ANNALS OF THE CHRYSENE FORMALISM

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

ORTHOGONAL REVIEW RECORD

Axiom-Free Invariants of the Cut-Set Lattice in the Chrysene Formalism: Admissible Maximal Chains, Obstruction Functionals, and Disjoint Optima on Hamiltonian Paths

Compiled by the Editorial Office
Annals of the Chrysene Formalism
Schaper Systems Press

Paper ID:	ACF-1.4-087
Issue:	Volume 1, Number 4
ISSN:	3144-2021
Review architecture:	Dual-agent orthogonal review by independent agents
Review focus:	Combinatorial foundations, redox accounting, independent verification of cut-set claims
Acceptance date:	2026-11-03
Online publication:	2026-11-03
Pages:	1761–1777
Document version:	v1.0
Last audited:	3 November 2026

Citation: Schaper, C. D. (2026). Axiom-Free Invariants of the Cut-Set Lattice in the Chrysene Formalism: Admissible Maximal Chains, Obstruction Functionals, and Disjoint Optima on Hamiltonian Paths. *Annals of the Chrysene Formalism*, 1(4), 1761–1777. ISSN 3144-2021. `https://chrysene.com/index.php/acf/article/view/87` • © 2026 Charles D. Schaper. Distributed under the CC Attribution-NonCommercial-NoDerivatives 4.0.

Scope of Review

The manuscript develops the discrete graph-theoretic, order-theoretic, and redox-thermodynamic foundations governing the ring-opening cleavage of the tetracyclic chrysene skeleton into unbranched C_{18} membrane lipid chains. It characterizes the cleavage poset on Hamiltonian-path cut sets, proves disjoint optima between the 1-yl terminal pattern and the activation minimum, establishes isomer invariance across tetracyclic polyhexes, and defines an oxidation-number state function for path-independent redox demand bounds. Concordance with upstream articles on endogenous carbon purity and the Article 83 electron ledger is examined.

Compiled disposition

Major revision. Dual-agent orthogonal review by independent agents. Referee A (mathematical, chemical, and concordance assessment) recommended accept with minor revisions, identifying only two editorial points. Referee B (independent recomputation and verification against the paper's own script and conventions) recommended major revision, with five major findings (Section 11 state vectors realized only on excluded cut sets; Strategy II not redox-neutral; inverted scholium; contradictory cleavage models; Proposition 7.6 void on the domain of Theorem 7.1) and numerous moderate and minor points. The combinatorial core of Sections 3–10 is sound and was independently reproduced. Section 11 and several supporting claims require rebuilding or withdrawal. The principal results (Theorems 6.1, 7.1, 8.5, 9.5) survive.

Contents

1 Stage summaries **3**

2 Full dual-agent review history (verbatim) **4**

1 Stage summaries

Dual-agent architecture

Referee A confined to mathematical assessment, chemical interpretation, and concordance with the surrounding corpus. Referee B confined to independent recomputation of every integer from the molecular skeleton, verification of the paper's script assertions against its stated conventions, and detection of inconsistencies between sections. Neither saw the other's report. The two reports diverge sharply: one accepts the paper's numerical claims and redox ledger; the other recomputes and finds Section 11's hard-coded vectors realized exclusively on the three cut sets excluded by Theorem 6.1, and finds the paper's sole claimed proof of its central activation bound void on the relevant domain.

Mathematical, chemical, and concordance findings (Referee A)

Referee A found the manuscript mathematically sharp, chemically grounded, and tightly integrated. The automorphism action, cleavage poset, disjoint optima theorem, isomer invariance, and oxidation-number state function are endorsed. Only two minor editorial revisions are requested (notation collision in a summary banner; Roman numeral formatting). Recommendation: accept with minor revisions.

Independent verification findings (Referee B)

Referee B independently reimplemented the definitions and reproduced every combinatorial integer in Sections 3–10. Theorem 8.5 stands. Section 11 does not: its oxidation-state vectors are realized only by the three all-fusion-fusion cut sets that Theorem 6.1 excludes; Strategy II is not redox-neutral on the clean cut sets; the scholium inverts the model; Definitions 4.6 and 11.4 are contradictory; Proposition 7.6 supplies no bound on the domain of Theorem 7.1. Additional moderate findings concern symmetry orbits, the isomer count (triphenylene), hollow assertions, and unverified claims. Recommendation: major revision. The combinatorial conclusions survive; two supporting sections do not.

2 Full dual-agent review history (verbatim)

The following is the complete text of the dual-agent orthogonal review by independent agents. Referee A is the mathematical/chemical/concordance assessment. Referee B is the independent recomputation and verification report.

Editorial note on the process

This article was reviewed under the journal's *AI-augmented peer review by orthogonal constraint*. Two referees read the submitted manuscript independently and without sight of each other's reports. Referee A was confined to mathematical assessment, chemical interpretation, and concordance. Referee B was instructed to recompute every numerical claim from the skeleton independently of the paper's verification script, to check assertions against the paper's own conventions, and to find defects. Both reports are reproduced in full below.

Referee A — mathematical, chemical, and concordance assessment

REFEREE'S REPORT

Annals of the Chrysene Formalism

Manuscript ID: ACF-2026-087 (Article 87)
Subject: *Symmetry, Cleavage Posets, Disjoint Hamiltonian Path Optima, Isomer Invariance, and Oxidation-Number State Functions for Chrysene-to-Lipid Ring Opening*
Recommendation: Accept with Minor Revisions

---

1. Executive Summary and Overall Assessment

This manuscript (Article 87) develops the discrete graph-theoretic, order-theoretic, and redox-thermodynamic foundations governing the ring-opening cleavage of the tetracyclic chrysene skeleton into unbranched C_{18} membrane lipid chains. The work is organized into three tightly integrated parts:

(i) **Symmetry, the Cleavage Poset $\mathcal{P}(K)$, and the Activation Bound (Section I; Propositions 3.4, 3.7, 5.1, 6.3, 7.6; Corollary 3.5; Theorems 6.1, 7.1):** Characterizes the C_{2h} graph automorphism action ($|\mathrm{Aut}(X)| = 2$, $\sigma(\Sigma) = \Sigma$) on the $|\mathcal{K}| = 27$ ring-opening cut sets (15 orbits) that reduce the chrysene graph (1,189 spanning trees across 252 skeletons) to a single unbranched C_{18} chain. Theorem 6.1 proves that a cut set admits no clean strict cleavage ordering if and only if it is contained entirely within the 5-edge ring-fusion set $\mathcal{F}$ ($\Pi_s^{\mathrm{cl}}(K) = \emptyset \iff K \subseteq \mathcal{F}$, $|\mathcal{F}| = 5$), leaving 24/27 clean cut sets and 54 strict (110 weak) clean orderings out of 27×24 (Proposition 6.3). Theorem 7.1 establishes the activation floor

$\min_K \alpha_{\min}(K) = 2 \leq \alpha_{\min}(K) \leq 3$ (split 12 and 12 across the 24 clean cut sets), bounded above by the poset parent count $\alpha(\pi) \leq \mathrm{pc}(K)$ (Proposition 7.6).

