ANNALS OF THE CHRYSENE FORMALISM

AI-AUGMENTED PEER-REVIEWED | VOL. 1, NO. 4 | ISSN 3144-2021 | OCTOBER 2026

REFEREE REPORT

A Finite Axiomatic Framework for the Interior Chrysene Lattice and Protocell Organization

Paper ID:	ACF-1.4-088
Issue:	Volume 1, Number 4
ISSN:	3144-2021
Review type:	Simulated independent mathematical assessment
Review focus:	Finite axioms, proofs, functional realizations and coordinated models
Review date:	11 October 2026
Publication pages:	1778–1807
Formatting date:	3 November 2026

Manuscript: Schaper, C. D. (2026). A Finite Axiomatic Framework for the Interior Chrysene Lattice and Protocell Organization. *Annals of the Chrysene Formalism*, 1(4), 1778–1807. Article 88. ISSN 3144-2021.

> **Scope and status of the record**
>
> This document preserves the simulated assessment of 11 October 2026. Its original recommendation, source locators and historical page references are retained. The publication formatting and the accompanying bibliography appendix do not constitute a new referee assessment or an editorial acceptance decision.

Recommendation: minor revision. The principal finite mathematical conclusions are supported by the definitions and proofs examined. Two local corrections are necessary to make the typing and material scope consistent. Neither requires withdrawing the finite existence theorem, the spectral and response constructions, the functional-equivalence results, the referenced address quotient, or the stated count and membrane bounds. The requests below are classified so that a response can distinguish a necessary correction from an optional clarification or an extension beyond this paper.

This is a simulated independent referee assessment, not an editorial decision or a claim of review commissioned by a journal. It evaluates the submitted mathematical article on its stated finite, constitutive-model scope. It does not impose laboratory validation as a premise for a formal existence theorem.

1 Version and extent of review

This report is relabelled for Article 88. The assessment, recommendation and review date are unchanged; the historical source paths, page references and hashes below identify the original Article 85 version reviewed, not a new review of the publication layout.

The reviewed version is the complete active closure of `Article85-Axiomatic-Revision/Outputs/Manuscript/main.tex`, including all ten section files, six appendices, macros and references, and its 42-page PDF, `Article85-Axiomatic-Manuscript.pdf`. All proofs in that closure were read. The supplied checker was read and rerun from an unchanged copy outside the original source tree. The 42 rendered pages were also examined for document-level consistency and legibility. The original sources were not edited.

The controlling comparisons were the supplied Article 40 concordance revision, `Pilot-03/Sources/Formalism2 (Article 40)/mainRevConcord.tex`, and the current revised Article 75 and Article 77 projects, `Article75-Revision/Outputs/Revised-Manuscript/main.tex` and `Article77-Revision/Outputs/Revised-Manuscript/main.tex`. The comparison focused on the native signature and finite replacements, the AAA reference action and retained-information counts, and the distinction between material ancestry, current molecular structure and compartment operation. Earlier referee reports were not used in place of reading the manuscript.

For unambiguous identification, the reviewed PDF's SHA-256 is

`2e4c5eac39d5aa2269897ee6659b3a48`
`c411265e3d3ccf3aae5c46bed233ec08.`

Page references below refer to that 42-page version. Stable source labels identify the relevant statements when pagination changes.

2 Mathematical assessment and substantive strengths

2.1 Finite axioms and recognition

The strongest organizing decision is the replacement of the infinite source theory by a separately quantified finite theory. F1–F12 preserve native molecular and interface information while explicitly changing global translation, algebraic, spectral and admissibility clauses. This avoids treating a finite patch as a model of a total nontrivial $\mathbb{Z}^3$ translation action. The statement that these are sufficient defining clauses, without an independence or minimality claim, is mathematically appropriate.

Theorem `thm:recognition` (pp. 4–5) correctly separates three obstructions: nonzero periods, collision of integrated coordinates, and incomplete coordinate-neighbor incidence. Root-path integration proves existence and uniqueness up to translation; injectivity and inducedness require the additional conditions stated. The warning that elementary plaquettes do not test all periods of an arbitrary directional graph is important. The boundary count and the domain-sensitive translation laws are correct, including the distinction between the domains of a composite partial shift and a direct shift.

