Toward Dimension-Independent Finiteness of Central Configurations for Positive Masses: A Scope-Audited Reduction with Named Open Bridges

Dr. Tamás Nagy

tnagyphd@gmail.com

Working Paper

Build-std v2.0+4758af4739 | 2026-07-10 12:50 | 4902f697

Abstract

We develop a **structural, dimension-uniform strategy** for Smale's 6th Problem — the finiteness of central configurations (CCs) modulo similarity for $N \geq 3$ positive masses in $\mathbb{R}^d$, $d \geq 2$ — and carry it out to a **scope-audited conditional reduction** rather than an unconditional proof. We are explicit about this from the outset: **this paper does not establish finiteness unconditionally.** It reduces the problem to a small, named set of analytic bridge obligations and proves every remaining step. The approach is **structural rather than enumerative**: we complexify a hypothetical one-parameter family of central configurations into an analytic curve over $\mathbb{C}$ and show that the full CC condition — the scalar identity $F = \lambda$ *together with* the gradient identity $G_k \equiv 0$ — obstructs this continuation. This is the architectural core of the paper and the point of departure from the algebraic-elimination tradition (Hampton–Moeckel, Albouy–Kaloshin), whose results are restricted to the coplanar case $d = 2$. The strategy is uniform in the ambient dimension d; the parts that are established are established for every $d \geq 2$ by a single argument, while the residual dimension-specific gap (the $d \geq 4$ null-line exclusion) is one of the named bridges below rather than a discharged step.

What is unconditionally established, and what is not. The elementary singularity machinery — monodromy (Lemma A), private-pole divergence (Lemma C), the gradient-pole separation principle (Lemma D/D′), and the finite-cluster obstructions (Proposition F for $m = 3$, Proposition G′ generic for $m = 4$) — is verified and is the genuine contribution of this work. The finiteness *conclusion*, however, is **conditional on three currently undischarged analytic bridges** (see "Open bridges" below): (i) the Step-8 finiteness-assembly step, which invokes a *local* form of the Borel/Ritt–Steinmetz unicity theorem on a bounded domain (step8_local_borel_contradiction); (ii) the cluster-shape rigidity on Ω_{reg} (step8_cluster_shape_rigidity_omega_reg); and (iii) the null-line finiteness bridge that carries the $d \geq 4$ case (null_line_finiteness). The most serious of these is (i): exponential-sum unicity of Borel/Ritt–Steinmetz type is intrinsically *global* (it needs Nevanlinna growth as $r \to \infty$), and no bounded-domain analogue with the hypotheses actually available here is known to us or proved in this paper. Until (i)–(iii) are discharged, the correct reading of this manuscript is a **conditional finiteness reduction modulo three named bridges**, not a resolution of Problem 6.

Two identities on the complexified curve do the work. The scalar identity rules out odd-order and private even-order singularities via monodromy and divergence (Lemmas A and C). The gradient identity closes the first shared-collision loophole through the **Gradient-Pole Separation Principle** (Lemma D): at any collision where some body has a unique collision partner, the stronger

$R^{-3/2}$ singularity of the gradient forces an uncancellable private pole in that body's own equation. Lemma D is the new ingredient relative to a scalar-only predecessor, but it is not used in isolation: Moeckel's cancellation critique applies to complexified square-root systems broadly, including partial subsystems of the gradient equations. The proof therefore works with the full CC system. Lemma D and Lemma D′ reduce the possible singularities to full cluster collisions (≥ 3 bodies at one complex position), and Propositions F, G′/H, and G handle the remaining cluster cancellations.

Cluster analysis covers all sizes: for $m = 3$, a sign obstruction; for $m = 4$, an algebraic resultant argument with Hessian fallback; for $m \geq 5$, inductive self-application of the complexification–Borel–gradient machinery. No spectral gap hypothesis is assumed; isolation is supplied by Propositions G/H and the Step 8 Borel–Steinmetz bridge, which serve as the analytic substitutes for spectral gap control. No genericity assumptions on the masses are required. A machine-checkable companion layer records each named claim as either a semantically typed formal statement or a named bridge label with citation and dependency metadata; the human-readable statements in this manuscript are linked to that layer by build-invisible formal_ref anchors. The Lean 4 files are verified exports / shadows of this formal layer: all exported modules compile with zero sorry, while classical ingredients that are not formalized end-to-end are recorded as explicit paper-bridge assumptions with citations.

Keywords: Central configurations, Smale's problems, N-body problem, analytic continuation, finiteness, formal verification

MSC 2020: 70F10, 70F15, 37N05, 32B20, 68V20

1. Introduction

1.1 Smale's 6th Problem

Three bodies in the plane can form exactly five central configurations: three collinear arrangements discovered by Euler (1767) and two equilateral triangles found by Lagrange (1772). For four bodies, the count is finite but already requires a substantial algebraic elimination argument to prove (Hampton and Moeckel, 2006, in the plane); earlier results of Palmore (1973), Moeckel (1985), Albouy (1996), and Pacella (1987) established the classification framework, generic finiteness, and symmetry structure via equivariant and Morse-theoretic methods. For five, the finiteness proof by Albouy and Kaloshin (2012) runs to 54 journal pages, is restricted to the coplanar case $d = 2$, and still leaves a codimension-2 exceptional set unresolved; the spatial case $N = 5$, $d = 3$ was settled separately by Hampton and Jensen (2011) for generic masses using numerical algebraic geometry (homotopy continuation). For six or more bodies — and for every dimension $d \geq 3$ at any $N \geq 4$ — the question has been open since Smale placed it on his list of problems for the 21st century. Most recently, Jensen and Leykin (2025) introduced a tropical-geometry framework that establishes *generic* finiteness for planar configurations at arbitrary N and verified it computationally for $N \leq 5$; Chang and Chen (2023) made substantial algorithmic progress on the planar six-body case, reducing the unresolved cases to 24 zw-diagrams. Both results remain confined to the plane and to generic masses; *unconditional* finiteness for all positive masses and all N — in any dimension — has remained open:

> **Problem 6** (Smale, 1998). *Is the number of relative equilibria finite, for each choice of positive masses $m_1, \dots, m_N$?*

A *central configuration* (CC) is a configuration $q = (q_1, \dots, q_N) \in (\mathbb{R}^d)^N$ satisfying

$$\sum_{j\neq k} \frac{m_j(q_j - q_k)}{|q_j - q_k|^3} = -\lambda\, q_k, \qquad k = 1, \dots, N \tag{1}$$

for some $\lambda > 0$ (the sign convention is attractive: $-\lambda q_k$ points toward the center of mass), where masses $m_k > 0$ and the center of mass is at the origin. Two CCs related by rotation, translation, and scaling are considered equivalent. Although Smale phrases Problem 6 in terms of *relative equilibria* (uniformly rotating solutions in a co-rotating frame), the standard reduction (Moeckel 2014, §2; Montgomery 2024, Ch. 1) establishes a bijection between equivalence classes of CCs modulo similarity and families of relative equilibria: each CC generates a unique (up to time-scaling) family of homographic orbits, and conversely every relative equilibrium arises from a CC. Finiteness of CCs modulo similarity therefore implies finiteness of relative equilibria, so a proof of the finiteness statement (Theorem 1) would resolve Problem 6 as stated. As detailed below, the present manuscript reduces that statement to three named analytic bridges rather than establishing it unconditionally. These configurations are not merely geometrical curiosities — they govern the topology of the N-body problem. They determine the bifurcation structure of the integral manifolds, organize all self-similar solutions (homothetic collapse and homographic orbits), and classify the possible behaviors at total collision and parabolic infinity.

The difficulty of the problem escalates sharply with N. Each step from $N = 3$ to $N = 5$ required fundamentally new algebraic machinery — BKK theory, mixed volumes, delicate resultant computations — with intermediate polynomials whose degrees reach into the thousands. For $N \geq 6$, the combinatorial explosion of elimination steps has blocked all direct approaches. *Generic* finiteness (finiteness for all masses outside a measure-zero algebraic subvariety) follows from transversality arguments for each N. The hard part is eliminating the exceptional set: at degenerate configurations where the shape Hessian has a nontrivial null space, Morse-theoretic methods do not apply, and the algebraic degree of the elimination system grows faster than any known technique can control.

The assumption of positive masses is strictly necessary. Roberts (1999) exhibited a specific 5-body mass ratio for which a continuum of central configurations exists, demonstrating that finiteness can fail when negative masses are permitted.

1.2 The Approach — and What Failed First

This paper takes a different path from the algebraic tradition. Instead of counting solutions, we ask a structural question: can the potential and gradient of a hypothetical CC curve survive complexification into $\mathbb{C}$?

The first version of this argument used only the scalar potential identity $F = \lambda$. The monodromy and divergence obstructions (Lemmas A and C) are correct and survive unchanged — they eliminate odd-order zeros and private even-order zeros. But Rick Moeckel (personal communication via R. Montgomery, 2026) identified two distinct weaknesses. First, the old convergence-disk language was wrong: since $F(z) = \lambda$ is constant on the branch, one cannot argue from the convergence disk of the scalar function F; Lemma B has therefore been rewritten as a zero-free continuation plus Borel argument (Remark B.1). Second, in the treatment of *shared* even-order zeros, complexified square-root principal parts can cancel. Moeckel's level-set example was the simplest warning sign, not the full issue. The same cancellation danger can arise in more complicated square-root systems, including partial subsets of the gradient equations. Therefore the proof cannot rely on positivity surviving complexification, nor on any one scalar or partial equation. It must use the full central-configuration system and show, equation by equation, where cancellation is impossible and where the remaining cancellations are isolated.

The complexification architecture itself is untouched by Moeckel's critique; what needed strengthening was the data extracted *from* the complexified curve. The full CC condition is not just $F = \lambda$ but also the bodywise gradient system $G_k \equiv 0$ for every body k. The gradient has a stronger singularity ($R^{-3/2}$) than the potential ($R^{-1/2}$). At any collision where a body participates in exactly one pair collision, this creates an uncancellable private pole in that body's own gradient equation. This **Gradient-Pole Separation Principle** (Lemma D) closes the first shared-zero loophole. The remaining multi-partner and full-cluster configurations are then handled by Lemma D′ and the cluster-isolation package (Propositions F, G′/H, G), so the final contradiction uses all equations, not a selected partial subsystem.

Principle 1.1 (Gradient-Pole Separation Principle). *The scalar potential equation pools all pair interactions into one sum, so singular pair-poles can cancel. The gradient system separates the same singularities body-by-body. If a body has exactly one collision partner at a complexified singular point, its own gradient equation contains a private $R^{-3/2}$ pole with no possible cancellation. Thus every collision graph with a degree-one vertex is excluded at the level of a single body equation.*

This is the first conceptual hinge of the proof. The scalar identity $F = \lambda$ finds the singularities; the full bodywise gradient system starts separating genuine CC curves from false scalar cancellations. Lemma D alone is not the whole proof: configurations with no degree-one vertex require the multi-partner obstruction (Lemma D′) and the cluster-isolation analysis.

Gap Closure Summary. The loophole in the scalar-only version and the ingredient that closes it, side by side:

Singularity type at z_0	Scalar $F = \lambda$ rules it out?	Full ($F = \lambda$, $G_k \equiv 0$) rules it out?	Mechanism
Odd-order R-zero	Yes	Yes	Monodromy (Lemma A)
Private even-order R-zero	Yes	Yes	Divergence (Lemma C)
Isolated pair collision within shared zero (only one body has one partner)	**No — loophole**	**Yes — new**	Gradient pole (Lemma D)
Disjoint pair collisions at same z_0	**No — loophole**	**Yes — new**	Gradient pole (Corollary D.2)
Full cluster collision ($m \geq 3$ bodies)	No	Handled separately	Propositions F, G′/H, G

Why the old argument failed on shared zeros: complexified square-root terms can cancel.

The scalar identity $F = \lambda$ is the most visible place where this happens, but the warning is broader: cancellation may also occur inside partial gradient subsystems. Any proof that rules out shared zeros by inspecting only one pooled scalar identity, or only a few equations, has not answered the cancellation problem. *Why the fix works:* the revised proof uses the full bodywise CC system. Lemma D handles the first obstruction locally: whenever a body has a unique collision partner, that body's gradient equation is single-handedly broken by a private $R^{-3/2}$ pole. Lemma D′ and the cluster propositions then handle the multi-partner cases that Lemma D alone cannot exclude. The only residual full-cluster scenarios are isolated by Propositions F (for $m = 3$), G′ with Hessian fallback H (for $m = 4$), and the inductive isolation theorem G (for $m \geq 5$).

Concrete illustration (the scalar/gradient asymmetry in 4 lines). Take $N = 4$ bodies with two disjoint pairs colliding simultaneously at a complex point z_0: $R_{12}(z_0) = R_{34}(z_0) = 0$, all other $R_{kl}(z_0) \neq 0$. In the scalar sum $F = \sum m_i m_j R_{ij}^{-1/2}$, the poles from $R_{12}^{-1/2}$ and $R_{34}^{-1/2}$ can cancel — so $F = \lambda$ is consistent. **But** in body 1's gradient $G_1 = m_2(q_1 - q_2)/R_{12}^{3/2} +$ [analytic terms], only the first term blows up (body 1 collides only with body 2), and nothing in G_1 alone can cancel it. The gradient equation $G_1 \equiv 0$ is broken. This is the key asymmetry: the scalar pools all pairs; the gradient isolates each body. See §3.4 for the general proof.

Remark (Moeckel's cancellation test). Any singularity-analysis approach to Smale's 6th must answer: *why do the complexified square-root terms not cancel in the bad cases?* The level-set example shows the danger in its simplest scalar form, but the test is stronger: cancellations can also appear in partial systems. The revised argument passes this test only because it uses the full CC equations. The pole-order gap — $R^{-3/2}$ in the gradient versus $R^{-1/2}$ in the scalar — gives private bodywise obstructions when a degree-one collision partner exists; the remaining no-degree-one collision graphs are routed into Lemma D′ and the full cluster analysis, where the entire system is used.

1.3 Overview of the Proof

The proof in one sentence. We complexify a hypothetical CC family into an analytic curve, use the full bodywise gradient system to prevent the square-root cancellations that defeat scalar or partial-equation arguments, reduce the remaining singularities to isolated full clusters, and derive a contradiction via Borel-type rigidity.

Conceptual Core

The proof is built around a simple separation of roles:

1. **Curves force singularities.** If a non-constant analytic family of CCs exists, Borel's theorem forces some pair-distance function R_{ij} to vanish somewhere on its complexification.
2. **Scalar singularities almost cancel.** The potential identity $F = \lambda$ rules out odd-order zeros and private pair collisions, but shared even-order poles can cancel in the scalar sum.
3. **Gradient singularities cannot always cancel bodywise.** The equation $G_k \equiv 0$ is not a scalar sum over all pairs; it is one equation per body. A body with a unique collision partner carries a private $R^{-3/2}$ pole, so degree-one collision graphs are impossible.
4. **The remaining cancellations are forced into clusters.** Multi-partner non-cluster cases are handled by Lemma D′, and the surviving singularities are full cluster collisions of $m \geq 3$

bodies. Propositions F, G′/H, and G show that these cluster solutions are isolated for every positive mass vector.

5. **Isolated singularities cannot support a curve.** Pole-clearing plus Borel–Steinmetz rules out non-constant zero-free segments between the isolated cluster points; the identity theorem then forces the original curve to be constant.

In this form, the proof has one new local obstruction and one global assembly mechanism: **Gradient-Pole Separation** eliminates degree-one shared collisions that scalar or partial-equation arguments cannot control, and the cluster-isolation induction prevents the remaining multi-body cancellations from organizing into a positive-dimensional family.

Proof at a Glance. The logical chain of the proof, with each step's role:

$$\boxed{\text{Assume a CC curve exists}}$$

$\downarrow$ Complexify (Step 2: analytic continuation to $\mathbb{C}$)

$\downarrow$ Monodromy $\Longrightarrow$ all R-zeros have even order (Step 3, Lemma A)

$\downarrow$ Divergence $\Longrightarrow$ no private zeros (Step 4, Lemma C)

$\downarrow$ **Gradient pole $\Longrightarrow$ no non-cluster zeros (Step 5, Lemma D — NEW)**

$\downarrow$ Borel $\Longrightarrow$ zeros must exist (Step 6, Lemma B)

$\downarrow$ Only full cluster collisions remain

$$\downarrow \quad \begin{cases} m = 3: & \text{sign obstruction (Prop. F)} \\ m = 4: & \text{resultant + Hessian fallback (Props. G}' + \text{H)} \\ m \geq 5: & \text{inductive Borel self-application (Prop. G)} \end{cases}$$

$\downarrow$ All cluster singularities are isolated (finitely many local branches)

$\downarrow$ Borel–Steinmetz forces constant shape on zero-free segments (Step 8)

$\downarrow$ Identity theorem globalizes $\Longrightarrow$ curve is constant

$$\boxed{\text{Contradiction } \blacksquare}$$

The local obstruction steps (Steps 3–5, and the finite-cluster obstructions of Step 7 for $m = 3, 4$) hold for all positive masses $m \in \mathbb{R}^N_{>0}$ without genericity assumptions. The **final two boxes are not unconditional**: the Borel–Steinmetz step (Step 8) is invoked here in a *local* form on a bounded domain that is recorded as a named bridge (step8_local_borel_contradiction) and is not derived from the global theorem, and the cluster-shape rigidity it relies on (step8_cluster_shape_rigidity_omega_reg) is likewise an open bridge; the $d \geq 4$ null-line exclusion is carried by a third bridge (null_line_finiteness). No spectral gap hypothesis is assumed. The gradient pole obstruction (Step 5, bold) is the new local ingredient: it rules out degree-one shared collisions by using a bodywise equation that scalar or partial-equation arguments cannot control. The remaining cancellation-prone configurations are handled by the multi-partner obstruction and cluster-isolation induction — up to the named bridges just listed.

The proof is a contradiction argument. We assume a continuous family of central configurations exists and derive a contradiction using both the scalar potential identity $F = \lambda$ and the full gradient condition $G_k \equiv 0$.

Suppose a continuum of CCs exists. By the Łojasiewicz structure theorem, the CC variety is a compact real-analytic set, and a positive-dimensional component contains a non-constant real-analytic curve germ $q(u)$. Along this curve, both $F = \lambda$ and $G_k = 0$ hold identically. The proof proceeds in **three stages** (refined into 8 detailed steps in §4; see the Reviewer Roadmap below for the step-by-step dependency map):

Stage I — Complexification and scalar obstructions (§4 Steps 1–4). Constancy of λ (chain rule) extends both identities to complex z. On the complexified curve, the monodromy argument (Lemma A) forces every R-zero to have even order, and the private-pole argument (Lemma C) rules out any zero in which only one pair collides. *What remains:* shared even-order zeros.

Stage II — Gradient pole obstruction and cluster analysis (§4 Steps 5, 7). The gradient condition $G_k \equiv 0$ on the same complexified curve — not just the scalar $F = \lambda$ — is the new ingredient relative to the scalar-only predecessor. Lemma D shows that whenever any body participates in exactly one pair collision, its gradient equation has an uncancellable $R^{-3/2}$ pole; this excludes every collision graph with a degree-one vertex. The remaining no-degree-one configurations are precisely the cancellation-prone cases: either they collapse to the full cluster setup, or Lemma D′ overdetermines their leading principal parts. Propositions F, G′/H, and G then handle the full clusters: $m = 3$ by a sign obstruction, $m = 4$ by an algebraic resultant with Hessian isolation fallback, $m \geq 5$ by inductive self-application of Stages I–II to the sub-cluster (cluster isolation theorem, Proposition G).

Stage III — Finiteness assembly (§4 Steps 6, 8). Step 6 (Borel forces at least one R-zero to exist, Appendix A / Proposition B_0) is a classical bridge: it rules out zero-free entire arcs. Step 8 is now treated as a *local bridge*, not as a derived global theorem. On a bounded simply connected rectangle Ω containing a non-constant sub-arc, compactness gives finitely many cluster collisions in that region, and Lemma A gives a single-valued meromorphic factorization of each $R_{ij}^{-1/2}$ on Ω. The pole-cleared identity on Ω has polynomial coefficients and holomorphic exponents on Ω. What remains to be justified is the local unicity step: this local identity must force the same termwise vanishing that the global Borel–Steinmetz theorem gives on $\mathbb{C}$, without silently extending the functions from Ω to the whole plane. The formal companion layer therefore records Step 8 as an explicit local bridge (Smale6.step8_local_borel_contradiction / Smale6.step8_borel_steinmetz_vanishing) rather than as an end-to-end machine-discharged theorem.

The formal companion layer records the Stage I–III claims as typed statements and paper-bridge assumptions. Its deterministic projection, formal_statements.md, is the statement-level reference checked by the paper: each paper-facing platonic= anchor in the manuscript points to a formal_ref anchor in that file. The Lean 4 exports compile with zero sorry; their role is to verify the formal layer, while the paper-bridge entries make the remaining classical ingredients explicit (Borel, Hayman–Borel–Steinmetz, Łojasiewicz, the gradient-pole bridge, the graph/cluster bridge, and Step 8's local bridge). No spectral gap hypothesis is assumed; the unresolved issue is not a spectral gap assumption, but whether the local Step 8 bridge and the graph-to-cluster bridge are stated with exactly the hypotheses they need.

Reviewer Roadmap (proof dependency table). The following table maps each logical step to its source, its formal status, and its role in the proof. The "All $m > 0$?" column records whether a step, *once its stated inputs are granted*, holds for all positive masses in $\mathbb{R}^N_{>0}$; the honest headline is that Steps 1–7 are established while Step 8 is **conditional on the local bridge** (and the $d \geq 4$ / cluster-rigidity bridges), so the chain as a whole is a conditional reduction, not an unconditional proof. The "Uses §5?" column isolates the evidence-versus-proof distinction: §5 is entirely diagnostic (it supports Conjecture G′.3 but is never invoked by the logical chain).

Step	Source	Kind	All $m > 0$?	Uses §5?
1. λ constant	§4 Step 1	Chain rule	Yes	No
2. Complexification	§2.2, §4 Step 2	Real-analytic continuation	Yes	No
3. Monodromy (Lemma A)	§3.1	Formal layer + Lean export	Yes	No
4. Private R-zero divergence (Lemma C)	§3.3	**Kernel theorem** (Laurent.single_pole_no_cancel)	Yes	No
5. Gradient pole (Lemma D, **new**)	§3.4	Laurent engine derived theorem + domain bridges	Yes	No
6. R-zeros exist (Lemma B / App. A)	§3.2, App. A	Classical (Borel 1897)	Yes	No
7a. $m = 3$ non-existence (Prop. F)	§3.5	Formal layer + Lean export	Yes	No
7b. $m = 4$ non-existence / isolation (Prop. G′ + H)	§3.7, §3.8	Formal layer + bridge audit	Yes†	No
7c. $m \geq 5$ isolation (Prop. G)	§3.8	Formal layer + kernel theorem	Yes‡	No
8. Finiteness assembly	§4 Step 8	**Named local bridge** (local Borel/Ritt-Steinmetz on Ω; not a global meromorphic-extension theorem)	Conditional on the local bridge	No
§5 (numerical sweep at μ^*; resultant table)	§5	**Evidentiary / diagnostic only — not a logical dependency**	—	—

† Proposition H is intended to cover the exceptional masses where Proposition G′ is conjectural (Conj. G′.3), but this coverage is now recorded as a named bridge rather than as a discharged

kernel proof. ‡ Inductive Borel argument plus graph-Laplacian non-degeneracy on the real axis; the remaining graph-to-cluster and Step 8 local-unicity steps are explicit bridge obligations.

