarrow \mathcal{S}_{\mathrm{inv}}.$$

Each arrow is established under the hypotheses stated in its owning manuscript. Those hypotheses remain part of the ancestry of the final theorem.

The appropriate logical form is therefore

$$\mathsf{H}_{\mathrm{I}} + \mathsf{H}_{\mathrm{II}} + \mathsf{H}_{\mathrm{III}} + \mathsf{H}_{\mathrm{IV}} + \mathsf{H}_{\mathrm{V}} \Longrightarrow \mathcal{S}_{\mathrm{inv}},$$

with the technical manuscripts providing the authoritative theorem statements and hypothesis lists.

### 11.1 What the five-paper chain claims

The program claims, within its stated domain, that finite Yang–Mills source data can be quantitatively certified; the certified data can be organized into a coherent scale branch at the theorem strength actually discharged, with coefficient-line descent and cofinality status left visible; the selected scale branch can be given a dynamically derived incision and realized as a positive-thickness regional system; the realized system can be reconstructed as a controlled thin system with an exact complete-ground kernel and selected physical hyperbolic handoff; and, on the certified terminal branches, the exact thin/hyperbolic handoff supports the stated Poincare, no-leakage, invariant-mass, and separately proved interaction conclusions.

The complete claim belongs to the composition. No intermediate stage is intended to stand in for the final theorem.

### 11.2 Hypotheses remain visible

A hypothesis verified in an earlier paper does not disappear when it is carried forward. Likewise, an ambient theorem imported from established mathematics should not be confused with a condition whose satisfaction must be proved inside the present construction.

The architecture therefore distinguishes

$$\text{external mathematical import} \neq \text{internal hypothesis verification.}$$

The former may be cited. The latter is part of the proof burden.

### 11.3 No hidden backward assumptions

The construction also enforces the rule that the endpoint may not determine its own ancestry.

Accordingly, the endpoint spectral conclusion is not assumed to obtain finite coercivity; a preferred endpoint sector is not used to choose the admissible branch; a thin theory is not assumed in order to justify positive-thickness realization; and the desired spectrum is not used to identify the reconstructed object after the fact.

Later stages may demonstrate that earlier structure survived. They may not create that structure retroactively.

### 11.4 Witnesses and general theorems are different claims

The explicit $SU(2)$ and $SU(3)$ calculations provide finite witnesses. They do not replace the general persistence arguments.

Thus

$$\text{finite witness} \neq \text{general reconstruction theorem.}$$

Similarly,

$$\text{finite numerical positivity} \neq \text{continuum invariant positivity without a persistence argument.}$$

The numerical and finite calculations establish concrete instances. The analytic chain determines what may be concluded from them.

### 11.5 Strict behavior is not assumed universally

The existence of strictifiable examples does not imply that all admissible sectors strictify.

The architecture explicitly allows the possibility that nontrivial higher compatibility data survive. The $SU(2)$ and $SU(3)$ examples serve as finite benchmarks for these different possibilities rather than universal definitions of two gauge-group classes.

### 11.6 The finite system is a source, not the endpoint system

The construction does not identify the finite source system with the reconstructed thin system.

Finite structures are certified and then transported through several nontrivial changes of mathematical type. The point of Papers II–IV is precisely to justify those changes.

Likewise, positive thickness is used constructively. The program does not require the assertion that physical spacetime possesses a fundamental nonzero thickness.

Its role is mathematical: positive thickness is a realization regime, not an ontological claim.

### 11.7 Higher carriers are mathematical structure

Where higher or non-strict carrier structure is required, its appearance should not be interpreted automatically as new particle content or a new independently observable physical field.

The carrier organizes descent and transport. Any stronger physical interpretation would require an argument separate from the five-paper reconstruction.

### 11.8 Spectral separation is not confinement

The terminal theorem is spectral. Confinement is a distinct infrared question.

**Spectral separation is not, by itself, a confinement theorem.**

In particular, the spectral result does not by itself establish an area law, flux-tube description, hadronic spectrum, or any other stronger confinement criterion. Any relation between the two must be established by additional mathematics.

### 11.9 No global scattering conclusion is imported

The spectral theorem likewise does not imply the existence of a conventional global scattering object.