The 27-cell witness in `thm:finite-existence` (p. 15) genuinely supplies a three-axis example rather than relying on an empty or singleton model. The native radius and diameter argument, ordered displacement invariants $(37, 13, 7)$, exact connector equations, injective cell placement, one graph-interior cell, 36 plaquettes and 8 cubes are consistent. The 54 boundary incidences are correctly kept distinct from 54 internal cell edges and from the 567 intramolecular carbon edges. The one-edge connectors are explicitly abstract structural data, so they do not carry an unsupported covalent-length claim.

2.2 Algebra metric and response

Proposition `prop:finite-generated-algebra` (p. 10) correctly uses localized shifts to obtain edge matrix units and connected paths to obtain all matrix units. Keeping the named generators and regional tensor subalgebras is essential: the bare full matrix algebra would not recover native geometry.

The standard-form argument and the trace-projection construction are sound. On $L^2(M_n, \tau)$ with $D = \lambda P_I$, the commutator seminorm on self-adjoint a is

$$\lambda \|a - \tau(a)I\|_{2,\tau}, \qquad d_D(\varphi_\rho, \varphi_\sigma) = \frac{\sqrt{n}}{\lambda} \|\rho - \sigma\|_{\mathrm{HS}}.$$

The factor $\sqrt{n}$ is correct for ordinary trace-one densities. The scalar-kernel proof applies on the standard-form representation, not on the defining n-dimensional representation. The explicit failure of the first-order condition is properly disclosed rather than hidden under the term "spectral."

The second construction, `prop:finite-generator-dirac` (pp. 11–12), is also valid. The named cell and local Pauli matrices are linearly independent; the isomorphism $M_n \otimes M_n^{\mathrm{op}} \simeq \mathrm{End}(L^2(M_n, \tau))$ then separates the right factors in the commutator. Their common commutant is scalar because they generate M_n. This gives a meaningful connection to the native finite generators while leaving the metric dependent on the supplied coefficients.

The edge-representation proof of weighted shortest-path spectral distance, `thm:graphmetric` (p. 16), is correct, including its singleton convention. The text appropriately keeps that metric distinct from both the matrix-state metric and the admissible-path metric. The conditional expectation and corner inclusion in Section 6.2 have the correct multiplicativity and unit qualifications.

The Schur-complement signs and operator types in F10 and `prop:finite-resolvent` are correct. The finite causal response has the stated upper-half-plane convergence, holomorphy and spectral denominators. In the witness, $\alpha_{xx}(z) = 2/(1 - z^2)$ and $\alpha_{yy}(z) = 4/(4 - z^2)$ give nonzero anisotropy from the supplied Hamiltonian and dipoles. The manuscript does not incorrectly infer those coefficients from port lengths.

2.3 Functional realizations and the operator programme

The definitions of encoded realization and exact quotient simulation retain domain equivalence, which is needed for repeated legal execution. Their composition theorem is valid. The path and four-cycle codes correct one error in each block on the declared domains, and $T_j = E_j u D_j$ satisfies the required state-intertwining equations. The distinction between a material edge and an edge in a supplied response network is explicit.

The path and cycle have equal complete static terminal maps, not merely equal currents at one voltage. The contextual-equivalence proof in `thm:contextual` (p. 25) is the correct elimination argument under its uniqueness hypothesis. It supports passive static replacement without asserting arbitrary dynamical, quantum or molecular equivalence.

The typed operator programme is usefully constructive. Restriction and gluing have the stated structural obstructions; read/write noninterference includes guards; and the split–join proof stays within its declared valence-upper-bound domain. Helical and sheet templates, rotor energy balance, a supplied catalytic network, resolvent perturbation, invariant-sector leakage and the encoded quantum channel are coherent mathematical examples. Their physical interpretations are correctly conditional. They should not be removed simply because the article does not also supply a laboratory mechanism.