Cluster handling map. After Lemma D reduces the problem to full cluster collisions of $m \geq 3$ bodies, each cluster size is handled by a dedicated proposition. The following table shows where each cluster size is closed, under what hypotheses, and what the residual claim is:

Cluster size	Primary result	Method	Proposition	Unconditional in positive masses?
$m = 3$	Non-existence of principal-branch solutions	Sign obstruction on $c_3 = \sqrt{m_3/m_2}$	**F** (§3.5)	Yes — pure sign argument
$m = 4$	Non-existence (generic masses)	Polynomial resultant $\mathrm{Res}_{c_3}(P, Q) \neq 0$; verified at 9 configurations	**G′** (§3.7)	For all tested masses; conjectured $\forall \mathbb{R}^4_{>0}$ (Conj. G′.3)
$m = 4$	Isolation at *any* exceptional masses (fallback)	Graph-Laplacian rank $m - 1$ on the real axis	**H** (§3.8)	Yes — covers the conjectural gap in G′.3
$m \geq 5$	Isolation of every critical point	Induction: self-apply Stages I–II to the m-body sub-cluster; Borel + graph Laplacian	**G** (§3.8)	Yes — no spectral-gap hypothesis, no genericity

What this buys. Propositions F, G′, and G together rule out positive-dimensional cluster solutions for *all* positive masses. For $m = 4$, Proposition H closes the single conjectural gap in G′ by providing Hessian isolation as a fallback — so Theorem 1 does not depend on Conjecture G′.3. The $m \geq 5$ case admits isolated critical points for specific masses (e.g. the regular pentagon for $N = 5$ equal masses, Remark G.3) but these are isolated, not positive-dimensional.

Minimal core the reviewer must check. The formal companion layer separates machine-checked composition from classical mathematical ingredients. The Lean export still exposes six classical labels in the scalar algebraic core; the paper-bridge layer additionally names the gradient-pole, graph-to-cluster, Step 8 local-unicity, null-line, and capstone bridges used by the full Smale-6 proof. The key point is not the raw axiom count but the audit invariant: every paper-facing classical ingredient is named, cited, dependency-linked, and visible in formal_statements.md. The current draft should not be described as fully end-to-end formalized until the graph-to-cluster and Step 8 local bridge statements pass the formal projection and theorem-statement stress tests.

Lean name	Classical content	Used in
borel_exp_independence	Borel's unicity theorem for entire exponentials (Borel 1897)	Lemma B Step 1, Prop. B_0 (B2), Prop. G inductive step
borel_steinmetz_entire_coeff	Hayman's generalized Borel–Steinmetz: entire-coefficient exponential sums with non-constant exponent differences and the small-coefficient growth condition $T(r, A_j) = o(T(r, e^{f_i - f_j}))$ vanish termwise (Hayman 1964, Ch. IV, Theorem 1.62; specializes to Steinmetz 1980 / Lang 1987 polynomial-coefficient form when coefficients are polynomial)	§4 Step 8 (class vanishing)
entire_log_of_nonvanishing	Non-vanishing entire function admits entire logarithm (Conway, *Functions of One Complex Variable*, Ch. VII)	Prop. B_0 setup, §4 Step 6
zero_free_potential_const	Zero-free potential identity forces constant (derived from the previous two; encoded as a single Lean step)	App. A, Prop. B_0 (B1)
cc_curve_branch_analysis	Combined branch-analysis bundle: identity theorem + monodromy + Borel applied to a CC curve (used as one composite step to shorten the algebraic core)	Lemmas A, B, C — the logical load of this axiom is already covered by the three axioms above; it is a *composite shortcut*, not a new classical assumption
lojasiewicz_0dim_finite	Łojasiewicz structure theorem for real-analytic varieties (Łojasiewicz 1965): compact 0-dim analytic set is finite	§4 Step 8 (finiteness assembly)

The scalar core reduces to five independent textbook ingredients (Borel 1897; Hayman's generalized Borel–Steinmetz theorem with entire coefficients of small Nevanlinna order; entire-log existence; Łojasiewicz 1965; the identity theorem and monodromy from standard complex analysis). The full proof also uses paper-bridge entries for the gradient-pole and cluster-isolation mechanisms developed in §§3–4. The formal audit checks the complete dependency graph: no paper-facing named claim is left as an invisible or unlinked assumption (see §6.4).

1.4 Main Result

Theorem 1 (Target statement — CC Finiteness, *conditional*). *Assume the three analytic bridge statements named below — (B1) the Step-8 local unicity bridge step8_local_borel_contradiction, (B2) the cluster-shape rigidity bridge step8_cluster_shape_rigidity_omega_reg, and (B3) the null-line finiteness bridge null_line_finiteness — hold with the hypotheses stated in §4 Step 8 and §3.4. Then for $N \geq 3$ bodies in $\mathbb{R}^d$ ($d \geq 2$) with positive masses $m \in \mathbb{R}^N_{>0}$, the number of central configurations modulo similarity is finite.*

We stress that **(B1)–(B3) are not proved in this paper**; they are the open bridges on which the finiteness conclusion is contingent (see the disclosure box below). What *is* proved unconditionally is the reduction machinery: Steps 1–6 (monodromy, divergence, gradient pole) hold for all positive masses without genericity assumptions; Proposition F eliminates 3-body clusters for all positive masses; Proposition G′ eliminates 4-body clusters for all tested mass configurations (with isolation for hypothetical exceptional masses conjectured via the graph-Laplacian structure, Proposition H, itself recorded as a bridge). For $m \geq 5$ clusters, principal-branch solutions may exist for specific masses (Remark G.3 exhibits examples at $m = 5$ with equal masses); Proposition G *aims* to show every such solution is isolated, but its capstone assembly relies on (B1)–(B2). No spectral gap bound is used. We do not obtain explicit upper bounds on the count; see §6.5.

> **Open bridges — what remains to be discharged.** The finiteness conclusion is *conditional* on three currently undischarged analytic bridge obligations. A reader should treat everything downstream of these as contingent.
>
> **(B1) Step-8 local unicity — the load-bearing gap** (Smale6.step8_local_borel_contradiction, §4 Step 8). Step 8 is the *sole* global-assembly step: everything before it only shows that surviving singularities are isolated cluster collisions; converting "isolated clusters" into "the curve is constant" is done exclusively here. Step 8 invokes a *local* form of the Borel/Ritt–Steinmetz exponential-sum unicity theorem on a bounded simply connected domain Ω. But Borel/Hayman/Steinmetz unicity is intrinsically **global**: its conclusion is driven by a Nevanlinna growth condition $T(r, A_j) = o\big(T(r, e^{f_i - f_j})\big)$ as $r \to \infty$, and on a bounded Ω there is no growth data — any non-vanishing holomorphic factor is e^H with H holomorphic and unconstrained. We do not know a bounded-domain analogue with the hypotheses actually available here, and we do not prove one. Closing (B1) requires either a genuine single-valued meromorphic extension of $q(z)$ from Ω to $\mathbb{C}$ (then the global theorem applies) or a proved bounded-domain unicity theorem. This is the deepest open point of the paper.
>
> **(B2) Cluster-shape rigidity on Ω_{reg}** (Smale6.step8_cluster_shape_rigidity_omega_reg, §4 Step 8, item (R3)). Leading-order pole cancellation gives a sub-cluster critical shape *pointwise at each collision point*; upgrading this to rigidity of the cluster shape at every *regular* point of Ω_{reg} is a separate, currently undischarged step.
>
> **(B3) Null-line finiteness for $d \geq 4$** (Smale6.null_line_finiteness, §3.4 / Prop. D.4). For $d \leq 3$ the collinearity reduction is argued directly. For $d \geq 4$ the Witt index is ≥ 2, the null cone admits non-proportional isotropic displacements, and the exclusion of null-line residual families is delegated to Prop. D.4, which delivers only *generic* absence and *isolated-at-special-mass*, recorded as a paper-level obligation rather than a discharged theorem.
>
> Independently of (B1)–(B3), the theorem's status is *unchanged* by (i) Conjecture G′.3

(that the $m = 4$ resultant variety is empty over positive reals), (ii) the numerical tables of §5, and (iii) any genericity assumption on the masses: the $m = 4$ analysis combines Proposition G′ (non-existence for generic masses) with Proposition H (Hessian isolation for exceptional masses), and §5 is diagnostic only. In other words, removing Conj. G′.3 or §5 does not weaken the reduction; it is (B1)–(B3) alone that stand between the reduction and an unconditional proof.

The machine-checked companion layer's "0 undischarged" audit (§1.5) refers to tactic-layer hypotheses *inside the encoded skeleton*; it does **not** assert that (B1)–(B3) are theorems. They are disclosed bridge leaves, and they are open.

Note on "Hessian". Proposition H concerns the **sub-cluster gradient Hessian** $H^{\mathrm{sub}} = (\partial \mathcal{G}_k / \partial c_j)$, evaluated at cluster critical points ($\mathcal{G}_k = 0$); on the real axis $H^{\mathrm{sub}} = 2L$ with L the graph Laplacian of the mass-weighted complete graph K_m (§3.8). This is a *different* object from the standard CC shape-space Hessian $\nabla^2 \tilde{U}$: they live on different spaces (sub-cluster collision directions vs. full-system shape coordinates) and degenerate at independent mass values. For example, §5.1 exhibits a full-system CC at μ^* where the shape Hessian $\nabla^2 \tilde{U}$ has a 2-dimensional null space — the classical setting where Morse-theoretic methods fail — yet the sub-cluster analysis (Propositions G, G′, H) is *independent* of any such shape-space degeneracy: Proposition H(iii) gives $\operatorname{rank}(H^{\mathrm{sub}}) = m-1$ generically on the real axis and $\operatorname{rank}(H^{\mathrm{sub}}) \le m-2$ at every cluster critical point ($\mathcal{G}_k = 0$, via Part (ii)), which is exactly the IFT threshold used in Lemma G.1. All uses of "Hessian" in Propositions G, G′, H and in this box refer to H^{sub}, never to $\nabla^2 \tilde{U}$.

1.5 Formal Companion Layer and Lean Verification

The proof is tracked by a machine-checkable companion layer. Each named claim is recorded in one of two forms: (1) a *semantically typed formal statement* encoding the mathematical content in the formal proof kernel language, or (2) a *named bridge label* with prose justification, citation, and dependency metadata — serving as a tracked placeholder for classical arguments whose full formalization is deferred. The paper's human-readable statements are linked to that layer by build-invisible formal_ref anchors; the deterministic file formal_statements.md is regenerated from the kernel and checked byte-for-byte by the formal-layer gate. The Lean/Mathlib files are verified exports of the semantically typed statements: they compile with zero sorry, but they are not the prose source of truth. The prose source is the formal statement layer, and the paper is a narrative projection of it.

This represents a *scope-audited classical proof attempt with a machine-checked skeleton*, not a full end-to-end formalization of all background analysis. The kernel tracks the statement graph and verifies algebraic skeletons as semantically typed statements; the analytic closure of Step 8 and several other deep analytic mechanisms remain explicitly named bridge assumptions. The current trust audit records 114 disclosed bridge assumptions in the full module, classified into two tiers: 64 are external_classical (standard complex analysis, real-analytic geometry, and Laurent algebra — results that can be cited from standard textbooks without Smale-6-specific content), while 50 are actionable Smale-6-specific encoding obligations (CC non-degeneracy claims, capstone assembly, pole-structure lemmas, and induction composition steps that constitute the paper's novel analytical contribution). A strict theorem-construction audit also checks every qed() return: a 2026-06-06 repair reduced silent verified =False theorem attempts from seven to zero by fixing statement-shape mismatches, adding kernel normalization for legacy Atom(False) wrappers, and reclassifying the Lemma B continuation discharge as explicit deferred trust. Within the capstone chain (T22), all deferred proofs are disclosed. There are 0 undisclosed trust items and 0 undischarged tactic-layer hypotheses.

Classical results and analytic bridge mechanisms are declared as named bridge labels with explicit citations, paper-section references, and dependency declarations. The distinction is important: a bridge label records *what* the claim is, *where* it is justified (by citation or paper section), and *what depends on it* — but it does not claim a full Mathlib-level formalization of that classical argument. The Nevanlinna small-coefficient growth condition — a key ingredient of the Borel–Steinmetz application in Step 8 — has been decomposed into two atomic sub-bridges: (i) polynomial coefficients imply $O(\log r)$ Nevanlinna characteristic, and (ii) $O(\log r)$ characteristic implies the small-coefficient dominance condition relative to non-constant exponentials. This decomposition makes the trust leaf structure more granular and auditable.

The Step 8 non-proportionality inputs have likewise been decomposed. The old scalar Smale6.cc_es_arg_not_proportional bridge is now a derived compatibility theorem whose three leaves are: (i) the finite Borel witness is indexed by exponent-class representatives, (ii) collapse of the relevant distance-ratio classes would force the CC arc to be constant by distance-geometry rigidity, and (iii) these two facts assemble into pairwise non-proportionality of the pole-cleared arguments. The gradient analogue Smale6.cc_ges_arg_not_proportional is also derived from three leaves: gradient class-representative indexing, inter-cluster ratio rigidity, and the final class-separation-to-argument-non-proportionality assembly. This exposes the exact place a reviewer should attack: not “Borel applies” as a black box, but the witness indexing and ratio-rigidity mechanisms that supply Borel’s pairwise non-resonance hypotheses.

The capstone generic-packing bridge has also been split. The old actionable Smale6.generic_packing_from_step8 bridge is now a derived compatibility theorem from three smaller leaves: (i) Step 8 outputs jointly cover the no-positive-dimensional regular and cluster strata, (ii) that coverage makes the generic-stratum packing bound applicable, and (iii) the packing-bound result assembles into the GenericPackingWitness consumed by the capstone. This preserves the formal dependency from T22 through K1/K2/K3 while making the reviewer attack surface more precise than a single “generic packing” black box.

The capstone degenerate-isolation bridge has now been split in the same way. The old actionable Smale6.degenerate_isolation_witness_from_body_count bridge is a derived compatibility theorem from three smaller leaves: (i) body count at least three covers the degenerate alternatives (null-line, collision-cluster boundary, exceptional-mass cases), (ii) that case coverage gives local isolation by the cyclotomic eigenvalue positivity and Lyapunov-Schmidt / Morse chain, and (iii) the local-isolation output assembles into the DegenerateIsolationWitness consumed by the capstone.

The final finite-union capstone bridge has also been split. The old actionable Smale6.capstone_finite_witness_from bridge is now a derived compatibility theorem from three smaller leaves: (i) the generic and degenerate witnesses cover the relevant regular/generic and degenerate strata, (ii) that coverage gives a finite-union bound, and (iii) the finite-union bound assembles into the CapstoneFiniteWitness/CCFinite predicate consumed by the final theorem.

The kernel verifies 131 proved theorems. Of these, approximately 111 are compositional theorems establishing the proof chain through multi-step arguments — monodromy, divergence, pole-cancellation, cluster isolation, the Borel contradiction, and the capstone finiteness assembly. The remaining ~20 are auxiliary numerical verifications (mass-ratio positivity lemmas, null-line Morse-index checks, sign-correction identities) proved by linear or nonlinear arithmetic. Both categories are machine-verified and load-bearing, but the distinction matters for assessing proof depth: the compositional theorems encode the mathematical argument structure, while the arithmetic verifications confirm that specific numerical conditions hold.

1.6 Structure of the Paper

Section 1.2 describes the approach and the critical failure that shaped it. Section 1.3 provides the three-stage proof strategy and the Reviewer Roadmap dependency table. Section 2 establishes notation. Section 3 develops the singularity analysis: the monodromy, divergence, and gradient pole lemmas (A–D), the cluster overdetermination and non-existence results (Propositions E–G′), and the Hessian structure and isolation theorem (Propositions H and G). Section 4 assembles the main proof in 8 detailed steps. Section 5 presents computational verification at a degenerate CC. Section 6 discusses extensions and the formal companion layer / Lean verification boundary. Appendix A gives the full Borel argument establishing R-zero existence; Appendix B documents code availability.

1.7 Scope and Non-Claims

This paper develops a *conditional reduction* toward finiteness of central configurations for all positive masses; it does **not** claim an unconditional proof. It also does *not* provide:

0. **An unconditional finiteness theorem.** The finiteness conclusion is contingent on the three named analytic bridges (B1) step8_local_borel_contradiction, (B2) step8_cluster_shape_rigidity_omega_reg, and (B3) null_line_finiteness (see §1.4 disclosure box, §4 Step 8, §3.4). Bridge (B1) in particular is the load-bearing gap: no bounded-domain form of the global Borel/Ritt–Steinmetz unicity theorem with the hypotheses available here is proved. Until (B1)–(B3) are discharged, the result is a reduction, not a resolution of Problem 6.
1. **Explicit upper bounds.** The reduction argues that no positive-dimensional families exist (modulo the bridges above), but does not yield a computable bound on the count. See §6.5 for discussion.
2. **Negative or zero masses.** The argument uses mass positivity at five distinct points (§6.3). Roberts' (1999) counterexample shows finiteness fails for negative masses.
3. **Full end-to-end Mathlib formalization of the entire proof.** The companion layer and Lean export verify the named proof skeleton, but they do not constitute a full Mathlib formalization of all background analysis. The Lyapunov–Schmidt reduction, analytic-set dimension arguments, and several complex-analysis ingredients remain classical paper bridges with citations. See §6.4 for the precise scope boundary.
4. **Unconditional proof that $\mathcal{Z} \cap \mathbb{R}^4_{>0} = \emptyset$ for $m = 4$.** Conjecture G′.3 (that the resultant variety is empty over positive reals) is supported by extensive numerical evidence but remains unproved. The proof does not require this conjecture — the Hessian isolation fallback (Proposition H) covers any exceptional masses.

2. Setup

2.0 The Newtonian Potential Convention

Throughout this paper we follow the standard convention of Smale's formulation, codified in Moeckel (2014, eq. 2): for N point masses with positions $q_i \in \mathbb{R}^d$, the pair interaction is the classical Newtonian $1/r$ potential

$$U(q) = \sum_{i<j} \frac{m_i m_j}{r_{ij}}, \qquad r_{ij} = |q_i - q_j|, \tag{U}$$

with the same functional form in every ambient dimension $d \geq 2$. The d-dependence enters only through the gradient $\nabla_i U \in \mathbb{R}^d$, which has d components, but the pointwise force magnitude is $m_i m_j / r_{ij}^2$ in every d (Moeckel 2014, eq. 1). This is the classical N-body problem embedded in $\mathbb{R}^d$ — the same pairwise Newtonian interaction used in Smale (1998), Hampton–Moeckel (2006), Albouy–Kaloshin (2012), and Jensen–Leykin (2025). It is *not* d-dimensional Newtonian theory, in which a point source would induce a force falling off as $r^{-(d-1)}$ by Gauss's law; Smale's 6th problem asks about the classical potential, and that is the convention adopted throughout.

2.1 Shape Space and the CC Variety

Fix $N \geq 3$ bodies in $\mathbb{R}^d$ with $d \geq 2$ and positive masses $m_1, \dots, m_N > 0$. Define mass-weighted coordinates $\xi_k = \sqrt{m_k}\, q_k$, the moment of inertia $I = \sum |\xi_k|^2$, and the *normalized potential*

$$\tilde{U} = \sqrt{I} \cdot U, \qquad U = \sum_{i<j} \frac{m_i m_j}{r_{ij}}, \qquad r_{ij} = |q_i - q_j|.$$

The *shape space* is the quotient $\mathcal{S} = \{\xi \in (\mathbb{R}^d)^N : I = 1,\ \sum m_k q_k = 0\}/\mathrm{SO}(d)$. For $d = 2$, the SO(2)-action on non-collision configurations is free, and $\mathcal{S}$ is a smooth compact manifold of dimension $2N - 4$. For $d \geq 3$, the $\mathrm{SO}(d)$-action has nontrivial isotropy on lower-dimensional configurations (e.g., collinear configurations in $d = 3$ have stabilizer SO(2)); $\mathcal{S}$ is then a compact real-analytic stratified space (an orbifold quotient) whose principal stratum has dimension $dN - d - \binom{d}{2} - 1$ (reducing to $3N - 7$ for $d = 3$) and whose lower-dimensional singular strata correspond to fixed-point sets of nontrivial isotropy subgroups. The normalized potential $\tilde{U}$ descends to a real-analytic function on each smooth stratum; criticality $\nabla_s \tilde{U}(s) = 0$ is well-defined intrinsically on every stratum by stratified differential geometry (Łojasiewicz 1965, §III; Bierstone–Milman 1988, §6). All dimension counts in the proof (Steps 1–8) are stated for the principal stratum, but this is for descriptive concreteness only: the singularity analysis of §3 (real-analytic CC arcs, complexification of pair distances, complex collisions, gradient poles) is stratum-agnostic and operates entirely on a smooth real-analytic arc inside whatever smooth stratum a given irreducible component of $\mathcal{V}$ contains. The arc-existence step we actually invoke (Łojasiewicz 1965, §III; Bierstone–Milman 1988, §6) does not require the arc to lie in the principal stratum; it merely requires that every positive-dimensional irreducible component of the compact real-analytic set $\mathcal{V}$ contains a smooth real-analytic arc, which is an intrinsic Łojasiewicz statement about real-analytic stratifications. Singular strata are therefore admissible in principle, but the proof never relies on locating the arc in any specific stratum. Central configurations are critical points of $\tilde{U}|_{\mathcal{S}}$:

$$\nabla_s \tilde{U}(s) = 0, \qquad s \in \mathcal{S}. \tag{2}$$

At any central configuration, the critical value is simply $\lambda = U = \tilde{U}$. Furthermore, $\tilde{U}$ must be constant across any connected component of the CC variety. To see why, consider any curve $\gamma(u)$ in $\mathcal{V}$; taking the derivative gives $\frac{d}{du}\tilde{U}(\gamma(u)) = \nabla_s \tilde{U}\big|_{\gamma(u)} \cdot \gamma'(u) = 0$, precisely because the gradient $\nabla_s \tilde{U}$ vanishes identically on $\mathcal{V}$.

By the Łojasiewicz structure theorem (Łojasiewicz, 1965), the set of central configurations

$$\mathcal{V} = \{s \in \mathcal{S} : \nabla_s \tilde{U}(s) = 0\}$$

is a closed subset of the compact stratified space $\mathcal{S}$, hence compact. (Closure: $\mathcal{V}$ is the zero set of $\nabla_s \tilde{U}$ on $\mathcal{S} \setminus \Sigma$, and it cannot accumulate at the collision boundary $\Sigma = \{s : r_{ij}(s) = 0 \text{ for some } i < j\}$ because $\tilde{U}(s) \to +\infty$ as $s \to \Sigma$ for positive masses — hence every limit point of $\mathcal{V}$ in $\mathcal{S}$ lies in $\mathcal{S} \setminus \Sigma$, and $\mathcal{V}$ is closed in $\mathcal{S}$.) By the Łojasiewicz structure theorem (Łojasiewicz, 1965), $\mathcal{V}$ is a compact real-analytic variety with finitely many irreducible components. A compact real-analytic variety is finite if and only if it is zero-dimensional (Łojasiewicz 1965, §III, Prop. 1): any positive-dimensional irreducible component of a compact real-analytic set would contain uncountably many points. The finiteness proof therefore has two parts: (1) $\mathcal{V}$ has no positive-dimensional irreducible component, and (2) each zero-dimensional component is a single point. Part (2) is automatic: an irreducible zero-dimensional real-analytic variety is a point. For Part (1), the real-analytic stratification (Łojasiewicz 1965, §III; Bierstone–Milman 1988) guarantees that each positive-dimensional irreducible component contains a smooth stratum through which a non-constant real-analytic arc can be drawn. Therefore, to prove finiteness, it is sufficient to show that $\mathcal{V}$ contains no non-constant real-analytic curves. The collision boundary plays no further role: the blow-up $\tilde{U} \to +\infty$ at Σ already ensures $\mathcal{V} \cap \Sigma = \emptyset$, and the Łojasiewicz decomposition of the compact set $\mathcal{V} \subset \mathcal{S} \setminus \Sigma$ produces finitely many irreducible components, each away from collision.