The five-paper construction does not require a global $S$-matrix as a primitive or as an automatic consequence of spectral closure. Any later scattering or asymptotic analysis must be formulated using the structures actually supported by the reconstructed theory and proved separately.

Thus

$$\text{spectral closure} \not\Rightarrow \text{global scattering theory.}$$

### 11.10 The route is not claimed to be unique

The program establishes one reconstruction architecture under its stated hypotheses. It does not claim that every mathematically valid construction of four-dimensional Yang–Mills theory must use identical finite packets, identical carrier language, identical positive-thickness geometry, or identical reconstruction machinery.

The claim is that this construction closes under its proved handoffs, not that it is the only possible route.

### 11.11 Later developments are not imported backward

The broader Geometry of Admissible Transport program has continued to develop beyond the version of the technology used in the five submitted Yang–Mills manuscripts.

Later formulations may clarify or generalize parts of the construction. They are not unstated premises of the submitted proof. The five-paper chain must therefore be evaluated from the mathematics actually present in its proof record and its explicit imports.

### 11.12 This document is not the proof

Architectural exposition is not technical verification.

A short paragraph here may correspond to a long argument in its owning manuscript. A single arrow may compress several lemmas, compatibility conditions, and analytic estimates.

The present document tells the reader what the arrow means and where it is established. The manuscripts establish it.

### 11.13 Compact scope ledger

The source-audited ownership may be summarized as follows.

| Statement | Status within the five-paper program |
|---|---|
| Marked finite Wilson-tube package with exact complete-ground identification and retained floor $\kappa_{\mathrm{cof}} > 0$ | Paper I |
| Coherent higher pseudosection with one common profile at full multiparameter scale | Paper II, under its stated amalgamation/extension hypotheses |
| Fixed-stratum $SU(2)$ depth exhaustion with strict and center-banded coefficient-line globalizations | Paper II; explicit witness with deliberately narrower quantifier scope |
| Dynamically derived regional obstruction carrier and incision | Paper III |
| Regional/Cauchy-compatible positive-thickness local-normal realization and thin-readiness ledger | Paper III |
| Sector-faithful thin observable/represented reconstruction, closed comparison–Mosco form, exact complete-ground kernel, and selected hyperbolic energy threshold | Paper IV |
| Poincare/translation implementation, joint PVM, invariant mass operator, non-ground null-orbit exclusion, sharp no-leakage, and positive invariant-mass threshold | Paper V on the certified terminal branches |
| Haag-type non-equivalence / non-quasi-free interaction witness | Paper V, separately from the mass-threshold theorem |
| Confinement, scattering completeness, or hadronization follows automatically | Not claimed |
| All branches or correspondence-valued sectors receive the same terminal theorem unconditionally | Not claimed |
| This architecture note substitutes for the detailed proofs | Not claimed |

Several source-native firewalls should remain visible throughout: the complete-ground projection, selected invariant state, GNS vector, zero-momentum projection, and zero-mass projection are distinct until identified; likewise $\kappa_{\mathrm{cof}}$, the Paper-II profile bound, $\kappa_{\mathrm{IV,thin}}$, $\Delta^2_{E,\mathrm{IV}}$, the Paper-V certified threshold, and the final invariant-mass constant have different ownership and scope.

The governing reading rule remains:

> Attribute to the construction exactly the conclusion established after its stated hypotheses, branch restrictions, carrier identifications, and handoffs—and no more.

### 11.14 External endpoint boundary

The source-audited Paper-V record makes one further distinction that this public architecture should preserve. The manuscript's terminal paper-owned spectral conclusions are not themselves presented as an unqualified declaration that the Clay Yang–Mills problem has been resolved. That stronger external claim is reserved to a separate end-to-end audit against the accepted axiomatic, nontriviality, gauge-group, and four-dimensional continuum requirements. The present document maps the paper-owned theorem chain; it does not substitute for that external problem-level audit.

# 12 How to Read the Five-Paper Program

The manuscripts form one directed proof, but readers may enter from different mathematical interests. For complete verification, the authoritative order remains

$$\mathrm{I} \longrightarrow \mathrm{II} \longrightarrow \mathrm{III} \longrightarrow \mathrm{IV} \longrightarrow \mathrm{V}.$$

## 12.1 Entry points

**Paper I.** Readers interested in finite Wilson-tube Hamiltonians, exact complete-ground identification, center marking, UV–IR–mixing control, and finite/cofinal coercivity should begin here.