2.4 Genetic and membrane coordination

The addressed quotient in `thm:address` (pp. 25–26) agrees with revised Article 75: $F^a R^d u$ is the unique zero-tag representative, the resident AAA type is fixed, and the retained-data counts are 64, 36 and 20. The transported reversal includes the within-pair exchange of C and G. The cited 61-state lower bound is correctly limited to joint preservation of the registry and designated standard output, rather than interpreted as an amino-acid count.

The fiber-product condition and the common projected-vector-field argument are correct sufficient compatibility statements. The model in `thm:common-model` (pp. 28–29) has the stated solution $\xi(t) = \tanh(k_m t)$, constant accessible water volume 200, and total volume 201. The capsule area–volume inequality, radius lower bound and declared scaled-lattice clearance are valid. The comparisons at time one distinguish the two addresses and the removal of the single feedback factor in this particular constitutive model.

The finite-count appendix's scalar theorem also survives scrutiny. The floor remainders give $p = Y + e$ with $|e| < 1/100$; the improved rate bound yields $B \leq 1000$; the concave lower comparison function gives the displayed hitting time; and the water barriers and energy reserve have the stated signs. A deactivation clock cannot advance while its active count is zero. Production hits have positive rate, and a deactivation hit has an active chain immediately before its reset. With the prescribed simultaneous-event order, $D \leq B \leq 1000$ rules out infinitely many such events. These facts support uniqueness and continuation of the declared count execution. The material interpretation needs the local scope correction R2 below.

3 Mandatory corrections

R1 Separate state preparation from operation dispatch

Location: `sections/genetics.tex`, `thm:address` and "An explicit dispatch and its retained information" (p. 26); its use in `sections/membrane.tex`, `thm:common-model`.

The theorem types $h : M \to U$ as a selector of an input operation of a functional system. The example then writes $h(n) = (m, r) = (1 - n_6, 1)$, whose displayed value is a state in $X = \{0, 1\}^2$, and calls this a preparation. The prose recognizes that preparation and the controlled update u are different, but the maps and their codomains are not separated. In a paper centered on typed operations, this is a local formal gap.

Minimal repair: use a distinct state initializer

$$h_{\mathrm{init}} : \{100110, 100111\} \longrightarrow X, \qquad h_{\mathrm{init}}(n) = (1 - n_6, 1).$$

If this is to instantiate the theorem's operation-valued dispatch, introduce preparation symbols $p_x \in U$ with total maps $f_{p_x}(z) = x$, and set $h_{\mathrm{op}}(n) = p_{h_{\mathrm{init}}(n)}$ on the selected address domain. On encoded component states use $T_{j,p_x}(y) = E_j(x)$. State explicitly whether other addresses are outside the dispatch domain or have a specified default operation. Either this construction or an equally explicit separation resolves the problem. The normal-form theorem, the two prepared values and the membrane trajectories need no numerical change.

R2 Keep the count specialization distinct from an atom-mapped realization

Location: `appendices/closed_loop.tex`, opening paragraph, `app:closed-loop` (pp. 38–39), read with `appendices/materials.tex`, `prop:splitjoin` and the productive stoichiometry (p. 38).

The count appendix says its preparation is represented by the "atom-preserving block" of the material appendix. The latter actually provides an element-balanced reaction formula and selected carbon-allocation descriptions, then explicitly says that applying the split–join proposition requires complete atom maps and admissible endpoint bond graphs. The present closure therefore has not established that stronger block description unconditionally. An element-count equality is insufficient to specify every current atom identity, bond graph and permitted intermediate state.

Minimal repair: describe the displayed block as an element-balanced or count-level specialization. Keep the exact carbon ledger, all stock/energy bounds, and the count execution. State separately that a complete atom-resolved realization is obtained only after supplying compatible labelled endpoint graphs, complete atom assignments, the permitted editing domain and the required finite resources. The general split–join proposition remains a valid sufficient construction once its hypotheses are supplied. A full new molecular certificate is not required for the stated count-level theorem; a conditional material lift is enough.

4 Useful clarifications

These suggestions improve the presentation but do not identify counterexamples to the central results.