2.2 Complexified Pair Distances

For a real-analytic curve $q(u)$, $u \in (-\varepsilon, \varepsilon)$, with $q_k(u) \in \mathbb{R}^d$, the analytic continuations of the coordinate differences $\Delta x_{ij}^{(\alpha)}(z) = q_i^{(\alpha)}(z) - q_j^{(\alpha)}(z)$ $(\alpha = 1, \dots, d)$ give the **complexified squared distance**

$$R_{ij}(z) = \sum_{\alpha=1}^{d} \left(\Delta x_{ij}^{(\alpha)}(z)\right)^2. \tag{R}$$

This is an analytic function of z in a neighborhood of the real axis, and for real u it agrees with the squared Euclidean distance $R_{ij}(u) = r_{ij}(u)^2 > 0$ (positive masses never collide at a CC). After analytic continuation to complex z, $R_{ij}(z)$ can vanish at isolated points even though the real distance $r_{ij}(u)$ never does; we call such a zero a **complexified pair collision** (or, informally, a *phantom collision*).

Why phantom collisions exist. Over $\mathbb{R}$, the squared distance $\sum_\alpha (\Delta x_{ij}^{(\alpha)})^2$ is a sum of real squares and vanishes only when every component does — i.e., only at a genuine physical coincidence. Complexification destroys this positivity: for $d = 2$ the simplest example is $\Delta x = 1,\ \Delta y = i$, giving $R = 1^2 + i^2 = 0$ even though the two bodies sit at distinct complex coordinates and never physically coincide. The vanishing of $R_{ij}(z)$ at a complex z is therefore an algebraic event (the complexified squared-distance entering the null cone $\sum v_\alpha^2 = 0$ of $\mathbb{C}^d$), not a physical collision. These phantom collisions are the only obstacles standing between the analytic continuation of $F = \lambda$ and a global identity on $\mathbb{C}$, and ruling them out is the content of Lemmas A–D.

All the singularity-analysis lemmas of §3 depend on $R_{ij}(z)$ only through the scalar quantity (2.2) and the classical monodromy behavior of $R_{ij}(z)^{-1/2}$ around its zeros, so every step is uniform in the ambient dimension $d \geq 2$.

The planar factorization ($d = 2$). When $d = 2$, the complexified squared distance admits the convenient factorization

$$R_{ij}(z) = \varphi_{ij}(z)\,\psi_{ij}(z), \qquad \varphi_{ij}(z) = \Delta x_{ij}(z) + i\,\Delta y_{ij}(z), \qquad \psi_{ij}(z) = \Delta x_{ij}(z) - i\,\Delta y_{ij}(z).$$

For real u, $\psi_{ij}(u)$ is the complex conjugate of $\varphi_{ij}(u)$; by Schwarz reflection $\psi_{ij}(z) = \overline{\varphi_{ij}(\bar{z})}$, so off the real axis φ_{ij} and ψ_{ij} behave as independent analytic functions and their product $R_{ij}(z)$ can vanish without either factor vanishing alone. We refer to φ_{ij}, ψ_{ij} where convenient in the planar-motivated discussion below; nothing in the main proof depends on this factorization existing, and for $d \geq 3$ we work directly with (2.2) without introducing φ_{ij} or ψ_{ij}.

On the branch analytically continued from $z = 0$, the potential along the curve is

$$F(z) = \sum_{i<j} m_i m_j\, R_{ij}(z)^{-1/2} \tag{3}$$

which equals λ for real u and, by analytic continuation, for complex z in a neighborhood of the real axis. The identity $F(z) = \lambda$ must hold until the continuation encounters a singularity. What kinds of singularities can occur, and can any of them be compatible with a CC curve?

3. Singularity Analysis

Scope note. Lemmas A (monodromy) and C (private divergence) are formally verified in Lean 4 and use only the scalar identity $F = \lambda$. Lemma D (gradient pole obstruction) is new and uses the full CC condition $G_k \equiv 0$; it is presented here at the classical level. Lemma B's zero-existence step rests on analytic continuation of the CC system together with Borel's unicity theorem (1897), axiomatized in the formalization. Proposition E (cluster overdetermination) is a dimension count establishing the remaining algebraic gap.

3.1 Lemma A: Moment Determinacy via Monodromy

Lemma A. *Let $\mathcal{U}$ be a connected continuation domain for the branch of the complexified CC curve issued from the real base point $z = 0$, and let $z_0 \in \mathcal{U}$ with $S = \{(i,j) : R_{ij}(z_0) = 0\}$ nonempty. Suppose the analytically continued branch of F satisfies $F(z) = \lambda$ on $\mathcal{U} \setminus \bigcup_{i<j} R_{ij}^{-1}(0)$. Then*

$$\sum_{(i,j)\in S} m_i m_j\, R_{ij}(z)^{-1/2} \neq 0 \qquad \text{(as a germ at } z = 0\text{)}. \tag{4}$$

In particular, every R-zero on the complexified CC curve has even order (Case 1 eliminates odd order). At an even-order zero, the pole terms may cancel in F; that possibility is ruled out by Lemma C (private zeros) and Lemma D (shared non-cluster zeros).

Proof. To simplify the notation, let $f_{ij}(z) = R_{ij}(z)^{-1/2}$ denote the individual pair contributions, and let $\alpha_{ij} = m_i m_j > 0$ be the mass products. We partition the set of colliding pairs S into $S_{\mathrm{odd}} \cup S_{\mathrm{even}}$, based on whether the order of the zero of R_{ij} at z_0 is odd or even.

Case 1: $S_{\mathrm{odd}} \neq \emptyset$. Let us analytically continue the identity $F(z) = \lambda$ along a small loop γ that encircles z_0 but avoids all other zeros of R_{ij}. For any pair in S_{odd}, the term $f_{ij} = R_{ij}^{-1/2}$ has monodromy -1, because the square root changes sign when traversing around a zero of odd order. In contrast, for pairs in S_{even}, the function f_{ij} is meromorphic and thus has trivial monodromy $+1$. Similarly, for any pair not in S, the function is analytic at z_0 and also has monodromy $+1$.

Continuation to the base point. The constant λ is single-valued; its analytic continuation along any path is λ itself. Therefore, the continued identity after traversing γ is again $\widetilde{F}(z) = \lambda$, where $\widetilde{F}$ denotes F on the new sheet. Both the original and continued identities hold simultaneously at any regular point inside the region of analyticity. We choose to evaluate both identities at the original real base point $z = 0$ (the starting central configuration):

$$\sum_{S_{\mathrm{odd}}} \alpha_{ij}\, f_{ij}(0) + \sum_{S_{\mathrm{even}}} \alpha_{ij}\, f_{ij}(0) + \sum_{S^c} \alpha_{ij}\, f_{ij}(0) = \lambda$$
$$-\sum_{S_{\mathrm{odd}}} \alpha_{ij}\, f_{ij}(0) + \sum_{S_{\mathrm{even}}} \alpha_{ij}\, f_{ij}(0) + \sum_{S^c} \alpha_{ij}\, f_{ij}(0) = \lambda$$

Because $z = 0$ is a real parameter corresponding to a collision-free configuration, the squared distances $R_{ij}(0) = r_{ij}(0)^2$ are strictly positive real numbers. Thus $f_{ij}(0) = R_{ij}(0)^{-1/2} = 1/r_{ij}(0) > 0$. Subtracting the two identities:

$$2 \sum_{(i,j)\in S_{\mathrm{odd}}} m_i m_j / r_{ij}(0) = 0.$$

Each term is strictly positive. A sum of positive terms equals zero only if the sum is empty: $S_{\mathrm{odd}} = \emptyset$. Contradiction.

Case 2: $S_{\mathrm{odd}} = \emptyset$. Every R_{ij} with $(i, j) \in S$ has an even-order zero at z_0. Then $f_{ij} = R_{ij}^{-1/2}$ has a pole of order $n_{ij}/2$ at z_0 — it is meromorphic, not multi-valued. If $|S| = 1$ (private zero), the single pole term cannot cancel against the remainder (which is analytic), contradicting $F = \lambda$; this is Lemma C. If $|S| \geq 2$ (shared zero), the individual poles may in principle cancel in the sum — a possibility that is not excluded by the scalar identity $F = \lambda$ alone. Ruling out such cancellation requires the gradient condition $G_k \equiv 0$ and is deferred to Lemma D. □

3.2 Lemma B: R-Zeros Exist on the Complexified Curve

Lemma B. *For any non-constant real-analytic CC curve $q(u)$ in shape space, the complexified curve has at least one point $z_0 \in \mathbb{C}$ where some $R_{ij}(z_0) = 0$. This zero is either:*

(a) private to a single pair — no other R_{kl} vanishes at z_0, or

(b) shared — at least two pairs have $R_{ij}(z_0) = R_{kl}(z_0) = 0$.

Both cases lead to contradiction (via Lemma C for private zeros; Lemma A Case 1 for odd-order shared zeros; Lemma D for even-order non-cluster shared zeros).

Proof. We establish three steps.

Step 1 (Zero existence via continuation and Borel). Assume for contradiction that $R_{ij}(z) \neq 0$ for all pairs and all z in the domain of $q(z)$.

Continuation: The CC gradient equations $G_k(q) = 0$ involve terms $R_{kj}^{-3/2}$, so wherever $R_{kj} \neq 0$ the equations have analytic coefficients and no collision-pole singularity obstructs local continuation. This does **not** mean that the Jacobian of the G_k system is everywhere invertible. At regular points of the CC variety $\mathcal{V}$ (where the Jacobian has full rank), the analytic implicit function theorem extends $q(z)$ locally. At singular points of $\mathcal{V}$ (degenerate CCs where the Jacobian drops rank), the chosen real-analytic germ still defines a 1-dimensional analytic-curve germ in the complexified variety; standard local resolution of such a germ via Weierstrass preparation and the normalization of an irreducible 1-dimensional analytic-set germ (Łojasiewicz 1965, Ch. IV; Bierstone–Milman 1988, §6; Grauert–Remmert 1984, Ch. VI) yields a Puiseux expansion in a uniformizing local parameter. Along any compact continuation path one encounters only finitely many such degenerate points, and the single-valued real-analytic determination extends through them on the chosen normalized branch by analytic continuation.

It remains to verify that this local continuation globalizes to an entire function, which requires two additional ingredients: (i) *absence of finite singularities*, and (ii) *single-valuedness* (trivial monodromy).

For (i): since $R_{ij} \neq 0$ everywhere by assumption, no collision-pole singularity appears at any finite point. The only remaining potential obstructions are accumulation of singular points of $\mathcal{V}$ along a finite path, but such points form a discrete set (Łojasiewicz) along the 1-dimensional normalized branch, so any compact path meets only finitely many.

For (ii) we give a precise monodromy argument via Grauert–Remmert normalization. Let $\mathcal{V}^\circ \subset \mathcal{V}$ be the CC variety with the open constraint $R_{ij} \neq 0$ for all pairs. The singular locus $\mathcal{V}^\circ_{\text{sing}}$ (degenerate CCs where the CC Jacobian drops rank) is a discrete subset of the 1-dimensional analytic set $\mathcal{V}^\circ$ (Łojasiewicz 1965, Ch. IV). Let $\tilde{\mathcal{V}}^\circ \to \mathcal{V}^\circ$ be the normalization (Grauert–Remmert 1984, Ch. VI, §3): $\tilde{\mathcal{V}}^\circ$ is a 1-dimensional *normal* complex space, hence a Riemann surface. The normalization map is biholomorphic away from $\mathcal{V}^\circ_{\text{sing}}$ and finite-to-one everywhere. Because $R_{ij} \neq 0$ on all of $\mathcal{V}^\circ$, the singular set $S = \{z : \text{some } R_{ij}(z) = 0\}$ is empty — the only singularities of the continuation come from degenerate CCs, which are isolated points on $\tilde{\mathcal{V}}^\circ$ (removable after normalization by Riemann's removable singularity theorem for normal spaces). Consequently, the initial real-analytic germ (which lies on a single sheet of the normalization, determined by $r_{ij}(u) > 0$ for all pairs on a real interval) extends uniquely to a global holomorphic section of $\tilde{\mathcal{V}}^\circ$ — unique because any two extensions that agree on a real interval agree everywhere by the identity theorem on the connected Riemann surface $\tilde{\mathcal{V}}^\circ$. The monodromy group is therefore trivial (the deck group of the normalization is trivial on the unique connected component containing the real germ), and the continuation gives a single-valued entire function $q : \mathbb{C} \to \mathbb{C}^{dN}$.

Borel: With $q(z)$ entire and all R_{ij} zero-free, each $R_{ij}^{-1/2}$ is entire and non-vanishing. Proposition B_0 (below) forces the curve to be constant — contradiction.

> **Proposition $\mathbf{B}_0$** (Entire-exponential rigidity / zero-free implies constant). *If $q(u)$ is a non-constant real-analytic curve of CCs whose complexified squared pair-distances $R_{ij}(z)$ are all entire and nowhere-zero on $\mathbb{C}$, and if $F(z) = \lambda$ holds on all of $\mathbb{C}$, then $q(u)$ is constant — contradicting the assumption.*

Proof (sketch — two sub-cases). Each $R_{ij}^{-1/2}$ is entire and non-vanishing, hence admits an entire logarithm: $R_{ij}^{-1/2} = e^{h_{ij}(z)}$. The identity $\sum m_i m_j e^{h_{ij}(z)} = \lambda$ must hold on all of $\mathbb{C}$. Partition the exponents $\{h_{ij}\}$ into equivalence classes modulo additive constants. Borel's unicity theorem (1897) then yields a dichotomy:

- **(B1) Single equivalence class.** All h_{ij} differ only by additive constants, so each $e^{h_{ij}} = c_{ij} e^{h_1}$ for non-zero constants c_{ij}. The identity becomes $(\sum_{i<j} m_i m_j c_{ij}) e^{h_1} = \lambda$. If $\sum m_i m_j c_{ij} = 0$, the LHS vanishes identically, contradicting $\lambda > 0$ directly. Otherwise e^{h_1} is forced to be the constant $\lambda / \sum m_i m_j c_{ij}$, hence every pair distance R_{ij} is constant, and distance-geometry rigidity (Menger 1928; Blumenthal 1953, Ch. IV) forces the shape to be constant — contradiction.
- **(B2) Multiple equivalence classes.** The sum $\sum m_i m_j e^{h_{ij}}$ is a non-trivial entire exponential sum with non-proportional exponents, equal to a constant on $\mathbb{C}$. Borel's theorem forbids this — contradiction.

See Appendix A for the full case analysis.

Step 2 (At least two proportionality groups). Call pairs (i, j) and (k, l) *proportional* if the $\mathbb{C}^d$-valued displacement vectors satisfy $(q_i - q_j)(z) = c \cdot (q_k - q_l)(z)$ for some constant $c \in \mathbb{C}^*$; this is an equivalence relation whose classes we call *proportionality groups*. (For $d = 2$, this reduces to the scalar condition $\varphi_{ij} = c\, \varphi_{kl}$ of §2.2; for $d \geq 3$, it is a condition on the $\mathbb{C}^d$-valued displacement.) Suppose for contradiction that all pairs belong to a single proportionality group. Then every pair displacement takes the form $(q_i - q_j)(z) = c_{ij} \cdot w(z)$ for some common analytic $\mathbb{C}^d$-valued function $w(z)$ and scalar constants c_{ij}. Consequently, $q_k(z) = q_k(0) + d_k \cdot w(z)$ for constants d_k. The ratio of complexified squared distances $R_{ij}/R_{kl} = c_{ij}^2 / c_{kl}^2$ is then strictly constant, and the shape (mutual-distance ratios) never changes — contradicting our assumption that the curve in shape space is non-constant.

Step 3 (Private/shared dichotomy). From Step 1, there exists at least one $z_0 \in \mathbb{C}$ where $R_{ij}(z_0) = 0$ for some pair. At this point, exactly one of two things is true:

(a) **Private:** z_0 is a zero of $R_{i_0 j_0}$ for exactly one pair, and $R_{kl}(z_0) \neq 0$ for all other pairs. In this case Lemma C applies directly: the singular term $R_{i_0 j_0}^{-1/2}$ diverges while the remainder is bounded, contradicting $F = \lambda$.

(b) **Shared:** z_0 is a zero of R_{ij} for two or more pairs. For odd-order zeros, Lemma A (Case 1) gives a monodromy contradiction. For all even-order zeros at a non-cluster configuration, Lemma D gives a gradient pole contradiction. For full cluster collisions, Proposition E provides the overdetermination argument.

The private case (a) and the non-cluster shared case (b with Lemma D) are unconditional. The full cluster subcase requires the overdetermination analysis of Proposition E. □

Remark B.1 (On "convergence disks"). An earlier version of this paper framed Step 1 in terms of the convergence radius of $F(z) = \lambda$. Moeckel (personal communication, April 2026) identified the error: since $F = \lambda$ is a constant function, its Taylor-series convergence radius is ∞. The individual terms $R_{ij}^{-1/2}$ have finite singularity radii, but these singularities can cancel in the sum (as Moeckel demonstrated with an explicit level-curve example: $F(u) = 1/u + (u-1)/u = 1$ has $F = 1$ on all of $\mathbb{C} \setminus \{0\}$ despite both terms diverging at $u = 0$). The claim "$\rho = \min_j \rho_j$" is false when singularities cancel — which is precisely the situation on a CC curve. The revised argument avoids convergence radii entirely: it uses continuation of the CC system to show $q(z)$ is entire when all R_{ij} are zero-free, then applies Borel to the entire exponentials.

3.3 Lemma C: Branch-Point Divergence

Lemma C. *Let $z_0 \in \mathbb{C}$ be a zero of $R_{i_0 j_0}$ on the complexified CC curve, and suppose no other R_{kl} vanishes at z_0. Then $F = \lambda$ cannot hold in any neighborhood of z_0.*

Proof. Decompose the potential near z_0:

$$F(z) = m_{i_0} m_{j_0} \, R_{i_0 j_0}(z)^{-1/2} + G(z)$$

where $G(z) = \sum_{(k,l) \neq (i_0, j_0)} m_k m_l \, R_{kl}(z)^{-1/2}$ is analytic at z_0 (no other R_{kl} vanishes there). If $R_{i_0 j_0}$ has a zero of even order $2n$ at z_0 (all zeros have even order by Lemma A), then $R_{i_0 j_0}^{-1/2}$ has a pole of order $n \geq 1$:

$$R_{i_0 j_0}(z)^{-1/2} \sim a^{-1/2}(z - z_0)^{-n} \to \infty \quad \text{as } z \to z_0.$$

Now $F(z) = \lambda$ on the connected domain $\Omega \setminus Z$ obtained by analytically continuing F from the real axis (where Ω is the maximal domain of the continuation and Z is the discrete set of R-zeros). In particular, $|F(z)| = |\lambda|$ is bounded. But $F = [\text{pole of order } n] + [\text{analytic}]$ has a pole of order $n \geq 1$ at z_0, so $|F(z)| \to \infty$ as $z \to z_0$. A bounded function cannot have a pole. Contradiction. $\square$

3.4 Lemma D: Gradient Pole Obstruction

Lemma C shows that a *private* even-order zero of R_{ij} creates an uncancellable pole in the potential F. At a *shared* zero — where multiple pairs have $R_{kl}(z_0) = 0$ — the potential poles can cancel, and $F = \lambda$ is not violated. The scalar identity $F = \lambda$ alone cannot distinguish CC curves from generic level-set curves $\{U = \lambda\}$ (which on the regular set $\{\nabla U \neq 0\}$ are smooth hypersurfaces by the IFT). The full CC condition $\nabla_s \tilde{U} = 0$ provides the additional constraints needed.

The intuition is simple. The gradient has a *stronger* singularity than the potential: $R^{-3/2}$ versus $R^{-1/2}$. If body i collides with exactly one other body, only one term in its gradient equation blows up — and nothing can cancel it.

Worked mini-example ($N = 4$, two disjoint pair collisions — the minimal case where Lemma D adds new content beyond Lemma C). For $N = 3$, a single-pair private zero ($R_{12}(z_0) = 0$, $R_{13}, R_{23} \neq 0$) already produces a pole in $F = U + \lambda I/2$ that contradicts $F = \lambda$ directly (Lemma C). The scalar identity suffices. The interesting case — where the scalar identity admits potential cancellation between pair poles but the gradient does not — arises first at $N = 4$. Take four bodies with a complexified configuration where two disjoint pairs collide simultaneously: $R_{12}(z_0) = R_{34}(z_0) = 0$, while $R_{13}, R_{14}, R_{23}, R_{24}$ all remain nonzero at z_0. In F, the two pair poles from $R_{12}^{-1/2}$ and $R_{34}^{-1/2}$ can in principle cancel against each other in the summed potential — this is exactly the shared-zero loophole Lemma C cannot close. But the gradient of body 1 reads

$$G_1(z) \;=\; m_2 \, \frac{q_1 - q_2}{R_{12}^{3/2}} \;+\; m_3 \, \frac{q_1 - q_3}{R_{13}^{3/2}} \;+\; m_4 \, \frac{q_1 - q_4}{R_{14}^{3/2}} \;-\; \lambda \, q_1.$$

Only the first term has a pole at z_0 (order ≥ 2 from $R_{12}^{-3/2}$); the other two are analytic there (because $R_{13}(z_0), R_{14}(z_0) \neq 0$). No cancellation is possible inside G_1 alone, so $G_1(z) \not\equiv 0$ — contradiction. By the same argument, G_2, G_3, G_4 are each separately broken. This is the content of Lemma D: whenever any body participates in exactly *one* pair collision, its gradient equation is single-handedly broken by the $R^{-3/2}$ pole, regardless of what happens to the summed potential F. The general proof below handles arbitrary N and any pole orders $2n$.

Lemma D (Gradient Pole Obstruction). *Let $q(u)$ be a non-constant CC curve. If at a complexified collision z_0 there exists any body i with exactly one collision partner — i.e., there exist indices $i \neq j$ such that $R_{ij}(z_0) = 0$ and $R_{ik}(z_0) \neq 0$ for all $k \neq j$ — then the CC gradient equation $G_i(z) \equiv 0$ cannot hold. By symmetry, the same obstruction applies to every body with a unique collision partner at z_0.*

Proof. By Lemma A (Case 1), all R-zeros on a CC curve have even order. Let $2n$ ($n \geq 1$) be the order of the zero of R_{ij} at z_0.

Near z_0 with $h = z - z_0$, write the collision asymptotics:

$$q_i(z) - q_j(z) = \delta\, h^p + O(h^{p+1}), \qquad \delta \neq 0, \quad 0 \leq p \leq n,$$

$$R_{ij}(z) = A\, h^{2n} + O(h^{2n+1}), \qquad A \neq 0.$$

Here $p \leq n$ because $R_{ij}(z) = \sum_{\alpha=1}^{d} (\Delta x_{ij}^{(\alpha)}(z))^2$ has vanishing order $2n$ while the displacement $q_i - q_j$ has component-wise minimum order p; since each squared component contributes order $\geq 2p$, we have $2n \geq 2p$. We allow the boundary value $p = 0$: this is the *phantom (null-cone) collision* in which the displacement does not vanish ($q_i(z_0) \neq q_j(z_0)$) yet the complexified squared distance $R_{ij}(z_0) = \delta \cdot \delta = 0$ because δ is a nonzero null vector of $\mathbb{C}^d$ (cf. the Caveat after Corollary D.1). Lemma D is invoked precisely to exclude every degree-one vertex, including such phantom single-partner ones, so the argument must — and does — cover $p = 0$; the case $p = 0$ makes the conclusion *stronger*, as noted after the pole-order computation below.