**Paper II.** Readers interested in cofinal scale selection, higher pseudosections, scale bordisms, coefficient-line descent, and the distinction between fixed-stratum depth witnesses and full multiparameter cofinality should begin here.

**Paper III.** Readers interested in the dynamically derived obstruction carrier and incision, regional/Cauchy realization, vacuum-gerbe/represented descent, local-normal positive-thickness systems, and thin-readiness data should begin here.

**Paper IV.** Readers interested in incised thin descent, comparison–Mosco closure, exact complete-ground kernel persistence, Gauss–Ward/collar reconstruction, and the selected physical hyperbolic handoff should begin here.

**Paper V.** Readers interested in Poincare implementation, translation generators, joint spectral measures, no-leakage, invariant mass, and the terminal interaction witness should begin here.

## 12.2 Verification requires interface checks

A full audit should ask not merely whether each paper contains correct mathematics, but whether the exported packet is actually admissible as the next paper's input.

| **Interface** | Verification question |
|---|---|
| **I–II** | Is the Paper-I Hamiltonian/complete-ground/Wilson/center/refinement packet explicitly type-identified with a Paper-II finite-good package, without branch switching or conflating $\kappa_{\mathrm{cof}}$ with a profilewise bound? |
| **II–III** | Does Paper III receive the branch actually certified by Paper II with its exact fixed-stratum or full-cofinal status and coefficient-line type, rather than importing an unproved global branch or prechosen incision? |
| **III–IV** | Does Paper IV reconstruct the regional thick diagram and thin-readiness data actually exported by Paper III, preserving incision, branch, Gauss–Ward, comparison, and complete-ground provenance? |
| **IV–V** | Does Paper V begin from the exact thin/hyperbolic certificate of Paper IV and then prove, rather than assume, the strong translation/Poincare, joint-spectrum, null-orbit, no-leakage, and invariant-mass identifications? |

A break at any one of these interfaces breaks the composed theorem even if neighboring theorem blocks are separately valid.

## 12.3 Vertical and horizontal readings

The five-paper order is the *vertical* reading of proof responsibility. A *horizontal* reading follows one object or invariant as its type changes.

For example, the protected low-energy structure must not be treated as one unchanged projection symbol. Its lineage passes through a complete-ground projection, selected branch data, realized regional projections, the exact thin complete-ground kernel, invariant-state/GNS data, zero-momentum identification, and finally zero-mass identification only after the relevant terminal theorems authorize those crossings.

Likewise, the coercive constants form a lineage rather than a single number copied through five papers.

These horizontal trajectories illustrate the type of witness lineage studied abstractly in *The Architecture of Long Mathematical Proofs*. The present document records which Yang–Mills manuscript owns each conversion.

## 12.4 Using the two architecture documents together

The distinction remains. *The Architecture of Long Mathematical Proofs* develops general proof-architecture methodology, with Yang–Mills used as a worked case study rather than validated by inclusion. The present document instead provides source-audited Yang–Mills manuscript ownership, theorem concordance, and interface contracts.

### 12.5 Recommended first reading

A useful first route is

$$\text{Sections 1–2} \longrightarrow \text{Sections 8–11} \longrightarrow \text{the manuscript relevant to the reader's question.}$$

Sections 8 and 9 are especially important after the fidelity pass: they show exactly where the Paper-II scale branch ends, where Paper-III incision begins, where Paper-IV hyperbolic ownership stops, and where Paper-V invariant spectral ownership begins.

### 12.6 Navigation cost versus proof cost

This overview reduces navigation cost, not proof cost. The source-native theorem statements and hypotheses remain authoritative. Its task is to let a reader determine, without guessing, what object is crossing an interface, at what certified strength, and which manuscript is responsible for the crossing.

## 13 Current Status and Public Availability

The five technical manuscripts constituting the Yang–Mills reconstruction program were submitted to the *Journal of Mathematical Physics* in August 2026 and are presently under editorial consideration.