(ii) **Disjoint Optima on Hamiltonian Paths and Isomer Invariance (Section II; Proposition 8.2; Theorems 8.5, 9.5; Corollary 8.6; Proposition 10.1; Conjecture 10.2):** Classifies the 7 substituent/terminal patterns $p(K)$ on the resulting C_{18} chain and proves that the unique 1-yl terminal pattern $p^* = (1, 10)$ (realized by 6 cut sets in 3 orbits) is disjoint from the unconstrained activation minimum: $p(K) = p^* \Rightarrow \alpha_{\min}(K) = 3 > \min_{K \in \mathcal{K}} \alpha_{\min}(K) = 2$, and $\alpha_{\min}(K) = 2 \Rightarrow p(K) \neq p^*$. Thus, requiring a 1-yl terminal lipid (H_{term}) exacts a structural price of +1 activation step, proving that the recorded 3-activation route is strictly optimal within H_{term} (Theorem 8.5, Corollary 8.6). Extending the analysis across all 5 four-ring isomers (4 cata-condensed chains plus triphenylene) and the k-ring homologous series ($2 \leq k \leq 4$, 8 skeletons), Theorem 9.5 and Proposition 10.1 prove that all 5 tetracyclic isomers have 24 clean cut sets with $K \subseteq \mathcal{F} \iff \Pi_s^{\mathrm{cl}} = \emptyset$, that chrysene uniquely minimizes strict clean orderings (54 vs. 60–62), and that $\min_K \alpha_{\min}(K) = \min_K |K \cap \mathcal{F}| = k - 2$ holds across all 8 skeletons.

(iii) **Oxidation-Number State Function and Redox Demand Bounds (Section III; Definitions 11.1, 11.4; Propositions 11.2, 11.10; Theorem 11.5; Corollary 11.7):** Defines the vertex-summed carbon oxidation number $ON(S) = \sum_{v \in V}(o_S(v) - h_S(v))$ and proves that it is a path-independent state function (Proposition 11.2) whose endpoint difference $ON(S') - ON(S)$ fixes the net redox demand independently of reaction mechanism, cut set K, or cleavage ordering π (Theorem 11.5). Moving from $ON_{\mathrm{core}} = -8$ to $ON_{\mathrm{ether}} = -34$ or $ON_{\mathrm{ester}} = -30$ establishes an exact redox floor of 26 e^-/unit (52 e^-/dimer) for the ether lipid and 22 e^-/unit (44 e^-/dimer) for the ester lipid, independently confirming Article 83's electron ledger (Corollary 11.7). Finally, Proposition 11.10 compares total electron traffic between oxidative cleavage (Route I, which climbs +14 or +16 units before descending, traffic(I) $\in \{54, 58\}$) and direct reductive cleavage (Route II, which descends monotonically, traffic(II) = 26, a 2.08–2.23× reduction in redox traffic).

The manuscript is mathematically sharp, chemically grounded, and tightly integrated with the surrounding protocell and membrane papers. Below, I provide detailed evaluations on **Mathematics**, **Chemistry**, and **Concordance**, followed by minor points for revision.

---

2. Mathematical Assessment

(i) **Automorphism Group Action and Orbit Counting on the Cut-Set Space $\mathcal{K}$ (Propositions 3.4, 3.7, 5.1; Corollary 3.5):** The reduction of the spanning-tree enumeration on the chrysene molecular graph X (1,189 spanning trees partitioned into 252 unlabelled/isomorphism skeletons, Proposition 5.1) to the 27-element space $\mathcal{K}$ of Hamiltonian-path cut sets (1 unbranched C_{18} chain per cut set $K \in \mathcal{K}$) is handled with precision. Because $|\mathrm{Aut}(X)| = 2$ and the inversion generator σ preserves the substituent/structural locus Σ setwise ($\sigma(\Sigma) = \Sigma$), Burnside's lemma / orbit decomposition on $|\mathcal{K}| = 27$ yields 15 orbits (3 σ-fixed cut sets and 12 size-2 orbits, $3 + 2 \times 12 = 27$). Corollary 3.5 immediately follows from this action, identifying the

two symmetric K-region bonds 11–12 and 5–6 as a single $\mathrm{Aut}(X)$-orbit.

(ii) **Order-Theoretic Characterization of Clean Cut Sets via the Cleavage Poset $\mathcal{P}(K)$ (Theorem 6.1; Propositions 6.3, 7.6; Theorem 7.1):** For each 4-edge cut set $K \in \mathcal{K}$ ($|K!| = 24$ permutations, giving $27 \times 24 = 648$ total ordered pairs (K, π)), the cleavage poset $\mathcal{P}(K)$ captures the partial-order precedence constraints required for clean sequential ring opening. Two structural results stand out:

- *Exact fusion-set obstruction (Theorem 6.1):* The set of strict clean linear extensions $\Pi_s^{\mathrm{cl}}(K)$ is empty if and only if every remaining cut edge after the initial K-region cut lies in the 5-element ring-fusion edge set $\mathcal{F}$ ($\Pi_s^{\mathrm{cl}}(K) = \emptyset \iff K \subseteq \mathcal{F}$, $|\mathcal{F}| = 5$). Since choosing 3 fusion edges from the 3 remaining rings gives exactly 3 such cut sets, $27 - 3 = 24$ cut sets admit clean orderings, yielding 54 strict and 110 weak clean orderings (Proposition 6.3).
- *Activation bound and parent-count inequality (Theorem 7.1, Proposition 7.6):* Across the 24 clean cut sets, the minimum activation functional $\alpha_{\min}(K) = \min_\pi \alpha(\pi)$ takes values exclusively in $\{2, 3\}$, partitioning the 24 clean cut sets equally into 12 cut sets (7 orbits) at $\alpha_{\min}(K) = 2$ and 12 cut sets at $\alpha_{\min}(K) = 3$. Proposition 7.6 relates $\alpha(\pi)$ to the poset parent count $\mathrm{pc}(K)$ via $\alpha(\pi) \leq \mathrm{pc}(K)$, showing that the inequality is strict on 12 cut sets.