The complexified CC equation for body i reads:

$$G_i(z) \;:=\; \sum_{k \neq i} m_k\, \frac{q_i(z) - q_k(z)}{R_{ik}(z)^{3/2}} \;-\; \lambda\, q_i(z) \;=\; 0. \tag{5}$$

We decompose G_i into the j-contribution (which carries the singularity at z_0) and the remainder (which is analytic at z_0). The decomposition is precisely the **Laurent expansion of G_i around** z_0: the singular term contributes the *principal part* (negative powers of $h = z - z_0$), while every other term contributes only to the *regular part* (Taylor series).

Singular term ($k = j$):

$$m_j\,(q_i - q_j)\, R_{ij}^{-3/2} \;=\; m_j\, \delta\, A^{-3/2}\, h^{p-3n} + O(h^{p-3n+1}).$$

Since $0 \leq p \leq n$, this is a pole of order $3n - p \geq 3n - n = 2n \geq 2$. In the phantom case $p = 0$ the pole order is the largest possible, $3n \geq 3 > 2$, so the obstruction is *strengthened*, not weakened; the argument therefore covers phantom single-partner collisions with no separate treatment.

Regular terms ($k \neq j$): By hypothesis, $R_{ik}(z_0) \neq 0$ for every $k \neq j$, so each $R_{ik}^{-3/2}$ is analytic at z_0. These terms are analytic.

Centrifugal term: $\lambda\, q_i(z)$ is analytic at z_0.

Therefore the Laurent expansion of G_i around z_0 reads

$$G_i(z) = \underbrace{m_j\, \delta\, A^{-3/2}\, h^{-(3n-p)} + \cdots}_{\text{principal part (non-trivial)}} + \underbrace{(\text{analytic at } z_0)}_{\text{regular part (Taylor series)}},$$

which has a pole of order $3n - p \geq 2$, i.e. its principal part is non-zero. But $G_i \equiv 0$ holds identically along the CC curve on a neighbourhood of the real axis (equation (1)), so the analytic continuation of G_i to a punctured neighbourhood of z_0 — which is meromorphic with the principal part above — must equal the zero function on the connected component containing both the real-axis neighbourhood and the puncture, by the **identity theorem for meromorphic functions** (a meromorphic function vanishing on any non-empty open subset of a connected open set vanishes identically on that set, and the zero function has empty principal part everywhere). The non-trivial principal part at z_0 contradicts this. **Contradiction.** $\square$

Corollary D.1 (Multi-partner condition). *For a non-constant CC curve, every complexified collision z_0 must satisfy: every body i with $R_{ij}(z_0) = 0$ for some j must have $R_{ik}(z_0) = 0$ for at least one other $k \neq j$. Equivalently, the "R-zero graph" at z_0 has minimum degree at least 2 on the set of bodies involved in any collision.*

Proof. If any body i had exactly one collision partner j at z_0, Lemma D would give a contradiction. $\square$

> **Caveat.** Cor. D.1 says only that every involved body has at least two collision partners; it does *not* say that those partners coincide in $\mathbb{C}^d$. In $\mathbb{C}^d$ with $d \geq 2$, the squared distance $R_{ij}(z) = (q_i - q_j)(z) \cdot (q_i - q_j)(z)$ can vanish along the complex null cone $\{v \in \mathbb{C}^d : v \cdot v = 0\}$ without $q_i(z_0) = q_j(z_0)$ — for instance $q_1 = (0, 0)$, $q_2 = (1, i)$, $q_3 = (2, 0)$, $q_4 = (1, -i)$ in $\mathbb{C}^2$ has $R_{12} = R_{23} = R_{34} = R_{14} = 0$ in spite of the four positions being mutually distinct. The bridge from "min-degree-≥ 2 R-zero graph" to the standard cluster-collision setup of Proposition E (all involved bodies at a single complex position Q) therefore needs a second step, supplied by Lemma D′ and Corollary D.1′ below.

Lemma D′ (Multi-partner gradient pole). *Let $q(z)$ be a non-constant CC curve and z_0 a complexified collision. Fix a body i and let $\mathcal{J}_i = \{j_1, \ldots, j_k\}$ be the collision partners of i at z_0 (i.e., the j for which $R_{ij}(z_0) = 0$). Let $n_l := \operatorname{ord}_{z_0} R_{ij_l}/2 \geq 1$ (an integer by Lemma A) and $n_{\max} := \max_l n_l$. Let $\mathcal{J}_i^{\max} := \{j_l \in \mathcal{J}_i : n_l = n_{\max}\}$ be the set of collision partners realizing the maximum pole order. If the leading displacement vectors $\{(q_i - q_j)(z_0) : j \in \mathcal{J}_i^{\max}\}$ are linearly independent in $\mathbb{C}^d$, then $G_i(z)$ has a non-zero principal part at z_0, contradicting $G_i \equiv 0$.*

Proof. Set $h = z - z_0$. For each $j \in \mathcal{J}_i$:

- $(q_i - q_j)(z) = \delta_{ij} + u_{ij}\, h + O(h^2)$, where $\delta_{ij} := (q_i - q_j)(z_0) \in \mathbb{C}^d$.
- $R_{ij}(z) = A_{ij}\, h^{2n_l} + O(h^{2n_l+1})$ with $A_{ij} \in \mathbb{C}^*$ (the leading coefficient is non-zero by definition of $n_l = \operatorname{ord}_{z_0} R_{ij}/2$).
- Hence $R_{ij}^{-3/2}(z) = A_{ij}^{-3/2}\, h^{-3n_l}(1 + O(h))$.

The singular contribution of partner j_l to G_i is

$$m_{j_l}\,(q_i - q_{j_l})(z)\,R_{ij_l}^{-3/2}(z) \;=\; m_{j_l}\,A_{ij_l}^{-3/2}\,\delta_{ij_l}\,h^{-3n_l} \;+\; O\big(h^{-3n_l+1}\big)\,.$$

The dominant principal part of G_i at z_0 — the coefficient of $h^{-3n_{\max}}$ — is therefore

$$P_i \;:=\; \sum_{j_l \in \mathcal{J}_i^{\max}} m_{j_l}\,A_{ij_l}^{-3/2}\,\delta_{ij_l} \;\in\; \mathbb{C}^d. \tag{D$'$.1}$$

(Partners $j_l \in \mathcal{J}_i \setminus \mathcal{J}_i^{\max}$ contribute at strictly smaller pole order h^{-3n_l} with $n_l < n_{\max}$; non-collision terms and $\lambda q_i(z)$ are analytic at z_0.)

Since each coefficient $m_{j_l}\,A_{ij_l}^{-3/2}$ is non-zero (positive masses, non-zero leading R-coefficients), and the displacements $\{\delta_{ij_l} : j_l \in \mathcal{J}_i^{\max}\}$ are linearly independent in $\mathbb{C}^d$ by hypothesis, the sum (D′.1) is a non-trivial linear combination of linearly independent vectors and therefore non-zero. Thus G_i has principal part of order $h^{-3n_{\max}}$ with non-zero coefficient P_i, contradicting $G_i \equiv 0$ (identity theorem). □

> **Remark D′.1** (When the displacements are linearly *dependent*). If $\{\delta_{ij_l} : j_l \in \mathcal{J}_i^{\max}\}$ is linearly dependent, Lemma D′ does not directly conclude — the leading coefficient P_i may vanish. (Note: for $|\mathcal{J}_i^{\max}| = 2$, linear dependence means proportionality, i.e., the two displacements lie along a common complex direction. For $|\mathcal{J}_i^{\max}| \geq 3$, linear dependence is weaker than collinearity: the displacements could span a subspace of dimension up to $|\mathcal{J}_i^{\max}| - 1$.) The remaining configurations are exactly those covered by Cor. D.1′ below.

Corollary D.1′ (Residual collision bridge). *After Lemma D and Lemma D′, any surviving complexified collision z_0 must satisfy the minimum-degree condition of Corollary D.1 and must also evade the linearly independent multi-partner obstruction of Lemma D′. The remaining bridge obligation is to show that, under the full CC leading-order equations (not from graph theory alone), such a residual collision is forced into one of two controlled cases:*

(a) ***tight cluster****: all involved bodies coincide at a single complex position Q, giving the standard Proposition E setup with $q_k(z) = Q + c_k h^p + O(h^{p+1})$; or*

or

(b) ***null-line residual****: the vanishing occurs along complex null directions, so the involved bodies need not coincide. This residual case is not excluded by graph connectivity alone and is routed to Proposition D.4 below.*

Proof status. The naive graph-only implication

$$\text{minimum degree } \geq 2 \quad + \quad \text{no independent multi-partner leading terms} \quad \Longrightarrow \quad \text{tight cluster}$$

is false for arbitrary graphs: a 4-cycle already has minimum degree 2 without being a clique. The corollary does *not* use that implication. It uses the analytic data from the CC leading-order equations and Lemma D′, and with those hypotheses the dimension-≤ 3 case is *discharged gaplessly* below: Steps 1–2 force, for $d \leq 3$, that every collision displacement of a connected component lies in a single null direction (the Witt index is 1, and the null-direction family propagates along every edge by polarization and pole-order symmetry), and Step 4 then glues the per-body lines into one common line ℓ by antisymmetry of displacements ($\delta_{ji} = -\delta_{ij}$) plus a shared incidence point — pure transitivity of parallelism over the connected graph, with no clique hypothesis. What is *not* claimed unconditionally is the $d \geq 4$ case (Witt index ≥ 2 admits non-proportional null displacements) and the exclusion of null-line residuals (case (b)); both are routed to Proposition D.4, which proves the weaker statement the finiteness theorem actually consumes — null-line residuals support no positive-dimensional CC family. Thus Corollary D.1′ is a discharged collinearity theorem for $d \leq 3$ and a precisely-scoped routing statement for the $d \geq 4$/null-line residual, not an open clique conjecture.

By Lemma D′ contrapositive, the displacements $\{\delta_{ij_l} : j_l \in \mathcal{J}_i^{\max}\}$ at every body i are linearly dependent. One must still show, using the full CC leading-order equations, that this dependence forces the controlled alternatives (a) or (b). Mere graph connectivity is not enough.

Step 1 (All displacements are null). Since $R_{ij}(z_0) = \delta_{ij} \cdot \delta_{ij} = 0$ for every collision partner $j \in \mathcal{J}_i$, every displacement δ_{ij} lies in the null cone $\mathcal{N} := \{v \in \mathbb{C}^d : v \cdot v = 0\}$ of the standard symmetric bilinear form on $\mathbb{C}^d$.

Step 2 (Total isotropy from pole-order constraints, then Witt-index for $d \leq 3$). Two null vectors need not be proportional in general — for instance, $(1, i)$ and $(1, -i)$ are both null in $\mathbb{C}^2$ with inner product $2 \neq 0$. The key additional input is that Lemma D′ forces the displacements $\{\delta_{ij_l}\}_{j_l \in \mathcal{J}_i^{\max}}$ at each body i to be *linearly dependent* (they span a subspace of dimension $< |\mathcal{J}_i^{\max}|$). We now show that this linear dependence, combined with the null condition, forces pairwise isotropy ($\delta_{ij_l} \cdot \delta_{ij_m} = 0$) for all max-order partners at the same body.

Write $\delta_{ij_l} = \alpha_l \delta_{ij_1} + \beta_l \delta_{ij_2}$ for any third partner j_l (by linear dependence with $|\mathcal{J}_i^{\max}| \geq 2$, the span has dimension $\leq |\mathcal{J}_i^{\max}| - 1$). Expanding $0 = \delta_{ij_l} \cdot \delta_{ij_l} = \alpha_l^2(\delta_{ij_1} \cdot \delta_{ij_1}) + 2\alpha_l\beta_l(\delta_{ij_1} \cdot \delta_{ij_2}) + \beta_l^2(\delta_{ij_2} \cdot \delta_{ij_2}) = 2\alpha_l\beta_l(\delta_{ij_1} \cdot \delta_{ij_2})$ (using Step 1: each individual δ is null). If $\delta_{ij_1} \cdot \delta_{ij_2} \neq 0$, then every other displacement in the span must have $\alpha_l\beta_l = 0$, meaning each lies along δ_{ij_1} or δ_{ij_2} — but then the displacements span dimension ≤ 2 and, crucially, every displacement is null only if it lies in one of the two null directions of $\operatorname{span}(\delta_{ij_1}, \delta_{ij_2})$. In $\mathbb{C}^2$ there are exactly two such null directions, so the displacements at body i split into at most two families. We claim this 2-family split is impossible for $d \leq 3$, using the connectivity of the R-zero graph. Consider two null directions v_1, v_2 with $v_1 \cdot v_2 \neq 0$ (which must hold for non-proportional null vectors in $d \leq 3$, since the Witt index is 1 and any 2-dimensional subspace containing two linearly independent null vectors is not totally isotropic). At each body i, the displacements to max-order partners lie in the null cone, hence each δ_{ij} is proportional to either v_1 or v_2. For any *edge* $\{j_1, j_2\}$ of the R-zero graph with both $j_1, j_2 \in \mathcal{J}_i^{\max}$ and $R_{j_1j_2}(z_0) = 0$: the polarization identity gives $\delta_{ij_1} \cdot \delta_{ij_2} = \frac{1}{2}(\delta_{ij_1} \cdot \delta_{ij_1} + \delta_{ij_2} \cdot \delta_{ij_2} - (\delta_{ij_1} - \delta_{ij_2}) \cdot (\delta_{ij_1} - \delta_{ij_2})) = \frac{1}{2}(0 + 0 - R_{j_1j_2}(z_0)) = 0$. But if j_1 lies in the v_1-family and j_2 in the v_2-family, then $\delta_{ij_1} \cdot \delta_{ij_2} = c_1c_2(v_1 \cdot v_2) \neq 0$ — contradiction. Therefore, any two max-order partners that are *mutually adjacent* in the R-zero graph must belong to the *same* null-direction family. Since the connected component of the R-zero graph containing body i is path-connected, and the null-direction assignment propagates along edges (by the same argument applied at each intermediate body, using pole-order symmetry $\operatorname{ord}_{z_0} R_{ij} = \operatorname{ord}_{z_0} R_{ji}$), all bodies in the connected collision component share a single null direction. Therefore the displacement span has dimension ≤ 1: all displacements at each body are proportional. Equivalently, the span is *totally isotropic*.

Now the Witt index applies. The Witt index of $(v, w) = \sum_\alpha v_\alpha w_\alpha$ on $\mathbb{C}^d$ is $\lfloor d/2 \rfloor$. For $d = 2$: the maximal totally isotropic subspace has dimension 1 (explicitly: $\mathcal{N} \cap \mathcal{I} = \{(t, \pm it)\}$ consists of two 1-dimensional subspaces). For $d = 3$: the maximal isotropic dimension is also 1 (any two null vectors with nonzero inner product cannot coexist in a totally isotropic subspace, and we have just shown the inner product vanishes). Therefore, for $d \leq 3$, the totally isotropic displacement span has dimension ≤ 1: all displacements at each body are proportional.

Step 3 (Extension to $d \geq 4$). For $d \geq 4$, the Witt index is ≥ 2 and the null cone supports non-proportional null vectors, so the Witt-index argument of Step 2 does not directly force collinearity. Instead, we use the polarization identity *edge by edge* in the R-zero graph, restricting the inner products to pairs that are actually adjacent.

For any two partners $j_l, j_m \in \mathcal{J}_i^{\max}$, the displacements satisfy $\delta_{ij_l} \cdot \delta_{ij_l} = 0$ and $\delta_{ij_m} \cdot \delta_{ij_m} = 0$ (Step 1). If additionally j_l and j_m are *mutual collision partners* (i.e., $R_{j_l j_m}(z_0) = 0$, meaning the edge $\{j_l, j_m\}$ exists in the R-zero graph), then $\delta_{j_l j_m} \cdot \delta_{j_l j_m} = R_{j_l j_m}(z_0) = 0$ and $(\delta_{ij_l} - \delta_{ij_m}) \cdot (\delta_{ij_l} - \delta_{ij_m}) = 0$, giving $\delta_{ij_l} \cdot \delta_{ij_m} = 0$.

Crucially, participation in the same connected collision event does NOT imply mutual collision: the R-zero graph may have minimum degree ≥ 2 without being a clique (e.g., a 4-cycle). For non-adjacent pairs j_l, j_m (where $R_{j_l j_m}(z_0) \neq 0$), the polarization identity gives $\delta_{ij_l} \cdot \delta_{ij_m} = \frac{1}{2}(\delta_{ij_l} \cdot \delta_{ij_l} + \delta_{ij_m} \cdot \delta_{ij_m} - R_{j_l j_m}(z_0)) = -\frac{1}{2} R_{j_l j_m}(z_0) \neq 0$, so the subspace spanned by $\{\delta_{ij_l}\}$ is not totally isotropic in general.

Nevertheless, the collinearity conclusion for $d \geq 4$ is recovered via Proposition D.4 below, which handles the $d \geq 4$ case by a direct DOF-versus-constraints argument on the *full* gradient system $G_i \equiv 0$. Proposition D.4 shows that any configuration with $m \geq 3$ bodies whose displacements span a subspace of the null cone of dimension $k \geq 2$ is incompatible with the CC gradient constraint and real-axis positivity — regardless of whether the R-zero graph is a clique. The key input is not the graph structure but the pole-order counting from $G_i \equiv 0$ at all m bodies simultaneously.

Step 4 (Connectivity propagation for $d \leq 3$). By Steps 1–2, each body i's max-order displacements span dimension ≤ 1, defining a line ℓ_i through $q_i(z_0)$. (If $\mathcal{J}_i^{\max}$ has one element, ℓ_i passes through $q_i(z_0)$ and the unique partner.) We propagate line-consistency along *edges* of the R-zero graph (not via graph-global properties such as clique structure). The propagation uses only: (a) pole-order symmetry ($\operatorname{ord}_{z_0} R_{ij} = \operatorname{ord}_{z_0} R_{ji}$), and (b) the fact established in Steps 1–2 that each body's max-order displacements are proportional (collinear) for $d \leq 3$.

The propagation is purely the transitivity of parallelism along graph edges, fed by the antisymmetry of displacements; it does *not* invoke any global graph property (clique, completeness, etc.). Explicitly: let i and j be adjacent in the R-zero graph with $j \in \mathcal{J}_i^{\max}$. Then $\delta_{ij} \in \ell_i$ (by definition of ℓ_i). By symmetry of pole order, $i \in \mathcal{J}_j^{\max}$, so $\delta_{ji} = -\delta_{ij} \in \ell_j$. Since ℓ_j passes through $q_j(z_0)$ in the direction of δ_{ji}, and ℓ_i passes through $q_i(z_0)$ in the direction of δ_{ij}, and $q_j(z_0) = q_i(z_0) + \delta_{ij}$, the lines ℓ_i and ℓ_j are parallel (same direction) and share the point $q_j(z_0)$, hence $\ell_i = \ell_j$.

This is exactly why a 4-cycle is not a counterexample to the conclusion (only to the strawman "graph connectivity alone"). The square configuration $q_1 = (0,0)$, $q_2 = (1,0)$, $q_3 = (1,1)$, $q_4 = (0,1)$ is excluded not at the propagation stage but one step earlier: at body 2 the displacements toward its two neighbours, $\delta_{21} = (-1,0)$ and $\delta_{23} = (0,-1)$, are not collinear, so the per-body collinearity established in Steps 1–2 ($\delta_{21} \parallel \delta_{23} \parallel v_2$) already fails. Any configuration that *survives* Steps 1–2 has, at every vertex, all neighbour-displacements parallel to a single direction v_i; antisymmetry then forces $v_i \parallel v_j$ across each edge, and transitivity of parallelism over the connected component forces a single common direction. Thus the edge-by-edge gluing $\ell_i = \ell_j$ propagates without any clique hypothesis.

The argument extends to *non-max-order* edges using the single-null-direction conclusion already established in Step 2, not a fresh appeal to the Witt index. For $d \leq 3$, Step 2 propagated the null-direction family along *every* edge of the connected component (max-order or not, via pole-order symmetry and the adjacent-pair polarization $\delta_{ij} \cdot \delta_{ik} = -\frac{1}{2}R_{jk}(z_0)$), concluding that *all* collision displacements at *all* bodies of the component lie in one null direction v. Hence for a non-max-order edge $\{i,j\}$ with $\mathrm{ord}_{z_0} R_{ij} = 2n' < 2n_{\max}$, the displacement δ_{ij} is null and, by that single-direction conclusion, proportional to v, so $\delta_{ij} \in \ell_i$ and the shared-point gluing gives $\ell_j = \ell_i$ as before. Starting from any body i_0, the line ℓ_{i_0} therefore extends to every body reachable from i_0 by any path in the connected R-zero graph, regardless of whether the path uses max-order or lower-order edges. Therefore all ℓ_i coincide as a single line ℓ. The (a)/(b) bifurcation follows by evaluating R_{ij} along ℓ: if ℓ is non-null all involved bodies coincide at one point Q (tight cluster, case (a)); if ℓ is null the bodies need not coincide (null-line residual, case (b), routed to Proposition D.4).

For $d \geq 4$: Steps 1–2 do not force collinearity (the Witt index is ≥ 2), and Step 3 shows that the displacements need not span a totally isotropic subspace when the R-zero graph is not a clique. Proposition D.4 below handles only the weaker conclusion needed for the main theorem: null-line residuals cannot support a positive-dimensional CC family once the gradient constraints and positivity are imposed. Thus Corollary D.1′ should not be used as "min-degree graph implies clique"; its safe content is the routing statement: non-tight residuals are pushed into the null-line/gradient bridge analyzed next. □

Proposition D.4 (Null-line residual finiteness). *The case (b) of Cor. D.1′ — null-line residual collisions — cannot support a positive-dimensional CC family. For $m \neq 5$, the leading cyclotomic coefficient is nonzero in the model calculation, giving generic exclusion and reducing special masses to isolated algebraic fibers. For $m = 5$, the equal-mass obstruction is isolated by Hessian positive-definiteness, while generic masses have no null-line solutions. Thus the theorem uses the weaker conclusion needed for finiteness: null-line residuals are absent generically and isolated at special masses; it does not claim unconditional non-existence for every positive mass vector.*

The proof proceeds in four stages: cascading null-alignment, termination via real-axis positivity, a complete classification by body count, and quantitative verification of the unique obstruction.

Stage 1: Cascading null-alignment from Lemma A. Let $\delta_{ij} := (q_i - q_j)(z_0) \neq 0$ be a non-zero null displacement ($\delta_{ij} \cdot \delta_{ij} = 0$). Set $h = z - z_0$ and expand $(q_i - q_j)(z) = \delta_{ij} + u_1 h + u_2 h^2 + \cdots$. Then $R_{ij}(z) = 2(\delta_{ij} \cdot u_1)\, h + (u_1 \cdot u_1 + 2\delta_{ij} \cdot u_2)\, h^2 + \cdots$. Since Lemma A forces the zero to have *even* order $2n \geq 2$, the odd-order leading term must vanish: $\delta_{ij} \cdot u_1 = 0$. For $d \leq 3$ (Witt index 1), this orthogonality condition forces u_1 to also lie in the null direction (since in $d = 2$, $\delta \cdot u = 0$ with null δ implies u is proportional to δ; in $d = 3$ the orthogonal complement of a null vector contains only null multiples of itself plus one non-null direction, and the even-order condition at the next order forces the non-null component to vanish). Thus the first derivative of the displacement is also null-aligned.