While the submissions remain under editorial consideration, the author has elected not to place the complete submitted manuscript set on this website and is awaiting further indication from the journal and AIP Publishing before determining the subsequent public-release path.

The reason for withholding the manuscript files is procedural. The five-manuscript long-form construction has already been written, while its source-audited proof status retains the conditional and branch-specific qualifications stated explicitly in Sections 2 and 11.

### 13.1 The existing manuscript and proof record

The present status should not be interpreted as indicating that the five-paper program is a proposal for future work.

The submitted manuscripts already carry the five principal responsibilities:

$$\begin{aligned}
\text{Paper I} :&\quad \text{finite coercivity},\\
\text{Paper II} :&\quad \text{scale branch selection and higher descent},\\
\text{Paper III} :&\quad \text{derived incision and positive-thickness realization},\\
\text{Paper IV} :&\quad \text{thick-to-thin reconstruction and hyperbolic handoff},\\
\text{Paper V} :&\quad \text{Poincare, no-leakage, and invariant spectral closure}.
\end{aligned}$$

The detailed estimates, constructions, descent arguments, reconstruction theorems, and final spectral analysis reside in those manuscripts.

Their present absence from the website reflects publication status, not the absence of a written technical manuscript chain. The source-audited status of that chain is the one recorded in this document: proved, conditional, and branch-scoped statements are not collapsed into a single undifferentiated claim.

### 13.2 Why this document can be public now

The present note serves a different scholarly function from the five manuscripts.

It exposes the dependency spine, the ownership of the principal theorem obligations, the import/export structure of the five papers, the major cross-cutting mechanisms, the explicit finite witness program, and the claim boundaries of the construction.

It does not reproduce the detailed technical proofs.

The distinction is therefore

$$\text{public explanation of the proof architecture} \neq \text{independent release of the submitted manuscripts.}$$

This allows the mathematical shape of the program to be visible while the manuscript set remains within the present editorial process.

### 13.3 Relation to *The Architecture of Long Mathematical Proofs*

The separate *The Architecture of Long Mathematical Proofs* belongs to a different scholarly layer. It is a general methodological manuscript about proof architecture—typed packets, interface contracts, witness lineages, branch and carrier firewalls, certification systems, and public reconstructibility. It uses a five-stage spectral reconstruction program as a worked architectural case study and explicitly states that this use does not validate the Yang–Mills analytic or physical theorems. The present document has the narrower, source-audited job of navigating this specific five-paper Yang–Mills record.

Its governing subject is the mesoscopic architecture of long distributed arguments: typed stages, interface contracts, witness transport, branch and carrier changes, certification, and public reconstructibility. The Yang–Mills five-stage chain is one worked example of that general methodology, not the object whose mathematical correctness the methodological manuscript claims to certify.

The present document, by contrast, follows the specific source-audited manuscript record

$$\mathrm{I} \to \mathrm{II} \to \mathrm{III} \to \mathrm{IV} \to \mathrm{V}.$$

Thus the three questions are different:

***The Architecture of Long Mathematical Proofs.*** How should long distributed proofs be architecturally audited?

**Present document.** Where, in this Yang–Mills program, is each obligation owned and which source-native theorem anchors the handoff?

**Papers I–V.** What is the detailed source-native verification?

These remain different scholarly objects.

### 13.4 What is not being reproduced here

This overview deliberately does not reproduce the full finite coercive derivations, the complete higher-descent proofs, the detailed carrier and incision constructions, the full positive-thickness realization, the complete thin-reconstruction argument, the complete operator-theoretic stability analysis, or the final detailed spectral proof.

Those belong to the submitted technical record. The present note makes their relationship visible.

### 13.5 Scholarly inquiries

Researchers with a scholarly interest in a particular manuscript, theorem, finite witness, or part of the construction may contact the author for further information.

This is simply a point of scholarly contact during the current editorial interval. It is not intended as a substitute for eventual public access to the manuscripts.

### 13.6 After public release

Once the full manuscripts are made publicly available, the role of this document will remain. Only the publication-status explanation will become historical.

The permanent public navigation can then be organized as

**Start Here**
*Architecture of the Five-Paper Yang–Mills Reconstruction Program*

$$\downarrow$$

Papers I–V and supporting materials.