(iii) **Disjoint Optima Theorem on Hamiltonian Paths and Polyhex Invariance (Proposition 8.2; Theorems 8.5, 9.5; Proposition 10.1; Conjecture 10.2):** Theorem 8.5 is a clean multi-objective optimization result on a finite combinatorial space: on the 24 clean cut sets, the terminal-position map $p : \mathcal{K} \to \{7 \text{ patterns}\}$ and the activation functional $\alpha_{\min} : \mathcal{K} \to \{2, 3\}$ have strictly disjoint level sets for $p^* = (1, 10)$ and $\alpha_{\min} = 2$:

$$\{K \in \mathcal{K} : p(K) = (1, 10)\} \cap \{K \in \mathcal{K} : \alpha_{\min}(K) = 2\} = \emptyset.$$

Specifically, all 6 cut sets (3 orbits) realizing $p^* = (1, 10)$ have $\alpha_{\min}(K) = 3$, proving that no 2-activation route can produce a 1-yl terminal chain. Even more striking is the topological universality established in Section II: across all 5 four-ring polyhex isomers (4 chains plus triphenylene), the clean-cut count (24) and the characterization $K \subseteq \mathcal{F} \iff \Pi_s^{\mathrm{cl}}(K) = \emptyset$ are invariant, while chrysene is the unique extremum minimizing the number of strict clean orderings (54 vs. 60–62, Theorem 9.5). Furthermore, across all 8 polyhex skeletons for $2 \leq k \leq 4$, the minimum activation bound coincides with the minimum fusion-edge intersection number, $\min_K \alpha_{\min}(K) = \min_K |K \cap \mathcal{F}| = k - 2$ (Proposition 10.1), motivating Conjecture 10.2 for general k.

(iv) **Exact Potential Theory on the Reaction Graph via the Oxidation-Number State Function (Proposition 11.2; Theorem 11.5; Proposition 11.10):** Defining $ON(S) = \sum_{v \in V}(o_S(v) - h_S(v))$ converts multi-step organic redox accounting into an exact discrete potential (zero-form) on the state space (Proposition 11.2). Because the telescopic sum $\sum_i \Delta ON_i = ON(S_{\mathrm{final}}) - ON(S_{\mathrm{init}})$ depends only on the boundary states ($ON_{\mathrm{core}} = -8$, $ON_{\mathrm{ester}} = -30$, $ON_{\mathrm{ether}} = -34$), Theorem 11.5 establishes path-independent lower bounds on electron consumption ($\Delta ON = -26$ for ether, -22 for ester) while Proposition 11.10 uses the L^1 total variation $\mathrm{traffic}(\gamma) = \sum_i |\Delta ON_i|$ along a path γ to quantify the non-monotonicity penalty of oxidative cleavage ($\mathrm{traffic(I)} \in \{54, 58\}$ vs. $\mathrm{traffic(II)} = |\Delta ON| = 26$).

3. Chemical Assessment

(i) **K-Region Initiation, Ring-Fusion Bonds $\mathcal{F}$, and Unbranched C_{18} Chain Generation (Corollary 3.5; Proposition 5.1; Theorem 6.1):** Chemically, opening a tetracyclic polyaromatic hydrocarbon (C_{18}) into a single linear octadecyl/octadecanoyl lipid chain requires cleaving exactly one bond per ring (4 bonds in total) without severing the carbon skeleton into disconnected fragments (so $X \setminus K$ must be a Hamiltonian path, 1 of the 1,189 spanning trees of chrysene, Proposition 5.1). Initiating cleavage at the most reactive K-region double bond (5–6 or its symmetry-equivalent 11–12, Corollary 3.5) leaves 3 further ring-opening cuts. Theorem 6.1 supplies the chemical rationale for why 3 of the 27 Hamiltonian-path cut sets fail to admit a clean sequential cleavage ordering: when all cut edges belong to the shared ring-fusion bond set $\mathcal{F}$ ($|\mathcal{F}| = 5$), no peripheral bond cleavage can isolate a target bond without prior activation interference.

(ii) **Chemical Meaning of the $(1, 10)$ Terminal Pattern and the $+1$ Activation Price of 1-yl Lipids (Theorem 8.5; Corollary 8.6):** For the resulting C_{18} chain to function as an amphiphilic membrane lipid with its polar headgroup at a terminal carbon (C1, yielding a 1-yl chain rather than a branched/internal secondary attachment), the substituent locants carried from the parent chrysene core must land at positions $(1, 10)$ along the C_{18} Hamiltonian path ($p^* = (1, 10)$, Proposition 8.2). Theorem 8.5 proves that this regioselectivity requirement (H_{term}) is incompatible with the 2-activation pathways ($\alpha_{\min}(K) = 2 \Rightarrow p(K) \neq (1, 10)$), forcing any synthesis of a 1-yl terminal lipid to traverse at least 3 activation steps ($\alpha_{\min}(K) = 3$). This proves that the 3-activation synthetic route is not an artifact of sub-optimal design, but is the exact chemical minimum dictated by the topology of chrysene.

(iii) **Stoichiometric Redox Balances and Oxidative vs. Reductive Ring-Opening Energetics (Theorem 11.5; Corollary 11.7; Proposition 11.10):** In formal carbon oxidation-state accounting, each C–O bond contributes $+1$, each C–H bond contributes -1, and each C–C bond contributes 0, giving $ON(S) = \sum_{v \in V}(o_S(v) - h_S(v))$ (Definition 11.1).