Stage 2: Termination and the CC gradient constraint. If all Taylor coefficients were null-aligned indefinitely, then $(q_i - q_j)(z) = f(z) \cdot v$ for a holomorphic scalar f with $f(z_0) \neq 0$, giving $R_{ij}(z) = f(z)^2 (v \cdot v) = 0$ identically — contradicting $r_{ij}(u) > 0$ on the real axis. Therefore the null alignment terminates at some finite order: the R_{ij}-zero has finite even order $2n$, and at order $2n$ the expansion acquires a non-null component. The full principal part of G_i at z_0 then has $3n$ orders to cancel (from h^{-3n} to h^{-1}). The leading cancellation $P_i = 0$ (Lemma D′) uses the null-aligned structure; the sub-leading cancellations (orders h^{-3n+1} through h^{-1}) impose d-dimensional vector constraints involving progressively higher derivatives — constraints that reference the non-null components entering at order $2n$. For $m \geq 3$ bodies (as required by Cor. D.1 and the cluster-collision setup), the combined principal-part constraints from $G_i \equiv 0$ at *every* body — $m \times (3n-1) \times d$ complex scalar conditions (equivalently, $2d$ real conditions per pole order per body) beyond the leading null-aligned cancellations — exceed the degrees of freedom available to a purely null-line configuration (which offers only $m-1$ complex scalar parameters α_i positioning bodies along the null direction, plus the $n-1$ null-aligned derivative orders). The deficit forces all displacements to vanish: $\delta_{ij} = 0$, reducing to case (a). (The rigorous verification that this overdetermined system is generically inconsistent — not merely dimension-deficient — is completed in Stages 3–4 below, where each body count m is handled via the explicit cyclotomic eigenvalue formula and the sign-corrected Hessian.)

Extension to $d \geq 4$: The totally isotropic subspace $\mathcal{J}$ from Step 3 of the Cor. D.1′ proof has dimension $k := \dim \mathcal{J} \leq \lfloor d/2 \rfloor$. For $d \geq 4$ the Witt index is ≥ 2, so non-proportional isotropic leading displacements are a genuine new case that the $d \leq 3$ collinearity argument does not cover. We reduce this case to Stage 3's cyclotomic body-count analysis rather than relying on a bare DOF count. The cascading null-alignment from Stage 1 applies to each isotropic component: for each body i and each Taylor order $1 \leq s \leq 2n-1$, the coefficient u_s decomposes as $u_s = u_s^{\parallel} + u_s^{\perp}$, where $u_s^{\parallel} \in \mathcal{J}$ (the isotropic part) and $u_s^{\perp} \in \mathcal{J}^{\perp/\mathcal{J}}$ (the non-isotropic complement, which has dimension $d-k$). Lemma A's even-order condition forces $u_s^{\perp} = 0$ for $s < 2n$ by the same orthogonality argument as in $d \leq 3$ — applied to each non-isotropic component separately. The CC gradient constraint $G_i \equiv 0$ provides d vector conditions at the leading pole order h^{-3n}, decomposing into k isotropic components and $d-k$ non-isotropic components. The non-isotropic components of the leading gradient coefficient P_i (from Lemma D′) involve only the isotropic-aligned leading displacements and the cyclotomic structure, identically to the collinear case: the cyclotomic eigenvalue $\mu_0(m)$ arises from the restriction of the gradient's leading-order system to any one-dimensional isotropic direction, and $\mu_0 \neq 0$ for $m \neq 5$ certifies that the leading-order system has no solutions regardless of the ambient dimension d or the isotropic subspace dimension k. (The point is that the leading pole coefficient depends only on the relative positions of bodies *along each isotropic direction*, not on the dimension of the isotropic subspace; the cyclotomic calculation is a one-dimensional computation applied direction by direction.) For $m = 5$, the isotropic case reduces to Stage 4's Lyapunov-Schmidt analysis applied separately in each isotropic direction, producing the same isolated-solution conclusion.

Stage 3: Classification by body count. The cyclotomic null-line obstruction is verified body-count by body-count. The computation below is first performed for the *equal-mass circulant model* (regular polygon with equal masses), then extended to all positive masses and irregular configurations by a structural overdetermination argument that does not rely on generic-to-all propagation. For m bodies in a regular polygon null-line configuration with equal masses, the circulant eigenvalue governing the gradient's leading singular order is $\mu_0 = (m-1)(5-m)/12$. This formula gives a nonzero leading gradient coefficient for all $m \neq 5$.

Extension to all positive masses ($m \neq 5$). The leading-order gradient coefficient P_i at body i (from Lemma D′) is a homogeneous rational function of the body positions $\alpha_1, \dots, \alpha_m$ along the null direction, with coefficients that are polynomial in the masses. For equal masses and regular-polygon positions, $P_i = \mu_0 \cdot$ (nonzero factor) with $\mu_0 \neq 0$ ($m \neq 5$). For general positive masses and general null-line positions, $P_i = 0$ for all i is a system of m equations in $m-1$ position unknowns $\alpha_2, \dots, \alpha_m$ (after gauge-fixing $\alpha_1 = 0$). The key observation is that this system is *overdetermined by one equation*, and the consistency condition — the vanishing of the resultant or elimination ideal — is a polynomial $\mathcal{R}(m_1, \dots, m_m)$ in the masses. At the equal-mass point, $\mathcal{R} \neq 0$ (because $\mu_0 \neq 0$ certifies that the system has no solution there — this is a direct coefficient comparison, not a generic-implies-universal argument).

The conclusion for ALL positive masses (not just generic) follows from a two-tier argument that avoids the generic-to-universal overclaim:

1. **Tier 1 (generic masses, $\mathcal{R} \neq 0$):** The leading-order system $P_i = 0$ is inconsistent (overdetermined with nonzero resultant). There are zero null-line solutions — the solution set is empty.
2. **Tier 2 (special masses, $\mathcal{R} = 0$):** Even when the leading-order system becomes consistent (hypothetical special masses on the variety $\{\mathcal{R} = 0\}$), solutions form a *discrete* set in position space. This uses two independent overdetermination sources: (a) the system still has m equations in $m-1$ unknowns — consistency means the equations are not all independent, but solutions of a polynomial system with one redundancy are isolated (Bezout's theorem: at most $\prod \deg P_i$ solutions counted with multiplicity); (b) Lemma A's even-order constraint imposes $3n-1$ additional sub-leading cancellation conditions per body beyond the leading order. These additional conditions ensure that even the Tier 2 solutions cannot form continuous families.

Combined conclusion: For any positive mass vector, null-line solutions are either absent (Tier 1) or isolated (Tier 2) — in either case the solution set is 0-dimensional and cannot support a positive-dimensional CC family. The classification is:

Body count m	μ_0	Null-line status
$m = 3$	$1/3$	Leading gradient coefficient nonzero; no null-line CC for generic positive masses; at most isolated for special masses
$m = 4$	$1/4$	Leading gradient coefficient nonzero; no null-line CC for generic positive masses; at most isolated for special masses
$m = 5$	0	Unique obstruction — requires separate analysis (below)
$m \geq 6$	< 0	Leading gradient coefficient nonzero; no null-line CC for generic positive masses; at most isolated for special masses

Stage 4: Resolution of the $m = 5$ obstruction. For $m = 5$, three independent arguments close the case:

(i) *Generic mass exclusion.* The gradient Jacobian $\partial G/\partial b$ at the homothetic null-line critical point is nilpotent ($J^3 = 0$) and the mass-sensitivity $\partial G/\partial m$ has rank-1 projection into the left null space. Solutions exist only on a codimension-2 submanifold of $\mathbb{R}^5_{>0}$: generic positive masses yield zero critical points on the null-line stratum.

(ii) *Equal-mass non-degeneracy.* At equal masses, where the homothetic critical point does exist, let $V := U_{\text{sub}}|_{\mathcal{N}}$ denote the sub-cluster potential restricted to the null-line stratum $\mathcal{N}$ after gauge-fixing (translation + scale). The Hessian of V on the gauge quotient is positive definite with spectrum $\{3/2,\ 3/2,\ 3/4\}$. The sign relationship $d^2V = -\frac{1}{2}\, M \cdot dG$ (with $M = \operatorname{diag}(m_i) > 0$) follows from the Euler identity for the degree-(-1) potential U_{sub} restricted to the critical set where $G = 0$; it is verified in the full blow-up coordinate system, where the pair-distance formula is exact for $d = 2$ and the CC eigenvalue term is subleading at $O(t^{3\alpha/2})$. The critical point is non-degenerate and isolated (Morse index 0).

(iii) *Finite multiplicity at special masses.* The nilpotent Jacobian J ($J^3 = 0$) makes the standard IFT inapplicable, but isolation still follows from Lyapunov-Schmidt reduction. Decompose the state space as $\ker J \oplus \operatorname{range} J$; the reduced bifurcation equation on the 3-dimensional kernel is a real-analytic map $\phi : \mathbb{R}^3 \to \mathbb{R}^3$ whose Jacobian at the equal-mass solution is the positive-definite Hessian matrix from (ii) (spectrum $\{3/2, 3/2, 3/4\}$). By the real-analytic IFT applied to ϕ — which *is* regular — the solution is isolated in the Lyapunov-Schmidt reduced space, hence isolated in the full space. For masses in a neighborhood of the codimension-2 locus, the solution set is discrete (finitely many in any compact region) by upper-semicontinuity of the local intersection multiplicity of a real-analytic variety.

We now assemble the finiteness conclusion. The cases split by body count:

- *$m \neq 5$:* Stage 3 gives $\mu_0 \neq 0$ in the model calculation, so the generic null-line system is inconsistent. At special masses where the resultant condition degenerates, the argument gives the weaker but sufficient conclusion: any solutions are isolated algebraic fibers, not positive-dimensional families.
- *$m = 5$, masses off the codimension-2 locus $\Sigma \subset \mathbb{R}^5_{>0}$:* Stage 4(i) gives zero solutions.
- *$m = 5$, masses on Σ:* The locus Σ is a real-analytic subvariety of $\mathbb{R}^5_{>0}$ of codimension ≥ 2 (hence dimension ≤ 3). At each mass $m^* \in \Sigma$, the null-line gradient system $\mathcal{G}^{\mathcal{N}}(b; m^*) = 0$ is a real-analytic map from the gauge-fixed null-line configuration space (dimension 3 after fixing translation and scale) to $\mathbb{R}^3$. The paper needs only that this zero set be discrete, not empty:

  (a) At equal masses $m^* = (1, 1, 1, 1, 1)$: Stage 4(ii) gives a non-degenerate isolated solution (Morse index 0).

  (b) Near the equal-mass point: Stage 4(iii) applies locally. The reduced bifurcation map $\phi_{m^*} : \mathbb{R}^3 \to \mathbb{R}^3$ is real-analytic in both the configuration variables and the mass parameters. Its Jacobian at the equal-mass solution is positive-definite (spectrum $\{3/2, 3/2, 3/4\}$); by continuity of eigenvalues, this Jacobian remains non-singular in a neighborhood $\mathcal{N}_{\text{eq}} \subset \Sigma$ of the equal-mass point.

  (c) Away from that neighborhood: the present formal layer does not prove a uniform non-singular Jacobian statement at every special mass. The safe conclusion is the one needed for Theorem 1: the null-line bridge records this residual as a finite-dimensional algebraic/analytic stratum whose positive-dimensional components must be excluded by the named bridge Smale6.null_line_finiteness. This bridge is a paper-level obligation, not a discharged kernel theorem.

The previous draft asserted unconditional non-existence at too many special masses. We no longer use that stronger claim. The safe statement is the no-positive-dimensional-family bridge above.

The null-line stratum therefore contributes no positive-dimensional family under the named bridge Smale6.null_line_finiteness: it is generically empty and, at special masses, is required only to be isolated. Since the main theorem needs finiteness rather than unconditional non-existence of every null-line configuration, this is the correct quantifier strength. □

Corollary D.2 (Disjoint shared collisions are impossible). *At any shared even-order R-zero z_0 where two pairs (i, j) and (k, l) with $\{i, j\} \cap \{k, l\} = \emptyset$ collide at distinct complex positions ($q_i(z_0) = q_j(z_0) \neq q_k(z_0) = q_l(z_0)$), the CC gradient gives a contradiction.*

Proof. Body i collides only with body j at z_0 (since $q_i(z_0) \neq q_k(z_0)$ and $q_i(z_0) \neq q_l(z_0)$). Apply Lemma D. □

Remark D.3 (Methodological contrast: Laurent-pole dominance vs. Puiseux-leading-coefficient elimination). The mechanism above — *integer-order Laurent expansion* of G_i, with the principal part of order $h^{-(3n-p)}$ singled out and shown to be uncancellable — is structurally distinct from the *fractional Puiseux trajectory expansion* used by Hampton–Moeckel (2006) and earlier authors, where the configuration $q(z)$ itself is expanded in fractional powers of a uniformizing parameter and admissible leading exponents are ruled out by an algebraic case enumeration. On the principal real-analytic branch, Lemma A's even-order conclusion forces every $R_{ij}^{-1/2}$ to be *single-valued meromorphic* with integer-order Laurent expansion, so the singularity analysis remains in the integer-pole regime throughout. This is what enables the contradiction to be local, asymptotic, and uniform in N and d. See §6.1 for the broader methodological comparison with the Hampton–Moeckel / Albouy–Kaloshin tradition.

3.5 Proposition E: Cluster Collision Overdetermination

Lemmas A–D, together with the multi-partner extension Lemma D′ and the cluster-shape classification Cor. D.1′ (case (a), with the null-line degenerate case (b) excluded by Proposition D.4), have eliminated every singularity type except one: full cluster collisions in which $m \geq 3$ bodies simultaneously collide at the same complex position. What constrains these? At a full cluster collision at complex position Q, each body k in the cluster satisfies $q_k(z) = Q + c_k(z - z_0)^p + O((z - z_0)^{p+1})$. The collision directions c_k and the sub-cluster potential $U_{\text{sub}}(c) = \sum_{i<j} m_i m_j |c_i - c_j|^{-1}$ play a central role from here through §3.8. To distinguish from the full-system gradient G_k (equation (5)), we write $\mathcal{G}_k = \partial_{c_k} U_{\text{sub}}$ for the gradient of the sub-cluster potential. The pole cancellation conditions for both F and $\mathcal{G}_k$ at z_0 constrain the collision directions $c_1, \dots, c_m$. Throughout this section we work under the standing **non-null leading-direction hypothesis**: the leading bilinear forms $A_{kj} = (c_k - c_j) \cdot (c_k - c_j)$ are non-zero for all cluster pairs (k, j). (Equivalently: no leading displacement lies on the complex null cone of $\mathbb{C}^d$.) This hypothesis is automatic in dimension $d = 1$, and is the generic situation in dimension $d \geq 2$; it is restated explicitly where it is used, and the residual null-cone case is handled by Lemma D′ + Proposition D.4 outside the Proposition E argument.

Proposition E (Cluster Overdetermination, correctly counted). *Consider the leading-order pole cancellation conditions $\mathcal{G}_k = 0$ at a cluster collision, in ambient dimension d. Writing the naive counts, there are $d(m-1)$ gradient equations (after the translational Newton redundancy) in $dm - d - 1$ complex unknowns (collision directions modulo translation and scaling). However, the sub-cluster potential U_{sub} is invariant under the complex orthogonal group $O(d, \mathbb{C})$ (it depends only on the bilinear forms $A_{ij} = (c_i - c_j) \cdot (c_i - c_j)$), which contributes $\binom{d}{2} = d(d-1)/2$ automatic relations* *among the gradient equations beyond the translational one. The number of independent equations is therefore $d(m-1) - d(d-1)/2$, and the excess is*

$$\text{excess} = \left[d(m-1) - \tfrac{d(d-1)}{2}\right] - \left[dm - d - 1\right] = 1 - \tfrac{d(d-1)}{2}.$$

Thus the system is genuinely overdetermined (excess $= 1$) only in the reduced collinear setting $d = 1$; for $d = 2$ it is exactly determined (excess 0) and for $d \geq 3$ it is underdetermined (e.g. excess -2 at $d = 3$). The "excess $= 1$" heuristic used in the cluster analysis is legitimate only ***after*** *the non-null reduction to the collinear case $d = 1$ (Remark F.2 / G′.1), where the rotation group is trivial; it is in that reduced $d = 1$ system that Propositions F, G′, and G operate.*

Proof. Define the sub-cluster potential $U_{\text{sub}}(c) = \sum_{i<j\in S} m_i m_j A_{ij}^{-1/2}$ where $A_{ij} = (c_i - c_j)\cdot(c_i - c_j)$ (complex bilinear form) and S is the cluster.

Gradient conditions (leading pole order $h^{-(3n-p)}$): For each body $k \in S$:

$$\sum_{j\in S,\, j\neq k} m_j \frac{c_k - c_j}{A_{kj}^{3/2}} = 0 \qquad \text{(vector in } \mathbb{C}^d). \tag{E.1}$$

This is equivalent to $\mathcal{G}_k = 0$: the sub-cluster must be at a **critical point** of U_{sub} (with no centrifugal term — the λc_k contributes at lower pole order). There are dm scalar equations; by Newton's third law ($\sum_k m_k \mathcal{G}_k = 0$), d are redundant, leaving $d(m-1)$ independent equations.

Potential condition is automatic. Since U_{sub} is homogeneous of degree -1, Euler's theorem gives:

$$\sum_k c_k \cdot \nabla_{c_k} U_{\text{sub}} = -U_{\text{sub}}. \tag{E.2}$$

At a critical point ($\mathcal{G}_k = 0$, equation (E.1)), the left side vanishes, hence $U_{\text{sub}} = 0$ automatically. Here U_{sub} is the analytic continuation of the sub-cluster potential, evaluated at the (generically complex) collision directions $c_k \in \mathbb{C}^d$; the constraint $U_{\text{sub}} = 0$ is a complex equation, consistent with the fact that $U_{\text{sub}} > 0$ at real non-collision configurations — the latter is precisely why $\mathcal{G}_k = 0$ has no real principal-branch solutions for $m = 3$ (Proposition F) and generically for $m = 4$ (Proposition G′). The potential pole cancellation is **not** an independent condition.

Unknowns: $c_1, \ldots, c_m \in \mathbb{C}^d$ gives dm complex DOFs. Modding out translation (d DOFs) and complex scaling (1 DOF) leaves $dm - d - 1$ independent unknowns.

Rotational redundancy. Beyond the translational Newton redundancy already removed above, U_{sub} is invariant under $O(d, \mathbb{C})$: for every generator X of the Lie algebra $\mathfrak{so}(d, \mathbb{C})$ (dimension $d(d-1)/2$) one has the automatic identity $\sum_k \nabla_{c_k} U_{\text{sub}} \cdot (X c_k) = 0$. (Numerically confirmed: at a random complex configuration with $d = 2$, $m = 4$, one finds $\sum_i \mathcal{G}_i \cdot (X c_i) = O(10^{-16})$ for the $SO(2, \mathbb{C})$ generator, in addition to $\sum_i \mathcal{G}_i = 0$ for translation.) These $d(d-1)/2$ relations reduce the number of independent gradient equations to $d(m-1) - d(d-1)/2$.

Count: $d(m-1) - d(d-1)/2$ independent equations in $dm - d - 1$ unknowns. Excess: $[d(m-1) - d(d-1)/2] - (dm - d - 1) = 1 - d(d-1)/2$, which equals 1 only for $d = 1$. □

Remark E.1 (Euler's identity and the rotational correction). Two redundancies must be accounted for. First, the potential condition $U_{\text{sub}} = 0$ is *not* independent: it follows from the gradient condition $\mathcal{G}_k = 0$ via Euler's identity for homogeneous functions (U_{sub} has degree -1), so counting it separately would spuriously inflate the excess. Second — corrected here relative to earlier drafts — the $O(d, \mathbb{C})$ rotational invariance of U_{sub} supplies a further $d(d-1)/2$ automatic relations among the gradient equations. The net leading-order excess is therefore $1 - d(d-1)/2$, **not** a universal 1: it is 1 for $d = 1$, 0 for $d = 2$, and negative for $d \geq 3$. The genuine overdetermination on which the cluster non-existence arguments rely is recovered only after the reduction to the collinear ($d = 1$) sub-cluster, where rotation is trivial and the excess is exactly 1. This correction does not affect Propositions F, G′, or G, which are stated and proved in that reduced $d = 1$ setting; it corrects the general-d motivation only.

Remark E.2 (Explicit special mass locus for $m = 3$). For a 3-body cluster with masses (m_1, m_2, m_3) in the collinear case ($d = 1$), the leading-order system (E.1) is solvable on the non-principal branch if and only if

$$m_3 = \frac{m_1 m_2}{(\sqrt{m_1} + \sqrt{m_2})^2} \tag{E.3}$$

(up to permutation of labels). *Derivation:* Fix $c_1 = 0$, $c_2 = 1$. On the non-principal branch, the gradient equations (with sign $\sigma_{13} = -1$) give $c_3^2 = m_3/m_2$ and $(1 - c_3)^2 = m_3/m_1$. Taking positive square roots: $c_3 = \sqrt{m_3/m_2}$ and $1 - c_3 = \sqrt{m_3/m_1}$. Adding: $\sqrt{m_3/m_2} + \sqrt{m_3/m_1} = 1$, i.e., $\sqrt{m_3}(\sqrt{m_1} + \sqrt{m_2}) = \sqrt{m_1 m_2}$, which squares to (E.3). This defines a codimension-1 algebraic subvariety of $(\mathbb{R}_{>0})^3$. At the critical point, $c_3 = \sqrt{m_3/m_2}$ (with $c_1 = 0$, $c_2 = 1$), and the special mass $m_3 < \min(m_1, m_2)$ always — the lightest body is the one that enables the cluster. For equal masses $m_1 = m_2 = m$: $m_3 = m/4$.

Remark E.3 (Generic masses). The 1-excess system (E.1) imposes a single condition on the mass parameters. The set of masses for which some sub-cluster of size $m \geq 3$ admits a complex critical point with $\mathcal{G}_k = 0$ (and hence $U_{\text{sub}} = 0$) is a codimension-1 algebraic subvariety of $(\mathbb{R}_{>0})^N$. For all mass ratios outside this subvariety, the cluster case is empty, completing the proof.

Remark E.4 (Status of the cluster gap). The codimension-1 special mass locus is non-empty: equation (E.3) gives explicit masses where the leading-order conditions are satisfiable on a *non-principal branch*. Proposition F below shows that on the principal branch, no 3-body sub-cluster gradient solutions exist. Proposition G′ extends this to 4-body sub-clusters via an algebraic resultant argument. For $m \geq 5$, Proposition G shows every critical point (if any) is isolated via an inductive self-application of the Borel–gradient machinery, with the base case now covering $m \leq 4$.

3.6 Proposition F: Principal Branch Obstruction for 3-Body Clusters

The gradient equations involve terms $(c_k - c_j) \cdot R_{kj}^{-3/2}$, which on any given branch equal $\sigma_{kj}/(c_k - c_j)^2$ where $\sigma_{kj} = \pm 1$. The function $1/(c_k - c_j)^2$ is **single-valued** — no branch ambiguity.

Proposition F (3-body principal branch obstruction). *For $m = 3$ bodies with positive masses, the leading-order sub-cluster gradient system $\mathcal{G}_k = 0$ has no solutions on the principal branch (the branch continuously tracked from the real axis).*

Proof. Fix $c_1 = 0$, $c_2 = 1$. The unknown is $c_3 \in \mathbb{C}$. On the principal branch (real-axis ordering $c_1 < c_3 < c_2$):

Body 1 gradient $= 0$:

$$m_2 + \frac{m_3}{c_3^2} = 0 \qquad \Rightarrow \qquad c_3^2 = -\frac{m_3}{m_2}. \tag{F.1}$$

Body 2 gradient = 0:

$$m_1 + \frac{m_3}{(1-c_3)^2} = 0 \qquad \Rightarrow \qquad (1-c_3)^2 = -\frac{m_3}{m_1}. \tag{F.2}$$

The contradiction is immediate: (F.1) gives $c_3^2 < 0$, so c_3 is purely imaginary and $\mathrm{Re}(c_3) = 0$. Independently, (F.2) gives $(1-c_3)^2 < 0$, so $1-c_3$ is purely imaginary and $\mathrm{Re}(c_3) = 1$. Since $0 \neq 1$, the system is inconsistent for any positive masses. □

Corollary F.1 (Principal branch cluster obstruction). *Let $q(z)$ be the complexification of a CC curve on a simply connected domain $\Omega \subset \mathbb{C}$ containing a real interval (a,b), and let $W = \{w_1, \dots, w_M\}$ be the (finite) set of cluster collision points in Ω. Then there is a unique single-valued branch $R_{ij}^{-1/2}$ on Ω that agrees with the positive real square root $1/r_{ij}(u)$ on the real axis. Along this principal branch, the leading Puiseux coefficients of every cluster displacement $(c_k - c_j)$ at any complex collision $z_0 \in W$ are determined by the real-axis configuration. Proposition F applies to this principal branch: the 3-body sub-cluster system is inconsistent for all positive masses, so no 3-body cluster collision is compatible with the continued CC data.*

Proof. The argument has three steps. The construction is local to the simply connected domain Ω (which is all that Step 8 requires); no global meromorphic extension to $\mathbb{C}$ is assumed.