The separate methodological manuscript *The Architecture of Long Mathematical Proofs* should be linked alongside this route as related proof-architecture methodology, not inserted as a theorem-bearing stage between the present guide and Papers I–V. The present guide tells a reader how this proof record is divided; the technical papers provide the verification; the methodological manuscript explains a general framework for auditing long distributed proofs.

This organization is independent of the temporary publication boundary.

## 14 Conclusion

The five-paper Yang–Mills reconstruction program is long because its terminal invariant statement is not allowed to determine its own ancestry.

At public-overview resolution the spine remains

$$\begin{aligned}\text{finite coercivity} &\longrightarrow \text{coherent branch selection} \longrightarrow \text{positive-thickness realization}\\ &\longrightarrow \text{thick-to-thin reconstruction} \longrightarrow \text{spectral closure.}\end{aligned}$$

The fidelity audit sharpens what lives inside those five labels.

### 14.1 The exact ownership boundaries matter

Paper I supplies the marked finite certificate and the retained coercive provenance.

Paper II organizes the scale direction, selecting a higher pseudosection at the theorem strength actually discharged and fixing the branch-appropriate coefficient-line descent type. It does not construct the spacetime incision.

Paper III derives the regional obstruction carrier and incision, builds the Cauchy-compatible positive-thickness realization, and exports the typed thin-readiness package. It does not construct the thin system.

Paper IV reconstructs the thin observable and represented systems, proves the closed comparison–Mosco form and exact complete-ground kernel, and exports the selected physical hyperbolic/energy certificate. It does not prove invariant mass or physical no-leakage.

Paper V constructs the branch-appropriate Poincare/translation and joint spectral structure, proves the terminal identifications and no-leakage statements, transfers the threshold to invariant mass, and separately closes the declared interaction witness.

These are not editorial differences. They are theorem ownership.

### 14.2 The proof is partly the proof of the interfaces

The long proof therefore depends on four exact contracts:

$$\mathrm{I} \to \mathrm{II}, \qquad \mathrm{II} \to \mathrm{III}, \qquad \mathrm{III} \to \mathrm{IV}, \qquad \mathrm{IV} \to \mathrm{V}.$$

At each contract the reader must be able to identify the carrier, branch, projection/kernel lineage, quantitative constant, and unresolved obligations that cross the boundary.

The architecture becomes false if a fixed-stratum witness is renamed as a full cofinal branch, if scale obstruction data are renamed as a spacetime incision, if a selected-frame energy threshold is renamed as invariant mass, or if comparison/projection stability is renamed as the terminal physical no-leakage theorem.

### 14.3 Typed provenance is the unifying principle

The long construction repeatedly changes mathematical type. The correct standard is therefore not "does the notation look continuous?" but

what object crossed, on what carrier and branch, under which theorem?

The protected sector, the incision, the several coercive constants, and the zero-mode/spectral objects each have a lineage. The endpoint is legitimate only when those lineages are transported rather than merely renamed.

### 14.4 Three complementary levels of exposition

The project consequently has three distinct scholarly layers:

**Papers I–V.** Source-native technical verification.

**This document.** Program-specific manuscript ownership, theorem concordance, and interface navigation.

***The Architecture of Long Mathematical Proofs.*** General methodology for auditing long distributed proofs.

The present note is therefore neither a substitute for the technical manuscripts nor a duplicate of the separate proof-architecture methodology.

### 14.5 One construction

In the most compact language, the program remains

$$\boxed{\textbf{certify} \longrightarrow \textbf{select} \longrightarrow \textbf{realize} \longrightarrow \textbf{reconstruct} \longrightarrow \textbf{close}.}$$

But after the fidelity pass, each verb now has a sharper meaning:

$$\begin{aligned} \text{certify}:&\ \text{marked finite complete-ground/coercive source,}\\ \text{select}:&\ \text{coherent scale branch and higher descent,}\\ \text{realize}:&\ \text{derived incision and positive-thickness regional system,}\\ \text{reconstruct}:&\ \text{thin system and selected hyperbolic energy handoff,}\\ \text{close}:&\ \text{Poincare, no-leakage, invariant mass, and terminal interaction scope.} \end{aligned}$$

That is the five-paper proof record at the resolution this document is meant to preserve.