- For the parent substituted chrysene core, $ON_{\text{core}} = -8$.
- Full ring opening and reduction to the saturated ether lipid ($\text{ON}_{\text{ether}} = -34$) requires a net 26-electron reduction per chrysene unit (52 e^- per two-chain dimer lipid), whereas stopping at the ester oxidation level ($ON_{\text{ester}} = -30$) requires 22 e^- per unit (44 e^- per dimer, Theorem 11.5, Corollary 11.7).
- Proposition 11.10 highlights the severe redox-efficiency penalty of classic oxidative ozonolysis/dioxygenase-style ring cleavage (Route I): oxidizing the aromatic rings up to dicarboxylic/carbonyl intermediates climbs $+14$ or $+16$ oxidation units before reducing all the way down to -34, requiring 54 or 58 total electron transfers (2.08–2.23× the minimal 26-e^- reductive Route II).

4. Concordance Assessment

(i) **Endogenous Chrysene Carbon Purity of the Membrane Lipid Chain (Proposition 5.1; Section II):** Concordance within the Chrysene Formalism requires that major structural biomolecules—including the protocell membrane lipids of Articles 62, 65, 72, 76, 77, and 83—derive their carbon skeleton endogenously from the 18-carbon chrysene ring system without importing exogenous fatty acids. Proposition 5.1 and Section II verify that the C_{18} lipid chain is an exact Hamiltonian path of the parent chrysene core (18 core carbons preserved 1–1 without carbon loss or exogenous carbon addition).

(ii) **Optimality of the Recorded Corpus Route and Exact Confirmation of Article 83's Redox Ledger (Theorem 8.5; Corollary 11.7):** Two results directly validate upstream articles in the corpus:

- *Validation of the recorded 3-activation cleavage pathway (Theorem 8.5):* By proving that $\alpha_{\min}(K) = 3$ is the strict lower bound across all cut sets yielding the terminal pattern $p^* = (1, 10)$, Theorem 8.5 certifies that the corpus's recorded 1-yl lipid cleavage sequence cannot be shortened by any alternative choice of cut set $K \in \mathcal{K}$ or ordering π.
- *Confirmation of Article 83's 44-electron / 52-electron dimer ledger (Theorem 11.5, Corollary 11.7):* By establishing from endpoint state functions alone that the ester lipid requires 22 e^-/unit (44 e^-/dimer) and the ether lipid requires 26 e^-/unit (52 e^-/dimer), Corollary 11.7 independently confirms the redox ledger of Article 83 (`A83`) regardless of intermediate mechanism.

(iii) **Why Chrysene Rather Than Another Tetracyclic Isomer (Theorem 9.5):** A recurring concordance question across the corpus is why chrysene is privileged among the five tetracyclic benzenoid isomers (tetracene, benz[a]anthracene, chrysene, benzo[c]phenanthrene, and triphenylene). Theorem 9.5 provides a quantitative combinatorial answer: while all 5 isomers share 24 clean cut sets and satisfy $K \subseteq \mathcal{F} \iff \Pi_s^{\mathrm{cl}}(K) = \emptyset$, chrysene uniquely minimizes the number of strict clean cleavage orderings (54 vs. 60–62), providing the tightest kinetic/combinatorial funnel among all tetracyclic polyhexes.

---

5. Minor Technical and Editorial Revisions

m1. Clarifying the Set-Cardinality Notation in the Section I Summary Banner ($\Pi_s^{\mathrm{cl}}(K) = \emptyset \iff K \subseteq \mathcal{F}$): In the Section I summary line for Theorem 6.1, the parenthetical condition is printed with a collided symbol as ($|\mathcal{F}| = 52|C| = 3$, Thm. 6.1) (where $|\mathcal{F}| = 5$ and the 3-element subset/cut count 3 ran together). Separating $|\mathcal{F}| = 5$ and the excluded cut-set count $|\{K \in \mathcal{K} : K \subseteq \mathcal{F}\}| = 3$ with a comma or semicolon will make the display immediately clear.

m2. Formatting of Route Identifier `I` vs. Lowercase `l` in Section III (`traffic(I)`): In the Section III summary line comparing the electron traffic of oxidative Route I and reductive Route II, the Roman numeral `I` in `traffic(I)` is typeset as a lowercase italic `l` ($traffic(l) \in \{54, 58\}$ vs. $traffic(II) = 26$). Standardizing to

Roman numeral `\mathrm{I}` (traffic(I) $\in \{54, 58\}$ vs. traffic(II) $= 26$) will match "Oxidative route I" in the line immediately below.

Report prepared for the Editorial Board, Annals of the Chrysene Formalism.

Reviewer disclosure
Referee B — independent recomputation and verification

This review was produced by a model instance that did not see the paper being drafted, working from the compiled PDF and the sources, and instructed to find defects. It is independent of the drafting in that sense and not independent of the author-side model in general. Its findings were then re-derived independently before being recorded here; every numerical claim below was recomputed from the skeleton rather than taken from the review or from the paper. Two of the review's claims did not survive that re-derivation and are noted where they occur.

The paper's own convention is that every number it prints is emitted by its verification script. That convention is what makes a review of this kind possible, and it is also what raises the severity of several findings below: where prose and script disagree, the paper's stated standard treats the disagreement as a defect of the first order rather than a matter of presentation.

Recommendation

Major revision.

The combinatorial core, sections 3 to 10, is sound. Every integer in it was reproduced from an independent implementation of the definitions: the order and size of the skeleton, the automorphism group, the bond classes, the chord containment, 1189 spanning trees, 27 chain-forming cut sets, 24 admitting a clean ordering, 54 strictly clean orderings, the activation distribution {2:12, 3:12}, the seven substituent patterns, the uniqueness of the terminal pattern, and the incompatibility theorem. I could not break any of them. The main result, Theorem 8.5, stands as stated and is the paper's real contribution.

Section 11 does not stand. Its four hard-coded oxidation-state vectors are realised by exactly three cut sets, and those three are precisely the ones Theorem 6.1 proves admit no clean ordering. Every figure in Proposition 11.6, Corollary 11.7, Table 5 and Ruling 13.2 is computed on a molecule the rest of the paper excludes, the "redox-neutral" result is an artefact of an undercount, and the closing scholium asserts the opposite of what its own model implies. The section must be rebuilt from a named cut set or withdrawn.