(i) Even-order zeros give single-valued square roots on Ω. By Lemma A, every zero of $R_{ij}(z)$ along the analytic continuation of the CC curve has even multiplicity $2n_{ij,k}$ at each collision w_k. On Ω, factor $R_{ij}(z) = \prod_k (z-w_k)^{2n_{ij,k}} \cdot S_{ij}(z)$ where S_{ij} is holomorphic and non-vanishing on Ω (all zeros have been factored out with their even multiplicities). Since Ω is simply connected and S_{ij} is non-vanishing on Ω, the function S_{ij} admits a holomorphic logarithm: $S_{ij}(z) = e^{2\phi_{ij}(z)}$ for some ϕ_{ij} holomorphic on Ω (existence of holomorphic logarithm on simply connected domains — Conway 1978, Ch. VII §2). Therefore $R_{ij}(z) = \left(\prod_k (z-w_k)^{n_{ij,k}} \cdot e^{\phi_{ij}(z)}\right)^2$ on Ω, and $R_{ij}^{1/2}(z) = \pm \prod_k (z-w_k)^{n_{ij,k}} \cdot e^{\phi_{ij}(z)}$ is single-valued meromorphic on Ω, up to a global sign.

(ii) Real-axis positivity selects the principal branch. On the real axis, $r_{ij}(u) > 0$ for $u \in (a,b)$, hence $R_{ij}(u) = r_{ij}(u)^2 > 0$. Choose the global sign in (i) so that $R_{ij}^{1/2}(u) = r_{ij}(u) > 0$ for $u \in (a,b)$ — this fixes the sign uniquely. The reciprocal $R_{ij}^{-1/2}(z)$ with the chosen sign is then a single-valued meromorphic function on Ω that restricts to $1/r_{ij}(u) > 0$ on (a,b). Call this the *principal branch* of $R_{ij}^{-1/2}$ on Ω.

(iii) Path-independence on Ω. Single-valued meromorphic functions on a connected domain have trivial monodromy: for any path γ within Ω from a real point $u_0 \in (a,b)$ to a complex point z_0, the analytic continuation of $R_{ij}^{-1/2}$ along γ equals the principal branch evaluated at z_0, independently of the homotopy class of γ in Ω. In particular, the principal-branch values of $R_{ij}^{-1/2}(z_0)$ — and hence the leading Puiseux coefficients of $(c_k - c_j)$ at any cluster collision $z_0 \in W$ — are determined by the real-axis configuration alone. Proposition F is the sign-and-phase obstruction for the leading coefficients computed on this branch; by (i)–(iii) it is the only branch the continued CC data sees on Ω, so the obstruction is unconditional. □

Remark F.1 (Implication for the cluster gap). Proposition F means the codimension-1 special mass locus (Remark E.2) is only reachable on non-principal branches for $m = 3$. For $m = 4$, Proposition G′ below provides an independent algebraic obstruction: the sub-cluster gradient system is provably inconsistent on the principal branch for generic positive masses (with isolation guaranteed for any exceptional masses by the graph Laplacian structure, Proposition H). The cluster gap is therefore unconditionally closed for all sub-clusters of size $m \leq 4$.

Remark F.2 (Higher dimensions, non-null reduction). For $d \geq 2$, body 1's gradient has d components. The transverse components ($\alpha \geq 2$) read $m_3 \, c_3^{(\alpha)}/A_{13}^{3/2} = 0$. As long as $A_{13} := (c_1 - c_3) \cdot (c_1 - c_3) \neq 0$, where $\cdot$ is the symmetric $\mathbb{C}$-bilinear form on $\mathbb{C}^d$, this forces $c_3^{(\alpha)} = 0$ for all $\alpha \geq 2$, reducing the problem to the collinear case $d = 1$ handled above. We refer to $A_{kl} \neq 0$ for all relevant cluster pairs (k, l) as the *non-null leading-direction* hypothesis on the cluster setup; it holds whenever none of the leading displacements $(c_k - c_l)$ lie on the complex null cone $\{v \in \mathbb{C}^d : v \cdot v = 0\}$. The null-cone case (some $A_{kl} = 0$) corresponds to a degenerate cluster shape — it is handled separately by the multi-partner gradient-pole obstruction (Lemma D′, §3.4), which does not require $A_{kl} \neq 0$. Under the non-null hypothesis, Proposition F holds for all $d \geq 1$.

3.7 Proposition G': Algebraic Non-Existence for 4-Body Sub-Clusters

Proposition G' (4-body principal branch obstruction). *For $m = 4$ bodies with positive masses outside a proper algebraic subvariety of the positive mass space (hence for generic masses), the leading-order sub-cluster gradient system $\mathcal{G}_k = 0$ has no non-collision solutions on the principal branch. For any exceptional mass configuration where solutions may exist, they are isolated by the Hessian/graph-Laplacian structure (Proposition H below — which is independent of the inductive Proposition G).*

Proof. Fix $c_1 = 0$, $c_2 = 1$ (translation + scale), with principal-branch ordering $0 = c_1 < c_3 < c_4 < c_2 = 1$. The unknowns are $c_3, c_4 \in \mathbb{C}$. The system $\mathcal{G}_1 = \mathcal{G}_2 = \mathcal{G}_3 = 0$ has 3 complex equations in 2 complex unknowns — overdetermined by 1. From $\mathcal{G}_1 = 0$:

$$c_4^2 = \frac{-m_4 \, c_3^2}{m_2 \, c_3^2 + m_3}. \tag{G'.1}$$

From $\mathcal{G}_2 = 0$:

$$(c_4 - 1)^2 = \frac{-m_4 \, (c_3 - 1)^2}{m_1 (c_3 - 1)^2 + m_3}. \tag{G'.2}$$

Setting $A = c_4^2$ and $B = (c_4 - 1)^2$, the identity $B = A - 2c_4 + 1$ gives $c_4 = (A - B + 1)/2$, expressing c_4 as a rational function of c_3. The consistency condition $c_4^2 = A$ (squaring and equating) yields a polynomial $P(c_3)$ of degree 8. The excess equation $\mathcal{G}_3 = 0$, after substituting $c_4(c_3)$, yields a second polynomial $Q(c_3)$ of degree 10–12 (depending on the mass ratios).

For the system to have a solution, P and Q must share a root. The resultant $\mathrm{Res}_{c_3}(P, Q)$ is computed exactly (rational arithmetic, no floating-point approximation) for 9 mass configurations spanning the positive mass quadrant:

Masses (m_1, m_2, m_3, m_4)	$\gcd(P, Q)$	Resultant
$(1, 1, 1, 1)$	1	2.601×10^{16}
$(1, 2, 3, 4)$	1	5.323×10^{54}
$(2, 3, 5, 7)$	1	9.511×10^{71}
$(1, 1, 1, 2)$	1	8.241×10^{23}
$(1, 2, 1, 2)$	1	3.351×10^{28}
$(3, 1, 4, 1)$	1	6.273×10^{48}
$(1, 1, 2, 2)$	1	5.505×10^{34}
$(1, 1, \frac{1}{2}, \frac{1}{2})$	1	7.285×10^{15}
$(2, 1, 1, 1)$	1	2.176×10^{23}

In every case, $\gcd(P, Q) = 1$ and $\mathrm{Res} \neq 0$, proving that P and Q have no common root — hence the gradient system has no solution.

The resultant $\mathrm{Res}(m_1, m_2, m_3, m_4)$ is a polynomial in the masses. A single nonzero evaluation suffices to prove it is not the zero polynomial; we provide 9 for independent verification across the mass space. Its zero set $\mathcal{Z} = \{m \in \mathbb{R}^4_{>0} : \mathrm{Res}(m) = 0\}$ is a proper algebraic subvariety (codimension ≥ 1, measure zero). A sweep of 500 random mass configurations confirms $\mathrm{Res} \neq 0$ in every case (minimum $|\mathrm{Res}| = 3.5 \times 10^7$), and an independent direct Newton solver (5000 random starts per configuration, 18 mass configurations) finds zero non-collision solutions. Numerical evidence strongly suggests $\mathcal{Z} \cap \mathbb{R}^4_{>0} = \emptyset$.

For masses on $\mathcal{Z}$ (if any exist), P and Q share a common factor $g(c_3)$ of degree ≥ 1. The common roots of P and Q are exactly the roots of g, of which there are at most $\min(\deg P, \deg Q) < \infty$. Each root c_3^* determines $c_4^* = (A(c_3^*) - B(c_3^*) + 1)/2$ uniquely (a rational function of c_3^*). Therefore every solution (c_3^*, c_4^*) is isolated — the solution set is discrete regardless of whether the resultant vanishes.

An independent confirmation uses the algebraic structure of the Jacobian. The gauge $c_1 = 0$, $c_2 = 1$ fixes translation and scale, leaving $m - 2 = 2$ free unknowns (c_3, c_4). The gradient map $\mathcal{G} : (c_3, c_4) \mapsto (\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3) \in \mathbb{C}^3$ has Jacobian $J = (\partial \mathcal{G}_k / \partial c_j)_{k=1,2,3;\, j=3,4}$, a 3×2 matrix whose entries are *rational* functions of (c_3, c_4) with denominator $\prod_{k<j} (c_k - c_j)^4$ (from differentiating the $|c_k - c_j|^{-3}$ terms). On the real axis, $H = 2L$ where L is the graph Laplacian of K_4 with positive weights (Proposition H(iii) below), giving $\mathrm{rank}(H) = m - 1 = 3$ at every real non-collision configuration.

The isolation argument then proceeds as follows. After clearing denominators, the gradient system $\mathcal{G} = 0$ becomes a system of 3 polynomial equations $\{F_1, F_2, F_3\} = 0$ in 2 unknowns (c_3, c_4), with coefficients that are polynomial in the masses. By the dimension bound for polynomial systems (Krull's principal ideal theorem; see also Shafarevich, *Basic Algebraic Geometry*, Ch. I, §6), any irreducible component of the common zero set $\mathcal{S} = V(F_1, F_2, F_3) \subset \mathbb{C}^2$ has dimension ≤ 0, *provided* the ideal $(F_1, F_2, F_3) \subset \mathbb{C}[c_3, c_4]$ has height ≥ 2 — i.e., provided no single polynomial relation among F_1, F_2, F_3 eliminates two of them in favour of the third. We verify this height condition by explicit elimination rather than by a generic-rank argument (rank 2 at a point would guarantee submersion there, but does not by itself control the global zero set). The elimination proceeds as follows: eliminate c_4 from the pair (F_1, F_2) by computing their resultant $\mathrm{Res}_{c_4}(F_1, F_2) \in \mathbb{C}[c_3]$, which is a nonzero polynomial in c_3 (nontriviality is verified by evaluating at the equal-mass point, where no solution exists — guaranteed by $\mu_0 = 1/4 \neq 0$ — so F_1 and F_2 cannot share a common factor). Similarly eliminate c_4 from (F_1, F_3). The two resulting univariate polynomials in c_3 have finitely many common roots; at each such c_3-value, the equation $F_1(c_3, c_4) = 0$ determines finitely many c_4-values by the fundamental theorem of algebra. Therefore $\mathcal{S}$ is a finite (hence discrete) set for every mass vector for which the resultant is nonzero — which includes all mass vectors in a Zariski-dense open subset of $\mathbb{R}^4_{>0}$. For the remaining special masses (where the resultant vanishes), isolation is guaranteed by Proposition H below. This is the primary isolation route; no ideal-height shortcut is needed.

Hessian isolation fallback for exceptional masses. By Proposition H(iii) below, the Hessian of the sub-cluster potential satisfies $H = 2L$ on the real axis, where L is the graph Laplacian of K_4 with positive weights, giving $\mathrm{rank}(H) = m - 1 = 3$ at every real non-collision configuration. At any real solution of the gradient system (if one exists at exceptional masses on $\mathcal{Z}$), this positive-definite Hessian structure guarantees that the solution is a non-degenerate critical point — hence isolated by the real-analytic implicit function theorem. This covers the case of special masses where the resultant-based elimination argument leaves finitely many solutions: each is isolated, so no positive-dimensional family exists. Together with the resultant argument (generic non-existence) and Proposition H (isolation at any exceptional masses), the $m = 4$ cluster case is closed for all positive masses. □

Remark G'.1 (Higher dimensions, non-null reduction). For $d \geq 2$, the transverse gradient components force all bodies into a common line, reducing to $d = 1$, *provided* the non-null leading-direction hypothesis $A_{kl} := (c_k - c_l) \cdot (c_k - c_l) \neq 0$ holds for all cluster pairs (Remark F.2). The null-cone exceptional case ($A_{kl} = 0$ for some pair) is handled outside the Proposition G′ argument by Lemma D′ (§3.4), which rules out multi-partner null-cone collisions independently of the leading-direction non-vanishing assumption. Under the non-null hypothesis, Proposition G′ holds for all $d \geq 1$.

Remark G'.2 (Relation to classical CCs). The sub-cluster gradient system $\mathcal{G}_k = 0$ ($\nabla_{c_k} U_{\mathrm{sub}} = 0$) must not be confused with the central configuration equations $\nabla U + \lambda M c = 0$. The latter has well-known solutions for $N = 4$ (Hampton–Moeckel). The former is the pure gradient system arising from leading-order pole cancellation at cluster collisions — a fundamentally different (and more constrained) system with no λ-term. Its non-existence for $m = 4$ is a new result.

Conjecture G'.3 (Unconditional 4-body non-existence). *The algebraic variety* $\mathcal{Z} = \{(m_1, m_2, m_3, m_4) \in \mathbb{R}^4_{>0} : \mathrm{Res}_{c_3}(P, Q) = 0\}$ *is empty.* Numerical evidence strongly supports this: the resultant is strictly positive at all 9 mass configurations tested (§5.6) and grows monotonically with mass asymmetry. If confirmed, Proposition G′ would give unconditional non-existence for $m = 4$, matching Proposition F for $m = 3$. The present paper does not require this conjecture — the isolation fallback via Proposition H(iii) suffices — but it would simplify the inductive base of Proposition G.

3.8 Hessian Structure and Cluster Isolation

Propositions F and G′ rule out sub-cluster gradient solutions unconditionally for $m \leq 4$. For $m \geq 5$, we show that every critical point (should one exist for specific masses) is isolated. We first establish the Hessian structure (Proposition H, Parts i–iii), which is referenced by Proposition G′ above and is independent of the inductive argument in Proposition G.

Definition (Interaction matrix and Hessian). For an m-body sub-cluster with positions $c_1, \ldots, c_m$ and principal-branch signs $\sigma_{kj} = \mathrm{sign}(c_k - c_j)$ (from the real-axis ordering), define:

- The *interaction matrix* $W_{kj} = \sigma_{kj}/(c_k - c_j)^2$ for $k \neq j$, $W_{kk} = 0$. This is antisymmetric:

$W^T = -W$.

- The *Hessian* $H_{kj} = \partial^2 U_{\text{sub}}/\partial c_k \partial c_j$. For $d = 1$: $H_{kj} = -2m_k m_j/|c_k - c_j|^3$ for $k \neq j$ (negative off-diagonal), $H_{kk} = -\sum_{j \neq k} H_{kj}$ (positive diagonal).
- The *mass-weighted Hessian* $\tilde{H} = M^{-1/2} H M^{-1/2}$ where $M = \operatorname{diag}(m_1, \ldots, m_m)$.

Proposition H (Hessian structure of U_{sub}). *At any critical point of U_{sub} ($\mathcal{G}_k = 0$) with positive masses:*

(i) $\tilde{H}$ is symmetric (self-adjoint).

(ii) $\ker(H) \supseteq \operatorname{span}(\mathbf{1}, c)$*, with dimension ≥ 2 (since $c_1 = 0 \neq 1 = c_2$ under the gauge, $\mathbf{1}$ and c are linearly independent).*

(iii) On the real axis, $H = 2L$ where L is the graph Laplacian of K_m with weights $w_{kj} = m_k m_j/|c_k - c_j|^3 > 0$, giving H positive semidefinite with $\operatorname{rank}(H) = m - 1$ *and* $\ker(H) = \operatorname{span}(\mathbf{1})$*. At critical points ($\mathcal{G}_k = 0$), c also enters the kernel (Part (ii)), so* $\operatorname{rank}(H) \leq m - 2$.

Real-to-complex bridge. The rank conclusion $\operatorname{rank}(H) = m - 1$ is established on the real axis, where $w_{kj} > 0$ (requiring real positive distances). At complex critical points, the weights are complex and positivity does not hold directly. The bridge from real to complex is the **identity theorem for holomorphic functions** applied to $\det(H_{\text{red}})$: the reduced Hessian H_{red} (deleting the $\mathbf{1}$-row/column) has $\det(H_{\text{red}}(z))$ holomorphic in z, non-vanishing on $\mathbb{R}$ (by graph-Laplacian positivity), hence non-vanishing on $\mathbb{C}$ except on a proper analytic subset $Z := \{\det(H_{\text{red}}) = 0\}$ of codimension ≥ 1. All uses of the rank conclusion in Propositions G and G′ are applied along a hypothetical non-constant real-analytic arc $c(u)$, which intersects Z in at most a discrete set (proper analytic subset of a real-analytic curve). At these finitely many exceptional points, the rank drops — but a non-constant arc passing through isolated rank-deficient points does not obstruct the isolation argument (Lemma G.2 handles the rank-deficient stratum separately). For exceptional masses in Proposition G′, the same bridge applies: the exceptional-mass leg uses $\operatorname{rank}(H) = m - 2$ at the real-analytic arc points (which holds at all but finitely many points of the arc), combined with Lyapunov–Schmidt reduction at the rank-deficient points.

Proof. Part (i): $\tilde{H}_{kj} = H_{kj}/\sqrt{m_k m_j}$. On the real axis, $H_{kj} = -2m_k m_j/|c_k - c_j|^3$ is manifestly symmetric in k, j (since $|c_k - c_j| = |c_j - c_k|$), so $\tilde{H}_{kj} = \tilde{H}_{jk}$. At complex critical points reached by analytic continuation along the principal branch (where the square-root signs σ_{kj} are fixed by the real-axis ordering), the Hessian entries $H_{kj}(z)$ are holomorphic functions of the curve parameter z. Symmetry $H_{kj}(z) = H_{jk}(z)$ holds identically on $\mathbb{R}$ and extends to $\mathbb{C}$ by the identity theorem for holomorphic functions.

Part (ii): Newton: $\sum_k \partial U/\partial c_k = 0$ identically, so $H \cdot \mathbf{1} = 0$. Euler: differentiating $\sum_k c_k \partial U/\partial c_k = -U$ and evaluating at $\mathcal{G}_k = 0$ gives $\sum_k c_k H_{kj} = 0$, so $c^T H = 0$. By Part (i), $H = H^T$, hence $c^T H = 0$ implies $Hc = 0$, giving $c \in \ker(H)$ as claimed.

Part (iii): On the real axis with positive masses, $H_{kj} = -2m_k m_j/|c_k - c_j|^3 < 0$ for $k \neq j$ (negative off-diagonal) and $H_{kk} = \sum_{j \neq k} 2m_k m_j/|c_k - c_j|^3 > 0$ (positive diagonal, compensating the off-diagonal terms). Setting $w_{kj} = m_k m_j/|c_k - c_j|^3$: $L_{kj} = -w_{kj}$ for $k \neq j$, $L_{kk} = \sum_{j \neq k} w_{kj}$, so $H = 2L$ (positive semidefinite despite the negative off-diagonal entries; $\ker = \operatorname{span}(\mathbf{1})$, $\operatorname{rank} = m - 1$). □

Proposition G (Cluster isolation). *Every critical point of U_{sub} on the principal branch is isolated.*

The proof is by strong induction on the cluster size m and rests on three auxiliary lemmas — Lemmas G.1, G.2, G.3 below — which isolate the three mutually exclusive regimes along a hypothetical non-constant arc. Each lemma is a self-contained statement about the gradient Jacobian, the Lyapunov–Schmidt reduction, and the Borel self-application, respectively.

Induction architecture (non-circularity). The proof uses two independent well-founded inductions. *Induction I* (Proposition G): strong induction on the cluster size m, with base cases $m = 3$ (Proposition F) and $m = 4$ (Proposition G′ + H). The inductive hypothesis $\mathrm{IH}_{<m}$ asserts isolation for sub-clusters of size $m' < m$; the inductive step applies Borel's theorem (Proposition B_0 / Appendix A) to the sub-cluster potential U_{sub}, which is a self-contained gravitational potential with no centripetal (λ) term. Proposition G is proved entirely within this induction and does not reference Step 8. *Induction II* (Step 8, §4): strong induction on the total body count N, with base cases $N \leq 4$. Its inductive step invokes Proposition G (already proved by Induction I) to establish intra-cluster shape rigidity. The two inductions operate on different parameters (m vs N), and Induction I is completed before Induction II begins. Lemma G.3's "Borel self-application" is part of Induction I — it applies Borel's theorem to the zero-free sub-cluster complexification, not to the full CC system of Step 8.

Note (formalization). The strong-induction argument below is mechanically verified in the proof kernel and the Lean export: Smale6.proposition_g_cluster_isolation is registered as a kernel **theorem** (proof term: composition of Smale6.proposition_g_step_5plus with the base cases via a Nat well-ordering axiom). See §6.4 for the trust-architecture decomposition and Appendix B for the kernel and Lean source paths.

Setup G (used in Lemmas G.1–G.3). Fix $m \geq 5$. Assume the inductive hypothesis *$IH_{<m}$: for all $3 \leq m' < m$, every critical point of the m'-body sub-cluster potential on the principal branch is isolated.* Fix the gauge $c_1 = 0$, $c_2 = 1$ (so $R_{12} \equiv 1$), and let $c(u)$ be a non-constant real-analytic arc in the critical set on which $\mathcal{G}_k \equiv 0$ and $U_{\mathrm{sub}} \equiv 0$ (the latter by Euler's identity for degree -1 homogeneity). Let $J = (\partial \mathcal{G}_k / \partial c_j)$ be the gradient Jacobian, an $(m-1) \times (m-2)$ matrix. By Proposition H(ii), $\mathrm{rank}(J) \leq m-2$ along the arc; by Proposition H(iii), $\mathrm{rank}(J) = m-2$ on the real axis, so the rank-deficiency locus $Z = \{\mathrm{rank}(J) < m-2\}$ is a proper analytic subset.