C1 Use one unambiguous term for graph interior. In `def:boundary` (p. 5), "a finite interior region" is defined by $\Lambda^\circ \neq \varnothing$, whereas the introduction and F6 allow a finite ICL to have empty graph interior because "interior" also refers to an enclosure. Replace the former phrase by "a region with nonempty graph interior." Keep the broader finite ICL class and the stronger box example.

C2 Complete the observation fields of the component example. At `thm:components` in `sections/components.tex` (p. 24), state the chosen observation set O and maps o and o_j. For example, let $o(m, r)$ record the fixed ordered-port reduct together with $(1 + m, 1 + r)$ and put $o_j = o \circ D_j$, using the already described offset-corrected geometric observation. Then $o_j E_j = o$ is immediate, and every field of `def:realization` is visibly instantiated. This does not change the coding or terminal-response calculation.

C3 Display a concrete shared-state witness briefly. For `thm:common-model`, one compact tuple listing the fixed lattice/carrier/component data and the coordinates (m, ξ, A, v), together with the subsystem read maps, would make the nonempty fiber product directly inspectable. The existing proof's common-coordinate explanation is sufficient in substance, but an explicit tuple would prevent readers from mistaking arbitrary subsystem products for a proved compatible realization. No smoothness assertion about a general fiber product should be added.

C4 Identify the contribution and certificate coverage precisely. The contribution is the typed finite framework, its simultaneous nontrivial witnesses, and the coordinated use of these structures. Path integration, Schur complements, matrix standard form, elementary error correction and Lipschitz error propagation are established mathematical techniques, not new general theorems simply because they appear here. A short statement of this distinction would strengthen the introduction. In reporting the checker, retain the present separation between enumerated instance checks and general proofs. The check count should not be described as the number of independently verified theorems.

5 Optional future extensions

The following are possible future directions, not requested revisions or conditions for acceptance. They require no additional action for this revision.

S1 Physical implementation. It would be inappropriate to make acceptance of the finite existence and compatibility theorems conditional on synthesis, membrane stability experiments, measured optical response, catalytic turnover, coherent preparation or a selected peptide fold. Those would test the additional common molecular realization, which the article expressly distinguishes from the finite model. Retain that distinction; do not replace the positive formal constructions with a demand for an empirical research programme.

S2 Stronger foundational classification. Logical independence of F1–F12, a minimal axiom list, a classification of every model, a universal uniqueness result, an infinite-lattice extension, or a canonical identification of the different metrics would be new mathematical projects. The manuscript neither needs nor currently proves them. In particular, the full-matrix result does not classify the enriched finite structures up to isomorphism.

S3 Complete biological control and mechanisms. The two-address control example need not implement complete genetic translation or a unique origin of the biological code. Nor must the helical, sheet, hinge, pocket or quantum-channel examples become a single demonstrated molecular machine. Such claims would require additional typed data and new results. These extensions may be left for future work while preserving the supported finite conclusions.

6 Reproducibility findings and conclusion

The unchanged checker, copied as `verification/verify_reviewed_manuscript.py`, completed successfully with **2,593 exact checks**. Its machine-readable output is `verification/verification_results.json`. The checks include native coordinate and incidence identities, box counts and paths, representative rational matrix identities, complete declared code-error domains, the addressed quotient counts, and exact rational bounds for the two operating examples. Several thousand entries arise from finite enumerations of related instances; they are not independent proof obligations of equal scope.

The tests corroborate the manuscript's explicit calculations. They do not construct the enormous full tensor matrix, prove a universal theorem by sampling small dimensions, certify a chemical preparation, or establish an infinite model. The general arguments for recognition, generation, standard form, separation, composition, contextual replacement and compatibility were assessed from their proofs. No failure of those arguments was found within their stated hypotheses.

The revised submission should answer R1 and R2 explicitly, record its choices on C1–C4, and leave the optional S1–S3 extensions outside its claims unless separately undertaken. Conclusions that already follow mathematically should be retained. With the two local corrections, the article provides a coherent constructive account of finite ICL structures and their specified functional realizations. My recommendation is therefore **minor revision**, with the supported principal conclusions retained.