Separately, the paper's single claim to have proved anything about its central bound — Proposition 7.6, cited three times as giving 1 where the enumeration gives 2 — is void on the domain of Theorem 7.1, and the claim must be withdrawn.

Neither failure touches Theorems 6.1, 7.1, 8.5 or 9.3. The paper's conclusions survive; two of its supporting sections do not.

Principal findings

R1. MAJOR **Section 11's state vectors are realised only by the three cut sets Theorem 6.1 excludes.** Block W10 hard-codes `SUB_CORE`, `STRAT_I`, `STRAT_II_PRE` and `STRAT_II` as per-carbon lists. `STRAT_I` installs a carbonyl on six carbons. Four cleavages cap eight carbon ends. Over the 27 cut sets the number of *distinct* endpoint carbons is 8 for twelve cut sets, 7 for twelve, and 6 for three; only in the last case does the endpoint set coincide with the six ring-fusion carbons, and those three cut sets are exactly the three with four fusion–fusion bonds, which Theorem 6.1(i) proves admit no clean ordering.

Recomputing ON per cut set, with each endpoint carbon taking a carbonyl:

$$\mathrm{ON}(\mathsf{I}) = \begin{cases} +4 & \text{3 cut sets (all-ff; no clean ordering)} \\ +6 & \text{12 cut sets} \\ +8 & \text{12 cut sets} \end{cases}$$

so the oxidative demand is 14 or 16 equivalents, not 12, for every cut set the paper is actually about. The number of oxygen atoms installed is 8 invariantly, which is what Definition 11.4's own prose says; the script installs 6.

Correction to the review. The review states the vectors "correspond to no cut set". They correspond to three. That is worse, not better: it localises the section on precisely the cut sets section 6 removes from consideration.

Fix. Build the state vectors in code from a named cut set drawn from `CUTS`, emit ON per cut set or per orbit, and assert against W4.

R2. MAJOR **Strategy II is not redox-neutral, and its model describes three cleavages rather than four.** `STRAT_II_PRE` leaves only the six fusion carbons without hydrogen, so it presents three isolated alkenes, not four; and an sp^2 ring carbon cannot carry two double bonds. Per cut set:

$$\mathrm{ON}(\mathsf{II}) = \begin{cases} -8 & \text{3 cut sets (all-ff)} \\ -5 & \text{12 cut sets} \\ -2 & \text{12 cut sets} \end{cases}$$

against ON(core) $= -8$. So strategy II is a net oxidation of 3 or 6 equivalents for every cleanly linearizable cut set, and is redox-neutral only on the three the paper excludes. The precursor value is -19 or -18, not -20; hydrogens added are 11 or 10, not 12; the per-dimer demand is 22 or 20, not 24; and the ledger fractions change accordingly.

Correction to the review. The review gives the corrected figures as constant across the twelve cut sets whose four bonds are disjoint. They are not constant over the twenty-four: twelve cut sets have eight distinct endpoint carbons and twelve have seven, giving two values at each stage, as above. The section's dependence on the cut set is therefore real and cannot be absorbed into a single corrected number.

The assertion `ONII == ONSUB` passes because twelve added hydrogens happen to equal the twelve C–O bonds of six carbonyls. It is an arithmetic coincidence of the undercount in R1, and it is the kind of assertion the corpus's own conventions are

meant to prevent: it checks a hand-written constant against another hand-written constant.

Consequently Proposition 11.6(iii), Corollary 11.7, the section 12 register entry asserting redox neutrality, and Remark 11.5's warning that a cheaper route "is claiming an arithmetic error" are all false. The last is unfortunate.

Corollary 11.7 additionally says the strategy "introduces no new reductant requirement" in the same sentence as "consumes 46.2 per cent of the ether budget". Consuming half a budget is a requirement. One of the two clauses must go.

R3. MAJOR **The section 11 scholium and Remark 11.9 invert their own model.** Both assert that strategy II resolves the activation obstruction because a carbon made sp^3 bears hydrogen. In `STRAT_II_PRE` the carbons left without hydrogen are exactly the endpoints of the bonds to be cut, and every carbon that gains a hydrogen is one whose bond is *not* being cut — necessarily, since leaving the target carbons sp^2 is what presents the target bond as an alkene. By Definition 4.6 the index counts steps at which both carbons of the bond being cut carry no hydrogen, so on this model a strategy-II route has index equal to its number of cuts, against a floor of 2. The strategy maximises the obstruction rather than resolving it.

Ruling 13.2 withdraws an interpretation of Article 83 section 13 on the strength of this claim. A precedence ruling therefore rests on an inverted inference and must be withdrawn with it.

R4. MAJOR **Definitions 4.6 and 11.4 use contradictory models of what a cleavage does.** Definition 4.6 says a cleavage caps both carbons with hydrogen, and that convention is the entire reason the activation index depends on the ordering rather than only on the cut set — it is the content of Remark 7.5, Remark 7.7 and Target 14.1. Definition 11.4 says a cleavage installs a carbonyl at both carbons and changes no hydrogen count. Both cannot hold. The paper never states capping as a hypothesis and never reconciles the two sections.

The difference is not cosmetic. Of the 24 cleanly linearizable cut sets, twelve have activation index strictly below the number of their bonds that are hydrogen-free in the parent, which is exactly the capping effect; under the section 11 convention no cleavage ever donates a hydrogen, the index collapses to a function of the cut set alone, and Definition 4.6's stated purpose disappears.

It is also wrong at the place that matters most. All six target-pattern cut sets cut at a substituted carbon — necessarily, since a terminal substituted position means that carbon has lost a bond — and capping that carbon with hydrogen is precisely what does not happen to a carbon bearing the *O*-linkage.

Fix. Adopt one model throughout, or state capping as an explicit hypothesis alongside H_{crit} and relabel the index-dependent results **[D]** rather than **[A]**.

R5. MAJOR **Proposition 7.6 gives 0, not 1, on the domain of Theorem 7.1.** The proposition's second clause is conditioned on $\mathsf{K} \subseteq \mathcal{F}$, which by Theorem 6.1(i) is exactly the set of cut sets Theorem 7.1 excludes. Over the 54 strictly clean orderings, 42 have an unobstructed first step and 12 an obstructed one. The best lower bound the proposition supplies on Theorem 7.1's domain is therefore 0.