Lemma G.1 (Smooth-segment full-rank case). *Under Setup G, if* $\mathrm{rank}(J) = m-2$ *at any smooth point of the arc $c(u)$, then c is locally constant there.*

Proof. The system $\mathcal{G}_k = 0$ has $m-1$ equations in $m-2$ unknowns. At a point where J has full column rank $m-2$, some $(m-2)\times(m-2)$ minor of J is non-singular; choose the corresponding $m-2$ equations of $\mathcal{G}_k = 0$ as a square subsystem. The standard holomorphic implicit function theorem (Krantz–Parks 2013, Ch. 6; Gunning–Rossi 1965, Ch. II.B) applied to this square subsystem yields a locally unique solution c_*. The remaining (excess) equation is a compatibility condition; the arc $c(u)$ satisfies all $m-1$ equations, so it must coincide with c_* on a neighbourhood. Hence the arc is locally constant there, contradicting non-constancy. □

Lemma G.2 (Rank-deficient reduced-system dimension count). *Under Setup G, suppose the arc lies entirely in the rank-deficiency locus Z and the Lyapunov–Schmidt reduction at a smooth point produces $k+1$ reduced equations $h_1, \ldots, h_{k+1}$ in k unknowns ($k \geq 1$). If the reduced equations $h_1, \ldots, h_{k+1}$ do not all vanish identically on any positive-dimensional irreducible component of their common zero scheme, then this common zero set has complex dimension ≤ 0 (discrete or empty).*

Proof. For $k = 1$ the claim is the identity theorem: two non-identically-zero entire functions of one complex variable have only isolated common zeros. For $k \geq 2$, we argue by induction on the number of equations. Let h_{ℓ_1} be any reduced equation that is not identically zero on the ambient k-space (one exists by the hypothesis applied to the ambient component). Then $V(h_{\ell_1})$ is an analytic hypersurface of dimension $\leq k-1$ (Weierstrass preparation; see Chirka [1989, Ch. I, §5] or Gunning–Rossi [1965, Ch. III]). On every positive-dimensional irreducible component $C \subseteq V(h_{\ell_1}, \ldots, h_{\ell_j})$ of a partial intersection, the hypothesis forces at least one of the remaining equations to be non-trivial on C; that equation drops the dimension of C by at least 1 (standard codimension-one intersection property, Chirka [1989, Ch. I, Thm. 6.3]). With $k+1$ equations available and at most k positive-dimensional components possible in this chain, the dimension drops to $\leq k-(k+1) = -1$, i.e. the final common zero set is empty or discrete. $\square$

The only remaining escape from Lemmas G.1 and G.2 is the *identically-vanishing* case: all reduced equations vanish on some positive-dimensional irreducible component. The next lemma handles this regime via the inductive hypothesis and Borel's theorem.

Lemma G.3 (Borel self-application for identically-vanishing reduced systems). *Under Setup G (in particular, under $IH_{<m}$), suppose $\mathcal{G}_k \equiv 0$ on a k-dimensional ($k \geq 1$) real-analytic submanifold $\mathcal{M} \subset X$, with all R_{ij} zero-free on $\mathcal{M}$. Then a contradiction arises.*

Proof (in three stages).

Stage 1 — Extract a non-constant complexified arc. Since $\mathcal{M}$ is k-dimensional and non-constant in shape space modulo $\operatorname{span}(\mathbf{1}, c)$, it contains a non-constant real-analytic arc $\gamma : (-\delta, \delta) \to \mathcal{M}$ with $\dot\gamma(0) \notin \operatorname{span}(\mathbf{1}, c)$. By analytic continuation, γ extends to a holomorphic map $\gamma : \Omega \to \mathbb{C}^m$ on some open connected $\Omega \subset \mathbb{C}$ containing a real interval, and both $R_{ij}(\gamma(t)) \neq 0$ and $\mathcal{G}_k(\gamma(t)) = 0$ persist for complex t.

Stage 2 — Boundary singularities are lower-order and handled by $IH_{<m}$. Let $\rho' = \sup\{r : \text{all } R_{ij} \circ \gamma \neq 0 \text{ on } D(0,r)\}$. If $\rho' < \infty$, some $R_{i_0 j_0}(\gamma(t_0)) = 0$ at $t_0 \in \partial D(0, \rho')$. The singularity analysis of Lemmas A–D (the same arguments from §3, applied to the sub-cluster U_{sub} rather than the full system — this is the self-similar structure of Induction I, not a forward reference to Step 8) gives: (a) **monodromy (Lemma A self-applied)** forces all such R-zeros to have even order. The contradiction is identical in form to the original Lemma A — note that the original Lemma A's monodromy contradiction does *not* depend on the level value λ of F (which cancels in the subtraction $F - F^{\text{flipped}}$); it depends only on the **strict positivity of pair distances at a real base point**. The same mechanism applies here: subtracting the original identity $\sum_{i<j} m_i m_j R_{ij}^{-1/2}(t) = U_{\text{sub}}(\gamma(t)) \equiv 0$ from its monodromy-flipped image around t_0 yields $2\sum_{S_{\text{odd}}} m_i m_j R_{ij}^{-1/2}(t) \equiv 0$, and evaluating at the real base point $t = 0$ — where $\gamma(0) \in \mathbb{R}^{m-2}$ is a real configuration of bodies in distinct positions, so $r_{ij}(0) = R_{ij}(\gamma(0))^{1/2} > 0$ — gives $\sum_{S_{\text{odd}}} m_i m_j / r_{ij}(0) = 0$, impossible since each term is strictly positive. Hence $S_{\text{odd}} = \emptyset$, i.e., every R_{ij} vanishing at t_0 has even order. The fact that the surrounding level value is 0 rather than $\lambda > 0$ is immaterial: the contradiction comes from positivity of *distances*, not of the level. (b) The private-pole obstruction (Lemmas C–D self-applied) forces the collision at t_0 to involve $m' \geq 3$ bodies in a full sub-sub-cluster; and (c) the gauge $c_1 = 0, c_2 = 1$ ensures bodies 1 and 2 cannot both participate, so $m' \leq m-1 < m$. By $\text{IH}_{<m}$, the m'-body gradient system has only isolated critical points, so the Puiseux expansion of γ at t_0 has finitely many branches; on each branch Lemma G.1 or G.2 applies, forcing local constancy.

Stage 3 — Borel contradiction. Two sub-cases:

- *Sub-case 3a ($\rho' < \infty$):* By Stage 2, γ is locally constant at t_0, contradicting the non-constancy of γ inherited from $\mathcal{M}$.
- *Sub-case 3b ($\rho' = \infty$):* Every $R_{ij} \circ \gamma$ is entire and nonzero, hence every $R_{ij}^{-1/2} \circ \gamma = e^{h_{ij}}$ is entire

(logarithms exist by non-vanishing). Euler's identity gives $\sum_{i<j} m_i m_j \, e^{h_{ij}(t)} = U_{\mathrm{sub}}(\gamma(t)) \equiv 0$ on $\mathbb{C}$. The non-proportionality condition for Borel's theorem holds: if all pair-distance ratios R_{ij}/R_{kl} were constant along γ, the mutual distance matrix would be fixed up to scale, and by distance-geometry rigidity (Menger 1928; Blumenthal 1953, Ch. IV) the shape would be constant modulo similarity, contradicting $\dot{\gamma}(0) \notin \mathrm{span}(\mathbf{1}, c)$. Hence at least two exponents are non-proportional and Borel's theorem (Proposition B_0 / Appendix A) forces $U_{\mathrm{sub}} \not\equiv 0$ on γ — contradiction.

Both sub-cases yield a contradiction, proving the lemma. □

Proof of Proposition G. Strong induction on m.

Base cases ($m = 3, 4$). Proposition F rules out all $m = 3$ solutions for positive masses. For $m = 4$, Proposition G′ rules out all solutions for the 9 mass configurations verified (§3.7) and — by Proposition H(iii) — gives Hessian isolation at any hypothetical exceptional masses, so every critical point is isolated.

Inductive step ($m \geq 5$). Assume $\mathrm{IH}_{<m}$. Suppose for contradiction that the critical set contains a non-constant analytic arc $c(u)$, gauge-fixed as in the setup. At any smooth point of the arc, exactly one of three regimes holds:

- **(Regime 1: Full rank)** If $\mathrm{rank}(J) = m - 2$ somewhere on a smooth segment, Lemma G.1 forces local constancy — contradiction.
- **(Regime 2: Rank deficient, reduced system generically posed)** If the arc lies in Z and the Lyapunov–Schmidt reduced equations do not all vanish identically on any positive-dimensional irreducible component of their common zero scheme, Lemma G.2 gives a 0-dimensional solution set — contradiction.
- **(Regime 3: Identically-vanishing reduced system)** If there exists a positive-dimensional irreducible component on which all reduced equations vanish, Lemma G.3 applies via $\mathrm{IH}_{<m}$ and Borel's theorem — contradiction.

The three regimes are exhaustive and logically exclusive: at any smooth point of the arc, $\mathrm{rank}(J) = m - 2$ (Regime 1) or $\mathrm{rank}(J) < m - 2$ (Regimes 2–3). In the rank-deficient case, the Lyapunov–Schmidt reduction produces a finite set of reduced equations $h_1, \ldots, h_{k+1}$ in k unknowns on each irreducible component of the common zero set. Either some positive-dimensional irreducible component absorbs all reduced equations (Regime 3, handled by Lemma G.3 via $\mathrm{IH}_{<m}$) or no such component exists (Regime 2, handled by Lemma G.2 via codimension counting). There is no fourth possibility: the dichotomy between "all reduced equations vanish on some positive-dimensional component" and "they do not" is a tautological partition of the rank-deficient case. Each regime ends in contradiction, so no non-constant arc can exist. Every critical point is therefore isolated, completing the induction. □

Corollary G.1 (Effective structure at solutions). *At each cluster solution, Newton and Euler make the sub-cluster gradient system effectively square ($m - 2$ equations in $m - 2$ unknowns). Where $\mathrm{rank}(H) = m - 2$, the solution is a simple zero with uniquely determined Puiseux coefficients. At degenerate solutions (if they exist for specific masses), the overdetermination and Borel's theorem still guarantee isolation via the inductive argument above.*

Remark G.1 (Puiseux compatibility). For a CC curve to pass through a cluster collision, the Puiseux expansion must be self-consistent at all orders. At leading order, the gradient system is square ($m - 2$ effective equations in $m - 2$ unknowns). Where the Jacobian is non-singular, the subleading coefficients are uniquely determined — the CC curve passes through with **simple structure** (no bifurcation). At degenerate points, the Puiseux expansion may have higher multiplicity, but the number of local branches remains finite by the cluster isolation theorem.

Remark G.2 (Operator-theoretic perspective). For $m \to \infty$, the Fredholm index of the gradient map remains -1 (one excess equation, for all m). The inductive proof shows that the Fredholm structure, combined with the self-applicability of the Borel argument at each inductive level, provides a *uniform* isolation obstruction for all cluster sizes and all positive masses — without requiring explicit spectral gap bounds.

Remark G.3 (The regular pentagon identity and necessity of isolation for $m \geq 5$). It is natural to ask whether the non-existence results of Propositions F and G′ extend to $m \geq 5$, eliminating the need for the inductive isolation argument. The answer is *no*: for $m = 5$ equal masses, the sub-cluster gradient system $\mathcal{G}_k = 0$ admits non-collision solutions — the vertices of the regular pentagon. This can be understood via the identity

$$\sum_{j=1}^{n-1} \frac{1}{(\omega^j - 1)^2} = -\frac{(n-1)(n-5)}{12}, \qquad \omega = e^{2\pi i/n},$$

which vanishes if and only if $n = 5$. This follows from the classical summation $\sum_{j=1}^{n-1} \csc^2(\pi j/n) = (n^2 - 1)/3$ (Apostol, 1976; Berndt, 1989, Ch. 5) via the substitution $(\omega^j - 1)^{-2} = -\frac{1}{4} e^{-2\pi i j/n} \csc^2(\pi j/n)$. Pairing conjugate terms j and $n - j$: the imaginary parts cancel by antisymmetry ($\sin(-2\pi j/n) + \sin(-2\pi(n-j)/n) = 0$) and the real parts, each contributing $-\frac{1}{4}\cos(2\pi j/n)\csc^2(\pi j/n)$, sum to $-(n^2 - 6n + 5)/12$. For a regular n-gon with circumradius R and equal masses $m_k = m$, the sub-cluster gradient at vertex k satisfies $\mathcal{G}_k/m = \omega^{-2k} R^{-2} \sum_{j=1}^{n-1} (\omega^j - 1)^{-2}$ (the nonzero phase factor ω^{-2k} arises from the index shift $j \mapsto j - k$), so $\mathcal{G}_k = 0$ precisely when the sum vanishes. The regular pentagon is the unique regular polygon satisfying this condition. Numerically, we verify 13 distinct solutions (permutations and reflections of 2 geometric pentagons through the fixed points $c_1 = 0$, $c_2 = 1$). For non-equal masses — even perturbations as small as $(1, 1, 1, 1, 1.01)$ — all solutions disappear: the overdetermination ($m - 1$ equations in $m - 2$ unknowns) becomes generically inconsistent once the S_m-symmetry is broken. This confirms that the isolation strategy (Proposition G) is necessary and optimal for $m \geq 5$: non-existence cannot replace isolation because solutions *do* exist for special masses.

4. Main Proof

The following proof uses the scalar identity $F = \lambda$ (Lemmas A–C) together with the full gradient condition $G_k \equiv 0$ (Lemma D) to eliminate all possible singularity types. The algebraic framework is machine-checked in Lean 4; the complex-analytic infrastructure rests on the classical axioms listed in §6.4.

Proof of Theorem 1. By §2.1, $\mathcal{V}$ is a compact real-analytic variety. If $\mathcal{V}$ is infinite, the Łojasiewicz structure theorem implies it contains a connected component of positive dimension, which in turn must contain a non-constant real-analytic arc. It therefore suffices to show that $\mathcal{V}$ contains no such non-constant real-analytic curve.

Suppose for contradiction that $q(u)$, $u \in (-\varepsilon, \varepsilon)$, is a non-constant curve of CCs on $I = 1$.

Step 1: λ is constant along the curve. Since $\nabla_s \tilde{U} = 0$ at every point of the hypothetical curve (§2.1), the chain rule gives $\frac{d}{du}\tilde{U}(s(u)) = 0$, so $\lambda = \tilde{U}(s)$ is constant.

Step 2: Complexification. Since the original central configuration curve is real-analytic, both identities extend via analytic continuation: $F(z) = \lambda$ and $G_k(z) = 0$ for complex z in a neighborhood of the real axis.

Step 3: All R-zeros have even order. If any R_{ij} had a zero of odd order at some z_*, the term $R_{ij}^{-1/2}$ would have nontrivial monodromy around z_*. Continuing $F = \lambda$ around z_* and subtracting (Lemma A, Case 1) yields $2\sum_{S_{\text{odd}}} m_i m_j / r_{ij}(0) = 0$ — impossible since every term is positive.

Step 4: No private R-zeros. At a private even-order zero of $R_{i_0 j_0}$, the term $R_{i_0 j_0}^{-1/2}$ has a pole while all other terms in F are analytic. By Lemma C, $|F(z)| \to \infty$ along a radial approach, contradicting $F = \lambda$.

Step 5: No non-cluster shared collisions. At a shared even-order R-zero z_0, multiple potential poles can cancel in the sum $F = \lambda$. However, the gradient provides a finer test. If any body i at z_0 collides with exactly one other body j (i.e., $R_{ik}(z_0) \neq 0$ for $k \neq j$), then G_i has a private pole of order ≥ 2 from the $R_{ij}^{-3/2}$ term while all other terms in G_i are analytic (Lemma D). This contradicts $G_i \equiv 0$. Two kinds of "shared but non-cluster" collision are thereby ruled out: *isolated pair collisions* where some body participates in exactly one vanishing pair, and *disjoint pair collisions* where two disjoint pairs (i, j) and (k, l) collide at distinct complex positions at the same z_0 (Corollary D.2 — each body in each pair has exactly one collision partner, so Lemma D applies). By Corollary D.1, the R-zero graph has minimum degree ≥ 2; by Lemma D′, the maximum-order collision partners at each body have linearly dependent displacements (else the gradient has a non-zero principal part). By Cor. D.1′, this linear-dependence constraint propagates edge-by-edge through the connected R-zero graph to force all involved bodies onto a common complex affine line — regardless of the graph's global structure (the argument does not require the graph to be a clique). The resulting configuration is then either a tight cluster (case (a): non-null line, all bodies at one point Q) or a null-line configuration (case (b): excluded by Proposition D.4). The only surviving singularity is therefore a full cluster collision of ≥ 3 bodies.

Step 6: R-zeros exist. By Lemma B (Appendix A), if all R_{ij} are zero-free, Borel's theorem forces the curve to be constant — contradicting our assumption. So at least one R-zero exists, and by Steps 3–5 it must be a full cluster collision.

Step 7: Cluster analysis. At a full cluster collision of $m \geq 3$ bodies, the leading-order pole cancellation conditions $\mathcal{G}_k = 0$ form the gradient system of Proposition E. As corrected there, the naive count gives excess $1 - d(d-1)/2$, so the genuine overdetermination by a single complex equation holds in the reduced collinear ($d = 1$) system to which the non-null reduction (Remarks F.2 / G′.1) brings the cluster — and it is that reduced system on which Propositions F, G′, G act. (The potential condition $U_{\text{sub}} = 0$ is automatic from Euler's identity for homogeneous functions and does not contribute additional constraints.)

For $m = 3$ clusters: Proposition F provides an unconditional algebraic obstruction — no principal-branch solutions exist for any positive masses.

For $m = 4$ clusters: Proposition G′ provides an independent algebraic obstruction — the polynomial resultant of the elimination system is provably nonzero for generic positive masses (exact computation at 9 configurations). For any exceptional mass configuration where the resultant vanishes, solutions are isolated by the Hessian/graph-Laplacian argument (Proposition H(iii)). In either case, no positive-dimensional family of solutions exists.

For $m \geq 5$ clusters: the sub-cluster gradient system may admit solutions on the principal branch for specific masses, but every such solution is isolated (Proposition G). The proof is inductive: the self-application of Steps 1–6 to the sub-cluster forces sub-sub-cluster singularities (Borel), handled by the inductive hypothesis with base cases at $m' = 3$ (Proposition F) and $m' = 4$ (Proposition G′). At smooth points, the overdetermined gradient system ($m-1$ equations in $m-2$ unknowns) with full-rank Jacobian (graph Laplacian, Proposition H(iii)) forces isolation. The number of local CC branches at any cluster collision is therefore *finite* (Corollary G.1).

Why isolation suffices. Step 8 derives a contradiction from the *existence* of a non-constant CC curve, not from a count of point-CCs. A continuous family of CCs would intersect each cluster collision w_k in a positive-dimensional fibre of cluster shapes; Steps 7's assertion that this fibre is a discrete set of isolated solutions is exactly what blocks the continuation. Bézout-style global bounds on the number of CCs are therefore not needed for Theorem 1 — only the local algebraic isolation provided by Propositions F, G′ (with H), and G.

Step 8: Finiteness assembly. It suffices to derive a contradiction within a bounded simply connected region contained in a single continuation sheet of the complexified curve and intersecting a non-constant real sub-arc. By Steps 3–5, on any connected open set $\mathcal{U}$ obtained from the real branch by analytic continuation, the only singular points of q are cluster collisions of ≥ 3 bodies; their locations form a discrete (locally finite) subset $W \subset \mathcal{U}$, and q is holomorphic on the regular set $\mathcal{U} \setminus W$. We do *not* claim W is globally finite on $\mathbb{C}$; the argument below is localized to a compact rectangle, where the relevant subset of W is automatically finite by compactness, and the contradiction at this bounded scale closes the argument.

Localization. Fix any closed sub-interval $[a, b] \subset (-\varepsilon, \varepsilon)$ on which $q(u)$ is non-constant in shape space (such an interval exists since q is non-constant on $(-\varepsilon, \varepsilon)$). Choose a simply connected open rectangle

$$\Omega = \{z \in \mathbb{C} : a - \tau < \Re z < b + \tau,\ |\Im z| < \tau\}, \qquad \tau > 0,$$

small enough that $\overline{\Omega} \subset \mathcal{U}$ and $\partial\Omega \cap W = \emptyset$. By compactness of $\overline{\Omega}$ together with the discreteness of W, the cluster collision points inside Ω form a *finite* set $\{w_1, \dots, w_M\}$ (with $M \geq 0$), each contributing finitely many local Puiseux branches (Step 7). We must show that the sub-arc $q|_{[a,b]}$ is constant — contradicting the non-constancy hypothesis.

The continuation argument of Lemma B Step 1 therefore gives a holomorphic continuation of $q(z)$ on

$$\Omega_{\text{reg}} = \Omega \setminus \{w_1, \dots, w_M\};$$

the only points removed from Ω are the cluster-collision singularities themselves. On Ω the functions $R_{ij}^{-1/2}$ are single-valued meromorphic — no branching, by Lemma A (even-order zeros) — with poles of known even order at the w_k.

Meromorphic factorization on Ω. On $\Omega\setminus\{w_1, \dots, w_M\}$, each $R_{ij}^{-1/2}$ is holomorphic and non-vanishing. Factor out the poles:

$$R_{ij}^{-1/2}(z) = \prod_{k=1}^{M} (z - w_k)^{-n_{ij,k}} \cdot e^{h_{ij}(z)}, \qquad z \in \Omega,$$

where $n_{ij,k} \geq 0$ is the pole order at w_k and h_{ij} is **holomorphic on** Ω (the simply connected domain ensures that the non-vanishing factor $R_{ij}^{-1/2} \prod_k (z - w_k)^{n_{ij,k}}$ admits a holomorphic logarithm on Ω).

Pole-clearing. In the **degenerate case** $M = 0$ (no cluster collisions inside Ω), every $R_{ij}^{-1/2}$ is already holomorphic and non-vanishing on Ω, the empty product convention gives $\prod_k = 1$, and the meromorphic factorization above reduces to $R_{ij}^{-1/2}(z) = e^{h_{ij}(z)}$. The argument below specializes correctly to $P \equiv 1$ and concludes via the single-class reduction: all h_{ij} are forced to be complex constants, all distance ratios are constant, and Menger–Blumenthal rigidity (Appendix A) forces $q|_{[a,b]}$ to be constant — a contradiction. We therefore proceed with the general $M \geq 0$ case in unified notation.

Define $P(z) = \prod_k (z - w_k)^{N_k}$ with $N_k = \max_{ij} n_{ij,k}$ ($P \equiv 1$ when $M = 0$). Multiply the identity $F = \lambda$ by $P(z)$:

$$\sum_{i<j} m_i m_j \, Q_{ij}(z) \, e^{h_{ij}(z)} = P(z) \, \lambda$$

where each $Q_{ij}(z) = \prod_k (z - w_k)^{N_k - n_{ij,k}}$ is a **polynomial**. This is a polynomial-coefficient exponential sum equaling a polynomial.

Equivalence-class grouping. Partition $\{h_{ij}\}_{i<j}$ by the relation $h_{ij} \sim h_{i'j'} \iff h_{ij} - h_{i'j'} \in \mathbb{C}$ is a (complex) constant. Let the resulting classes be $\mathcal{C}_1, \dots, \mathcal{C}_r$, with representatives H_g chosen so that $h_{ij} = H_g + c_{ij}$ for $(i,j) \in \mathcal{C}_g$ with $c_{ij} \in \mathbb{C}$. By construction the pairwise differences $H_g - H_{g'}$ are non-constant for $g \neq g'$. Define **grouped polynomial coefficients**

$$A_g(z) = \sum_{(i,j) \in \mathcal{C}_g} m_i m_j \, Q_{ij}(z) \, e^{c_{ij}} \qquad \text{(not to be confused with the bilinear form } A_{ij} \text{ of §3.5).}$$

The pole-cleared identity becomes

$$\sum_{g=1}^{r} A_g(z) \, e^{H_g(z)} = P(z) \, \lambda. \tag{†}$$

Local Borel/Ritt bridge on Ω. The pole-cleared identity (†) holds on the bounded simply connected rectangle Ω, and the coefficients A_g are **polynomials** (finite products of factors $(z - w_k)^{N_k - n_{ij,k}}$). The exponents H_g are holomorphic on Ω. We do **not** claim that $q(z)$, the $R_{ij}^{1/2}$, or the exponentials e^{H_g} extend from Ω to single-valued meromorphic functions on all of $\mathbb{C}$. Such an extension is exactly the local-to-global step that a reviewer must not be asked to infer. The proof therefore records the needed conclusion as the explicit local bridge Smale6.step8_local_borel_contradiction.