The sentence "Proposition 7.6 is the only part of Theorem 7.1 that is proved rather than enumerated, and it gives 1 where the enumeration gives 2" is a non-sequitur. It appears three times: in the introduction, at the end of section 7, and in Target 14.1,

where it is used to size the paper's principal open problem. The paper's only claim to have proved anything about its central bound is void, and the honest statement is that nothing about Theorem 7.1 is currently proved.

R6. MODERATE **Remark 7.5's "coincidence" is contradicted by the paper's own data.** The remark says the agreement of $\min_{\mathsf{K}} |\mathsf{K} \cap \mathcal{F}|$ with $\min_{\mathsf{K}} \alpha_{\min}$ is "a coincidence of this skeleton". Recomputed on all seven skeletons of Table 4:

Skeleton	k	$\min \lvert\mathsf{K} \cap \mathcal{F}\rvert$	$\alpha_{\min}$
naphthalene	2	0	0
anthracene	3	1	1
phenanthrene	3	1	1
tetracene	4	2	2
benz[*a*]anthracene	4	2	2
chrysene	4	2	2
benzo[*c*]phenanthrene	4	2	2

They agree on every skeleton the paper computes, across three ring counts and four fusion topologies, and Table 2 prints the two columns side by side without comment. "A coincidence of this skeleton" is not supported by anything in the paper. The data instead suggest $\min_{\mathsf{K}} \alpha_{\min} = \min_{\mathsf{K}} |\mathsf{K} \cap \mathcal{F}| = k - 2$, which is the natural strengthening of Conjecture 10.2 and a plausible route to Target 14.1. The paper dismisses its most suggestive observation.

R7. MODERATE **Section 6's framing of Proposition 6.3 is inverted.** A step contributes to the index only if both its carbons are hydrogen-free *and* uncapped, so the parent count $pc(\mathsf{K})$ is an *upper* bound on α, not a lower one, and Proposition 6.3 implies nothing about $\alpha_{\min}$ in either direction. The paragraph following it calls it "the crude form of the obstruction" whose sharp form "is smaller than Proposition 6.3 would suggest", while Remark 7.5 says the sharp value equals it. Both cannot be right and neither is correctly derived.

R8. MODERATE **Nothing in the paper quotients by** $\mathrm{Aut}(\mathsf{X})$**, though section 4 argues at length that one must.** The action of σ on $\mathfrak{K}$ decomposes $27 = 3 + 12 \times 2$ into 15 orbits; since 27 is odd at least one fixed cut set is forced, and the paper does not note this. Up to symmetry there are 15 cut sets, the target pattern is realised by 3 distinct experiments rather than 6, and the activation-minimal family is 7 rather than 12. Remark 4.5 insists that on a centrosymmetric substrate locants are not the invariant, and then every count in the paper is a locant count. The recorded cut set of Remark 8.8 is one of a mirror pair whose partner is another row of Table 1 with identical invariants; neither Remark 8.8 nor Ruling 13.3 says so.

The symmetry also supplies the cheapest consistency check the paper is missing: every column of Table 1 must be σ-invariant. I ran it — zero mismatches of $(|\mathrm{ff}|, \mathrm{weak}, \mathrm{strict}, \alpha_{\min}, p)$ across all 15 orbits. No block performs it.

R9. MODERATE **There are five four-ring cata-condensed isomers of** $C_{18}H_{12}$**, not four.** Triphenylene is $C_{18}H_{12}$, has four rings, and is cata-condensed — no carbon lies in three rings. It is *branched*, so it is not a chain and the construction of Notation 9.1 cannot build it. Notation 9.1 correctly says "chain isomers"; Theorem 9.3, the Table 2 caption, the abstract and the section 12 register all drop the qualifier and claim completeness over cata-condensed isomers.

Built by hand and run through the same analysis: 18 carbons, 21 bonds, cyclomatic number 4, $|\,\mathrm{Aut}\,| = 6$, six fusion carbons, three chords, *six* fusion–fusion bonds, 27 cut sets, 24 clean, 60 clean orderings, $\min|\mathsf{K}\cap\mathcal{F}| = 2$, $\alpha_{\min} = 2$, and the Theorem 6.1 identity holding.

Two claims in Remark 9.4 fail once it is in scope: chrysene no longer uniquely attains the most cut sets, since triphenylene also has 27; and "the most fusion–fusion bonds, 5, tied with benzo[*c*]phenanthrene" is false, triphenylene having six. Chrysene's unique minimisation of clean orderings survives, $54 < 60 < 62$.

Against that, the paper forfeits free evidence: triphenylene is the only available test of Conjecture 10.2 outside the chain family, and it passes. That is exactly the generalisation Remark 10.3 says it is nervous about.

R10. MODERATE **The check-C1 claim in section 12 is not implemented.** Section 12 states that "check C1 of Article 82 is enforced in both directions, so a value emitted and not used, or used and not emitted, fails the build". The script opens exactly three paths: its own source, and its two output files. It never reads a section source. No such check exists. Run by hand, the check passes — all 82 emitted macros are used and none is used unemitted — so the claim is true of this build and false as a statement about enforcement, which is the precise failure mode the paper's own appendix warns against.

R11. MODERATE **Six proofs misdescribe what the script asserts, and two of the assertions cannot fail.**

- Theorem 7.1's proof says the block "asserts that the stated minimum is attained". The assertion compares `min(...)` against membership in the list it was taken from. It cannot fail.
- Theorem 8.5's proof calls clause (ii) an independent check. It follows from two assertions already made on the same fields of the same single pass. "The minimum and maximum over that family" are the first and last elements of a one-element sorted list.
- Theorem 9.3's proof says clause (ii) "is asserted per isomer". W8 asserts only a cardinality identity. The set equality the clause claims is checked for chrysene alone. It does hold for all four isomers, but a counterexample with matching cardinalities would pass.
- Proposition 11.3's proof cites W10 for path-independence. W10 contains no assertion bearing on it. The proposition is a definitional triviality carrying an **[A]** label and a block reference that does not support it.
- Remark 5.3 says W5 "asserts the consequences that the article relies on". W5 asserts counts derived from one greedy SSSR; nothing tests invariance across SSSR choices. The review reports having enumerated every SSSR of all 188 reachable states and found the relevant data invariant, so the remark's substance is true and unverified by the script.
- Proposition 6.6's proof, and the script comment beside it, claim the break condition protects the weak count. Since `ok_strict` is set false whenever `ok_weak` is, the second conjunct of `not ok_weak and not ok_strict` is dead and no strict failure ever breaks the loop. The counts are correct; the safeguard described cannot fire.

R12. MODERATE **Convention 3.2 overstates what is recomputed.** It says four

imported integers are "recomputed here from the skeleton and asserted against the published value". Three are. The alkane-skeleton count 252 is emitted and never recomputed; the script's own comment says so.

R13. MODERATE **The claim about the recorded ordering is never checked.** `AUTHOR_ORDER` appears nowhere in the verification script. Remark 8.8 asserts that the recorded order is clean under both criteria and attains $\alpha_{\min}$ for its cut set; the script locates the cut set and reports its invariants, and the ordering it exhibits is whichever permutation the enumeration reached first. The claim is true — I checked it — but it is prose with no emitted value behind it, against the paper's stated convention, and the datum is dead weight that a wrong edit would not catch.

R14. MODERATE **Hypothesis 12.2 hedges where a computation settles it.** The hypothesis says that under the weak criterion the bound "would then need recomputation over a larger set of orderings and could only decrease". Recomputed under w: all 27 cut sets admit a clean ordering and the minimum activation index is still 2. The three that become clean are the all-ff ones, whose patterns are $(3, 12)$, $(3, 12)$ and $(7, 12)$ — none is the target — so the target minimum is still 3 and the excess still 1.

Corollary 7.2 and Theorem 8.5 are therefore robust to H_{crit}, and the paper could say so unconditionally. As written it leaves its central claim contingent on an admittedly unsettled hypothesis for no reason, and the hedge implies the opposite of the truth.

R15. MINOR 81 **is off by one.** It is emitted before 7, so it records 81 while the generated file contains 82 `newcommand` lines. In a paper whose convention is that every printed number is emitted by the script, the script misreports its own output size.

R16. MINOR **Hard-coded constants and hollow assertions.** `Vorderingspercut` is typed as 24 rather than computed as 4!, and collides numerically with three other macros whose value is also 24, so a swapped macro would be invisible. `Visoamincommon`, `Vconjectureoffset` and `Vbuilderchecks` are hand-typed. The series test is written with the literal 2 rather than the macro. `assert all(len(c) == CYC for c in CUTS)` compares the enumeration's output to its own construction and cannot fail, so Definition 5.1's "necessarily $|\mathsf{K}| = 4$" is asserted vacuously. `assert set(_sizes) <= {1,2}` cannot fail once $|\operatorname{Aut}| = 2$ is asserted. In `assert HADD == 12 and HADD % 2 == 0` the second conjunct follows from the first. In the terminal-pattern filter the second disjunct is dead, since canonicalisation forces $p_1 + p_2 \leq 19$.

R17. MINOR **Theorem 6.1(i)'s "holds for any skeleton" is overclaimed.** The proof uses that a degree-three vertex lies in two or more rings. In a general graph a degree-three vertex may lie in no ring. The step is valid for cata-condensed benzenoids, which is all the paper needs, but the remark should say so.

R18. MINOR **Block attribution and status-label errors.** Table 6's caption credits the hypothesis register to W10; it is emitted outside every block. Section 2 says the [**A**] assertions are "enumerations over graphs on at most 18 vertices, performed in W1–W11", which includes W10, described by the script's own docstring as arithmetic on a bookkeeping convention. Proposition 11.6 therefore carries an [**A**] label Convention 2.1 does not license.

R19. MINOR **Proposition 9.2 checks eight skeletons and lists seven.** The eighth is

benzene, with $|\operatorname{Aut}| = 12$, which appears in no table and is not a benzenoid chain with $k \geq 2$. Either list it or exclude it from the count.

R20. MINOR **Remark 4.5's "mono- against bis-opening" presupposes something false.** No chain-forming cut set contains both K-region bonds — zero of 27 — so within the article's own universe bis-opening is not an alternative. The chemical point may be sound; as written it is a claim about cut sets and it is false.

R21. MINOR **Corollary 7.2 drops the qualifier the paper insists on.** Definition 4.6 and Remark 4.7 stress that the index is evaluated at the moment of cleavage. Corollary 7.2 reads "at which neither carbon of the bond being cut bears a hydrogen", which as written is a statement about the parent, that is, about the different and oppositely-directed quantity of finding R7.

R22. MINOR **Remark 9.5's anecdote misreports which check fired.** The remark says the ring-count error was caught because the automorphism orders "come from outside", and the conclusion repeats the moral as a recommended practice. By the remark's own account what caught it was the chrysene cut-set counts failing to reproduce Article 62's 27 and 24 — imported corpus integers. W8's assertion on carbon and bond counts would also have rejected a fourteen-vertex chrysene at once. The moral drawn is not the moral of the story told.

Checked and found sound

- Every combinatorial integer in sections 3 to 10 reproduces from an independent implementation: 18 carbons, 21 bonds, cyclomatic number 4; six fusion carbons; bond classes 5/8/8; three chords properly contained in the five fusion–fusion bonds; 1189 spanning trees; 27 cut sets; 24 clean; 54 strictly clean orderings and 110 weakly clean; the index distribution $\{2{:}12,\ 3{:}12\}$; seven patterns; $(1, 10)$ the unique terminal pattern, and unique over all 27 and not only the 24; six cut sets realising it, all of index 3; excess 1.
- Proposition 4.3 and Corollary 4.4 in full: $|\operatorname{Aut}| = 2$; σ an involution with no fixed vertex; eleven edge orbits, ten of size two and one of size one, summing to 21; the two K-region bonds in one orbit; $\sigma(\Sigma) = \Sigma$.
- Theorem 6.1(i)'s proof is correct for cata-condensed benzenoids, and the equivalence holds as a genuine set equality on chrysene, all four chain isomers, phenanthrene and triphenylene.
- Remark 3.3 earns its place: no cut set is contained in the chords, so Theorem 6.1(i) would be vacuous for $\mathcal{C}$, and the fusion–fusion bound of Proposition 6.3 would be false for it. W1's proper-subset assertion is the right one to make.
- Theorem 8.5, the paper's main result, in both clauses.
- The series and isomer tables as computed; $\alpha_{\min} = k - 2$ on all seven chains and on triphenylene; chrysene uniquely minimising clean orderings at 54 against 62 for each of the other three chains.
- Plumbing: the Laplacian integrality assertion, the escaping of data labels, the numeric locant sort, the letters-only and V-prefixed macro-name assertions, the

isomorphism and spanning-tree cross-checks of the builder, and the point-group-predicted automorphism orders. The block index has teeth: the review reports inserting an undeclared block and the run failing by name.

Recommended disposition

@p0.055p0.10p0.77@

Ref	Severity	Recommended disposition
R1	MAJOR	Rebuild W10 from a named cut set; emit ON per cut set or per orbit.
R2	MAJOR	Withdraw redox neutrality; restate all section 11 figures; drop the "no new reductant requirement" clause.
R3	MAJOR	Withdraw the scholium's convergence claim and Ruling 13.2, or re-derive both from a model in which the cut carbons bear hydrogen.
R4	MAJOR	Adopt one cleavage model, or state capping as a hypothesis and relabel index-dependent results [**D**].
R5	MAJOR	Withdraw the "gives 1" claim in all three places; state that nothing about Theorem 7.1 is currently proved.
R6	MOD.	Delete "of this skeleton"; record the agreement on all seven skeletons and either conjecture the identity or state that no implication is claimed.
R7	MOD.	Rewrite the paragraph after Proposition 6.3: pc is an upper bound on α.
R8	MOD.	Add an orbit block; assert σ-invariance of Table 1; report 15, 3 and 7 beside the locant counts.
R9	MOD.	Insert "chain" in Theorem 9.3, the caption, the abstract and the register — or add triphenylene and restate Remark 9.4. The second is the better paper.
R10	MOD.	Implement check C1 in the script, or delete the claim.
R11	MOD.	Replace the hollow assertions; correct the six proof descriptions.
R12	MOD.	Say three of four, or recompute the fourth.
R13	MOD.	Assert `AUTHOR_ORDER` and emit its three facts.
R14	MOD.	Compute the weak-criterion minimum, emit it, and state that the bound is 2 under either criterion.
R15	MINOR	Emit the counts after all other emissions.
R16	MINOR	Compute the constants; remove the assertions that cannot fail.
R17	MINOR	Add the cata-condensed hypothesis to the remark.
R18	MINOR	Correct the attribution; relabel Proposition 11.6.
R19	MINOR	List benzene or exclude it from the count.
R20	MINOR	Say what bis-opening means, or drop the clause.
R21	MINOR	Restore "in the substrate as it then stands".
R22	MINOR	Align the anecdote with the check that actually fired.

What a revision would not close

Three things should be said plainly, because a revision that addressed every finding above would still not reach them.

Nothing about Theorem 7.1 is proved. After R5 the paper has no argument for its central bound at all, only an enumeration. The enumeration is legitimate under Convention 2.1 and it is not in doubt. But Target 14.1 is the whole of the paper's mathematical ambition, and the revision will leave it exactly where it was, minus a claim that should not have been made.

Conjecture 10.2 remains thin, and R6 and R9 cut both ways on it. The agreement of $\alpha_{\min}$ with $\min |\mathsf{K} \cap \mathcal{F}|$ on every computed skeleton, and the passage of triphenylene, strengthen it; the fact that the paper could not see either is a reason to doubt how hard it has been tested. Four rings is not many, and the five-ring question of Target 14.4 is where the conjecture will be decided.

The relation between the index and the chemistry is still a hypothesis. H_{crit} is open, and R4 shows the paper does not have a single consistent model of what a cleavage does to a carbon. R14 shows the *bound* is robust to the criterion, which is worth having; it does not show the index is the operative cost. A reader should take Theorem 8.5 as a statement about the skeleton, which it is, and not as a statement about an experiment.

Every numerical claim in this report was recomputed from the skeleton independently of the paper's verification script. Where this report and the paper disagree on a number, the disagreement is stated and the recomputation given; where this report and the originating review disagree, the review is corrected in place. The findings on sections 3 to 10 are that they are correct.

Recommendation

Major revision. The combinatorial core (Sections 3–10), including Theorems 6.1, 7.1, 8.5 and 9.5, is sound and was independently reproduced. Section 11 must be rebuilt from named cut sets drawn from the verified enumeration or withdrawn; the redox-neutrality claim, the inverted scholium, and the claim that Proposition 7.6 proves a bound of 1 on the domain of Theorem 7.1 must be withdrawn. Contradictory cleavage models must be reconciled. The verification script's hollow assertions and unverified claims should be corrected to match the paper's stated conventions. The two minor editorial points from Referee A should also be addressed. No further computation is required for the surviving combinatorial results.

References

[Art83] C. D. Schaper. Related redox ledger (Article 83). *Annals of the Chrysene Formalism*, 2026.

Revision history

Version	Date	Notes
v1.0	2026-11-03	Review record compiled with Article 87.

Living Document Notice

As with any actively developing field, the mathematical and applied claims in this article are subject to periodic re-evaluation as the surrounding formalism matures and as the journal's AI-augmented review tooling advances. Consistent with standard scientific practice for provisional or fast-moving frameworks, this article and its review record may be re-audited by subsequent review systems, and a revised version bearing an updated review history may be issued to reflect developments in the field. Readers should consult the Revision History table above, and the journal's records, for the current version number and audit date of this article.

Editorial Disclaimer

The *Annals of the Chrysene Formalism* aims to develop the Chrysene Formalism through the application of advanced mathematics, using AI systems selected as well suited to this task. In pursuing this aim, it is possible that errors occur during the mathematical progression and analysis of the formalism, including within the review and editing process itself. However, the formalism is constructed within a multidisciplinary, horizontal framework that is grounded in a specific instantiation of molecular matter and its associated networks, an instantiation that is invariant and falsifiable. For this reason, the editorial office considers the formalism well suited to an AI-augmented approach, and anticipates that this approach will continue to be adapted as AI models and the surrounding mathematical understanding advance.