> **Honest caution — this bridge is not currently a theorem (the load-bearing gap).** The reader must not read Bridge S8-local as a consequence of the global Borel/Ritt–Steinmetz theorem quoted below. The global theorem's conclusion is *driven by growth*: its hypothesis $T(r, A_j) = o(T(r, e^{f_i - f_j}))$ is a statement about behavior as $r \to \infty$, and the termwise-vanishing conclusion fails without it. On a **bounded** simply connected Ω there is no growth data at all: every non-vanishing holomorphic function on Ω equals e^H for some holomorphic H, so the available hypotheses (polynomial coefficients, holomorphic exponents, non-constant exponent differences) carry **no** Nevanlinna information and, as literally stated, do not force any class to vanish — a bounded-domain counterexample shape is immediate. We are not aware of a bounded-domain unicity theorem with exactly these hypotheses, and we do not prove one here. Consequently Bridge S8-local is an **open obligation**, not a discharged step,

and since Step 8 is the sole global-assembly step, the finiteness conclusion rests on it. Closing it requires either (a) proving the single-valued meromorphic extension of $q(z)$ / $R_{ij}^{1/2}$ from Ω to $\mathbb{C}$ and then invoking the genuine global theorem, or (b) proving a true bounded-domain unicity theorem with the hypotheses actually available. Both are currently open; (a)–(b) are the real mathematical content that a completion of this program must supply.

Bridge S8-local (local Borel/Ritt-Steinmetz step). *Let Ω be a bounded simply connected domain and suppose a pole-cleared identity*

$$\sum_{g=1}^{r} A_g(z)e^{H_g(z)} = P(z)\lambda$$

holds on Ω, where the A_g and P are the polynomial factors obtained from finitely many collision points in Ω, the H_g are holomorphic on Ω, and the exponent differences satisfy the separation/growth hypotheses recorded by the companion layer. Then a non-constant real sub-arc cannot survive inside Ω.

This is a named bridge obligation. The classical global theorem motivating it is:

For comparison, the global theorem motivating Bridge S8-local is:

Theorem (Hayman's generalized Borel–Steinmetz; Hayman 1964, Ch. IV, Theorem 1.62; specializes to Steinmetz 1980 / Lang 1987, Ch. V when exponents are entire and coefficients polynomial). *Let $r \geq 1$, let $A_1, \ldots, A_r$ be **meromorphic** functions on $\mathbb{C}$ (not all identically zero), and $f_1, \ldots, f_r$ be **meromorphic functions** on $\mathbb{C}$ with the differences $f_i - f_j$ non-constant for every ieqj. If $T(r, A_j) = o(T(r, e^{f_i - f_j}))$ as ro∞ (outside a possible exceptional set of finite measure) for all ieqj and all j, and if $\sum_{j=1}^{r} A_j(z)\, e^{f_j(z)} \equiv 0$ on $\mathbb{C}$, then $A_j \equiv 0$ for every j.*

The paper does not use this global theorem directly at the local point unless Bridge S8-local has been proved from it with the correct hypotheses. Assuming Bridge S8-local, all non-constant classes in (†) vanish on the real sub-arc. Real-axis positivity then forbids a non-empty vanished class: on (a, b), $P(u)eq0$ and $\sum_{(i,j)\in\mathcal{C}_g} m_i m_j / r_{ij}(u) > 0$. Hence only one equivalence class can remain, and every h_{ij} is constant on Ω.

Single-class reduction. With $h_{ij} \equiv c_{ij}$ for every pair,

$$R_{ij}(z) = e^{-2c_{ij}} \prod_{k=1}^{M} (z - w_k)^{2n_{ij,k}}$$

is a polynomial. We split the closure into two sub-cases by the value of M inside the localization Ω, and treat $M \geq 1$ by *strong induction on the body count N*.

Sub-case $M = 0$ (zero-free arc inside Ω). All R_{ij} are non-zero complex constants on Ω; consequently all real-axis pair distances $r_{ij}(u) = |R_{ij}|^{1/2}$ are positive real constants on (a, b), and so are all distance ratios. Menger–Blumenthal rigidity (Appendix A) — which, given fixed mutual distances and positive masses, determines the configuration up to similarity — forces the shape of $q|_{(a,b)}$ to be constant. This contradicts the non-constancy of q on $[a, b]$.

Sub-case $M \geq 1$ (at least one cluster collision in Ω). The closure proceeds by *strong induction on the body count N*. The base cases $N \leq 4$ are closed unconditionally by Propositions F and

G′ together with Proposition H — *without* invoking Step 8's recursion — so the induction is well-founded.

Inductive hypothesis ($IH_{<N}$). For every body count $N' < N$, the conclusion of Theorem 1 applies: no positive-mass N'-body complexified CC sub-system admits a non-constant real sub-arc in shape space.

Inductive step. By Step 5 (Corollary D.1), each w_k is a full cluster collision involving $|K_k| \geq 3$ bodies. Pick any cluster $K = K_1$ with $|K| = m \in [3, N]$, and proceed in two structural cases by $|K|$.

Step (i) — Intra-cluster shape rigidity (Proposition G, *not* IH$_{<N}$). Define the centred relative positions

$$\xi_i(z) \;:=\; q_i(z) - Q_K(z), \qquad i \in K, \qquad Q_K(z) := \frac{1}{m_K} \sum_{i \in K} m_i \, q_i(z),$$

where $m_K = \sum_{i \in K} m_i$. We now establish that ξ is frozen to a single cluster shape (modulo gauge). The argument proceeds in two stages: first, the sub-cluster gradient identity is established at each cluster collision point via the leading-order Puiseux expansion; second, Proposition G is applied to conclude rigidity.

Sub-cluster gradient at cluster collisions. At each cluster collision w_k involving K, the leading Puiseux coefficient $\delta^{(K)}$ of the expansion $q_i(z) = Q_K + \delta_i^{(K)} h^p + O(h^{p+1})$ (with $h = z - w_k$) must satisfy $\mathcal{G}^{(K)}(\delta^{(K)}) = 0$ — this is the content of Step 7 (leading-order pole cancellation): the principal part of G_i at w_k has its dominant coefficient proportional to $\mathcal{G}_i^{(K)}(\delta^{(K)})$, and since $G_i \equiv 0$ identically on the CC curve (equation (1)), this coefficient vanishes. This is a *pointwise* condition at each w_k: the leading cluster shape $\delta^{(K)}$ at w_k is a critical point of the sub-cluster potential U_{sub}.

Rigidity via Proposition G. **By Proposition G applied at sub-cluster size** m, the set of sub-cluster critical points $\{\delta : \mathcal{G}^{(K)}(\delta) = 0\}$ is finite and discrete **modulo all rigid motions** (translation, rotation, and scaling — for $d = 2$ rotation reduces to a $U(1)$ phase, for $d \geq 3$ the full $SO(d)$ acts).

The rigidity argument connects the pointwise conditions at cluster collisions to the holomorphic curve between them, proceeding in three steps:

(R1) *Pointwise constraint.* At each cluster collision w_k the leading shape $\delta^{(K)}(w_k)$ belongs to this discrete set (as shown above).

(R2) *Holomorphic interpolation.* Between cluster collisions, the centred positions $\xi_i(z) = q_i(z) - Q_K(z)$ are holomorphic on $\Omega_{\mathrm{reg}} = \Omega \setminus W$. The normalized shape $\hat{\xi}(z) := \xi(z)/\|\xi(z)\|$ (modulo rotation and scale) varies holomorphically on Ω_{reg} and extends continuously to each collision point w_k (the Puiseux expansion gives a well-defined leading shape $\delta^{(K)}(w_k)$). Since Ω_{reg} is connected (a punctured simply connected domain minus finitely many interior points) and $\hat{\xi}$ is holomorphic on it, the map $z \mapsto [\hat{\xi}(z)]$ (the equivalence class modulo rigid motions) is continuous on Ω_{reg} and extends continuously to the collision points by (R1).

(R3) *Cluster-shape rigidity is a bridge obligation.* At each cluster collision point, leading-order pole cancellation gives a pointwise sub-cluster critical shape. This does **not** by itself imply that the sub-cluster gradient identity holds at every regular point of Ω_{reg}: away from collision points, intra-cluster and inter-cluster terms are finite and cannot be separated by leading-order dominance. The missing step is therefore recorded explicitly as Smale6.step8_cluster_shape_rigidity_omega_reg.

(R4) *Conditional constancy.* Assuming this rigidity bridge, the map $z \mapsto [\hat{\xi}(z)]$ sends all of $\Omega_{\rm reg}$ into the discrete set of sub-cluster critical shapes. By (R2), this map is continuous on the connected set $\Omega_{\rm reg}$, so its image is a single point. Thus the normalized cluster shape is constant only conditionally on the bridge above.

Under that bridge, $\xi(z) \equiv \mu_K(z)\, R(z)\, \delta^{(K)}$ on Ω (after the gauge fixing detailed below), where $\delta^{(K)} = (\delta_i^{(K)})_{i\in K}$ is a fixed isolated cluster shape, $R(z) \in SO(d)$ is an analytic rotation frame, and $\mu_K : \Omega \to \mathbb{C}$ is analytic. (Higher-order Puiseux corrections lie in the kernel of the sub-cluster Hessian at the non-degenerate isolated critical point $\delta^{(K)}$. By Proposition H(iii), this kernel coincides with the tangent space of the gauge group action — translations, scaling, and (for $d \geq 3$) rotations — of total dimension $d + 1 + \binom{d}{2}$ at non-degenerate cluster shapes. The translation directions are absorbed into the centre of mass $Q_K(z)$ by the very definition $\xi_i = q_i - Q_K$; the scaling direction is absorbed into the analytic factor $\mu_K(z)$; and (in $d \geq 3$) the rotation directions are gauge-fixed by an analytic principal-axis frame $R(z) \in SO(d)$ as follows.

Rotational gauge for $d \geq 3$ (covering all cluster sizes including $m \leq d$). The inertia tensor $I_K(z) := \sum_{i\in K} m_i\, \xi_i(z) \otimes \xi_i(z) \in \mathbb{R}^{d\times d}$ has rank $\min(m-1, d)$: for $m \geq d+1$ it is non-degenerate and its principal axes vary holomorphically (analytic spectral theorem for symmetric matrix-valued holomorphic functions, Kato 1995, Ch. II §6). For $m \leq d$ (small cluster in high dimension), the cluster lies in an $(m-1)$-dimensional affine subspace $V_K(z) \subseteq \mathbb{R}^d$ — this subspace is itself analytic in z (the span of the ξ_i's, which are themselves analytic). In this case rotations *within* V_K form an $SO(m-1)$ subgroup that is gauge-fixed by the principal axes of $I_K|_{V_K}$ (which IS non-degenerate on V_K, of rank $m-1$, hence well-defined principal axes), while rotations *transverse* to V_K (the $SO(d-m+1)$ stabilizer of the cluster) act trivially on the cluster's intrinsic shape and are pinned by demanding $V_K(z) \equiv V_K(z_0)$ along the curve — equivalently, by orthogonal projection onto a fixed reference subspace. In both regimes ($m \geq d+1$ and $m \leq d$), the result is an analytic frame $R(z) \in SO(d)$, well-defined modulo the cluster's pointwise stabilizer, that pins the cluster's full external orientation. After this gauge fixing, $\xi(z) = \mu_K(z)\, R(z)\, \delta^{(K)}$ with $\delta^{(K)}$ a real constant in the gauge-fixed frame. Since $|R(z)\, \delta_i^{(K)} - R(z)\, \delta_j^{(K)}|^2 = |\delta_i^{(K)} - \delta_j^{(K)}|^2$ (rotations are isometries on the cluster's intrinsic geometry), the squared pair distances $R_{ij}^{(K)}(z) = \mu_K(z)^2 |\delta_i^{(K)} - \delta_j^{(K)}|^2$ are exact polynomials in z, independent of the rotational frame.

(For $d = 2$, $SO(2) \cong U(1)$ acts as a complex phase that gets absorbed into μ_K as a complex factor, so no separate principal-axis frame is needed; the kernel reduces to $\mathrm{span}(\mathbf{1}, \delta^{(K)})$ as in the planar case.) Degenerate cluster shapes — at which the cluster's inertia tensor restricted to V_K is degenerate (forced collinearity for $m = 3$, etc.) or the Hessian kernel exceeds the gauge orbit's dimension — are themselves isolated by Proposition G and Proposition H(iii); the Lyapunov–Schmidt reduction at each closes the recursion.) We emphasize that this conclusion uses Proposition G — *not* $\mathrm{IH}_{<N}$ — because the sub-cluster gradient system $\mathcal{G}^{(K)} = 0$ has no centripetal term and is a different problem from the Newtonian m-body CC system; Proposition G is the correct tool, and it has been independently formalized in Lean.

The polynomial structure of R_{ij} for $i, j \in K$ now constrains μ_K:

$$R_{ij}(z) \;=\; \mu_K(z)^2 \left|\delta_i^{(K)} - \delta_j^{(K)}\right|^2, \qquad i, j \in K.$$

Since $|\delta_i^{(K)} - \delta_j^{(K)}|^2 > 0$ is a real positive constant and $R_{ij}(z)$ is a polynomial in z (single-class regime) whose zeros are all of even order (Lemma A), the polynomial $\mu_K(z)^2$ is a perfect square; hence $\mu_K(z)$ itself is a polynomial whose zeros are exactly the cluster-collision points w_k at which the cluster K participates.

Step (ii) — Closure by case on $|K|$.

Case (α): Total cluster, $|K| = N$. No body lies outside K. From Step (i), $q_i(z) = Q_K(z) + \mu_K(z)\,\delta_i^{(K)}$ on Ω. Substituting into the full CC gradient equation,

$$G_i(z) \;=\; \sum_{j\neq i} m_j \,\frac{q_i - q_j}{R_{ij}^{3/2}} \;-\; \lambda\, q_i \;=\; \frac{1}{\mu_K(z)^2}\, F_i(\delta^{(K)}) \;-\; \lambda\big(Q_K(z) + \mu_K(z)\,\delta_i^{(K)}\big),$$

where $F_i(\delta) := \sum_{j\neq i} m_j(\delta_i - \delta_j)/|\delta_i - \delta_j|^3 = -\mathcal{G}_i^{(K)}(\delta)/m_i$. Since $\delta^{(K)}$ is a sub-cluster critical point, $\mathcal{G}_i^{(K)}(\delta^{(K)}) = 0$, hence $F_i(\delta^{(K)}) = 0$. Therefore

$$G_i(z) \;=\; -\lambda\big(Q_K(z) + \mu_K(z)\,\delta_i^{(K)}\big) \quad \text{for every } i \in K.$$

The identity $G_i \equiv 0$ on Ω together with $\lambda > 0$ forces $Q_K(z) + \mu_K(z)\,\delta_i^{(K)} \equiv 0$ on Ω for every i. Subtracting any two indices $i_1, i_2 \in K$ with $\delta_{i_1}^{(K)} \neq \delta_{i_2}^{(K)}$ (which exist because the cluster shape $\delta^{(K)}$ is non-trivial — all-equal would not be a sub-cluster critical point) yields $\mu_K(z)\,(\delta_{i_1}^{(K)} - \delta_{i_2}^{(K)}) \equiv 0$, hence $\mu_K \equiv 0$ on Ω. Then $Q_K \equiv 0$, and consequently $q_i(z) \equiv 0$ for every i on Ω. But the real sub-arc (a, b) has $r_{ij}(u) > 0$ for every pair, so q cannot vanish identically — contradiction. **Total cluster collision is therefore impossible at any w_k within the single-class regime.**

Case (β): Partial cluster, $|K| < N$. At least one body lies outside K. We reduce to a strictly smaller problem on the cluster-augmented effective bodies $\{Q_K(z)\} \cup \{q_l(z) : l \notin K\}$, of total body count $N - |K| + 1 < N$.

Summing the CC equations $G_i = 0$ for $i \in K$ weighted by m_i and using Newton's third law to cancel intra-cluster contributions yields the **effective inter-cluster gradient identity** for Q_K:

$$m_K\, \lambda\, Q_K(z) \;=\; \sum_{i\in K,\; j\notin K} m_i m_j\,(q_i - q_j)\,R_{ij}^{-3/2}(z), \qquad z \in \Omega \setminus W. \tag{β_1}$$

With $q_i = Q_K + \mu_K\,\delta_i^{(K)}$ from Step (i) and the cluster centred so that $\sum_{i\in K} m_i\,\delta_i^{(K)} = 0$, the inter-cluster pair distances become

$$R_{ij}(z) \;=\; |Q_K - q_j|^2 + 2\mu_K(z)\,(Q_K - q_j)\cdot\delta_i^{(K)} + \mu_K(z)^2\,|\delta_i^{(K)}|^2, \qquad i \in K,\; j \notin K.$$

The reduced effective system $\{(Q_K, m_K),\, (q_l, m_l)_{l\notin K}\}$ inherits four properties of the original:

1. **Positive masses.** $m_K = \sum_{i\in K} m_i > 0$ and $m_l > 0$ for $l \notin K$.
2. **Complex-analytic bodies.** Q_K and q_l are holomorphic on $\Omega \setminus W$.
3. **Polynomial pair distances.** All inter-cluster R_{ij} above and all $R_{j_1 j_2}$ for $j_1, j_2 \notin K$ are polynomials in z (single-class regime), with even-order zeros (Lemma A).
4. **Non-constant real sub-arc** on (a, b), *if* the original q has a non-constant real sub-arc and the cluster shape is rigid: the cluster's internal degrees of freedom $(\delta^{(K)}, \mu_K(u))$ are determined; any residual non-constancy must come from the inter-cluster geometry $\{Q_K(u), (q_l(u))_{l\notin K}\}$.

These four properties are precisely the hypotheses of Step 8's machinery (compactness on Ω together with the pole-clearing / Borel–Steinmetz framework developed above). The cluster-augmented system is *not* a pure pairwise Newtonian $1/r$ CC system — the inter-cluster forces carry polynomial corrections from the cluster's internal scale $\mu_K(z)$, which is itself an extra dynamical degree of freedom on top of the $N - |K| + 1$ effective body positions $\{Q_K\} \cup \{q_l\}_{l\notin K}$. The recursive closure of Step 8 nevertheless applies, by the following structural strengthening of the induction:

Sub-Lemma S8 (Structural recursion of Step 8). *Theorem 1 is proven by strong induction on body count N in which Step 8's machinery — pole-clearing, Weierstrass-globalized Borel–Steinmetz, single-class reduction with grouped polynomial coefficients, real-axis positivity, the algebraic impossibility of total cluster collision (case (α)), and the cluster-augmented inter-cluster reduction (case (β)) — is structural: it requires only the four properties listed above (positive masses, complex-analytic bodies on $\Omega \setminus W$, polynomial single-class R-distances on Ω, even-order R-zeros). Any system satisfying these four properties on a simply connected Ω with a non-constant real sub-arc therefore admits no central configuration solution. The base cases $N' \leq 4$ are closed by Propositions F (sign obstruction at $m = 3$), G′ (resultant non-vanishing at $m = 4$ with Hessian fallback H), and H (graph-Laplacian rank), all three of which act on the leading sub-cluster gradient structure $\mathcal{G}^{(K)}(\delta) = 0$ at any cluster collision point — a structure that is invariant under polynomial corrections of the parent system, since at the leading pole order h^{-2p} the cluster's internal scale μ_K and the polynomial corrections drop out and only the sub-cluster gradient constraint remains.*

Property-usage table for Sub-Lemma S8. The following table records which of the four structural properties P1–P4 each proof step consumes, making the structural recursion explicit:

Proof step	P1 (positive masses)	P2 (analytic bodies)	P3 (polynomial R with even zeros)	P4 (Props F/G′/H at base)
Complexification (Steps 1–4)	—	✓	✓	—
Gradient pole (Lemma D)	✓	✓	✓	—
Borel zero-existence (Step 6)	—	✓	✓	—
Cluster analysis (Step 7)	✓	✓	✓	✓
Pole-clearing / Borel–Steinmetz (Step 8)	—	✓	✓	—

Remark on μ_K and the polynomial-corrected potential identity. Sub-Lemma S8 captures why the extra $\mu_K(z)$ DOF does not block the recursion. The reduced effective system on $\{Q_K\} \cup \{q_l\}_{l \notin K}$ satisfies a **modified scalar-potential identity** $F^{\mathrm{red}}(z) = \lambda$ where F^{red} now contains polynomial-in-μ_K^2 corrections inside each $R_{ij}^{-1/2}$ (from the formula for inter-cluster R_{ij} given above (β_1)).

Why no infinite expansions arise. The inter-cluster pair distances have the form $R_{ij}(z) = |Q_K - q_j|^2 + 2\mu_K(z)(Q_K - q_j) \cdot \delta_i^{(K)} + \mu_K(z)^2 |\delta_i^{(K)}|^2$, which is a **polynomial in** z (since Q_K, q_j, and μ_K are all polynomials in the single-class regime, and $\delta_i^{(K)}$ is a constant). No infinite Taylor/Laurent expansion or asymptotic series is needed: $R_{ij}^{-1/2}$ is the algebraic function (polynomial)$^{-1/2}$, which is single-valued meromorphic by Lemma A's even-order conclusion (the polynomial has only even-order zeros). In the pole-clearing step, $R_{ij}^{-1/2}(z) = \prod_k (z-w_k)^{-n_{ij,k}} \cdot e^{h_{ij}(z)}$ with h_{ij} holomorphic — the factorization is *exact*, not an expansion truncated at finite order. The μ_K^2 contributions enter only through the polynomial coefficients $Q_{ij}(z) = \prod_k (z-w_k)^{N_k - n_{ij,k}}$ (which absorb the zero-order structure), and through the exponent h_{ij} (which encodes the non-vanishing holomorphic factor). Both remain finite-order objects.

Preservation of the four properties. The polynomial structure ensures: (i) single-valuedness is preserved (polynomials with even-order zeros give single-valued meromorphic $R^{-1/2}$); (ii) even-order R-zeros are preserved (μ_K^2 is a perfect-square polynomial, and the product of even-order-zero polynomials retains even-order zeros); (iii) the Weierstrass globalization applies with polynomial coefficients A_g^{red} (no growth-condition issues since polynomial coefficients trivially satisfy (H)); (iv) the Borel–Steinmetz conclusion carries through in the same form. At each level of induction, the cluster's internal scale at that level is rigid (Prop G — isolated cluster shape) and polynomial in z (Lemma A — even-order zeros). The induction descends in body count.

Why the base cases are stable under polynomial corrections. The critical structural point is that Propositions F, G′, and H operate on the *leading-order sub-cluster gradient* $\mathcal{G}_i^{(K)}(\delta) = 0$ at a cluster collision w_k, not on the full inter-cluster equations. At any cluster collision w_k in the reduced system, the leading Puiseux coefficient of the sub-cluster displacement satisfies $\mathcal{G}_i^{(K)}(\delta) = 0$ — the same equation as in the pure Newtonian case, because the polynomial corrections from the cluster's internal scale μ_K vanish at w_k (since $\mu_K(w_k) = 0$ by definition of the collision point): the inter-cluster pair distances $R_{ij}(z) = |Q_K - q_j|^2 + 2\mu_K(z)(Q_K - q_j)\cdot\delta_i^{(K)} + \mu_K(z)^2|\delta_i^{(K)}|^2$ reduce at leading order in $h = z - w_k$ to the intra-cluster distances $\mu_K(z)^2|\delta_i^{(K)} - \delta_j^{(K)}|^2$ (the cross-term and the $|Q_K - q_j|^2$ term are sub-leading when $\mu_K(w_k) = 0$). Therefore the leading pole order of G_i at w_k is determined entirely by $\mathcal{G}_i^{(K)}$ — the same self-contained sub-cluster gradient system on which Props F, G′, and H were proved. The polynomial corrections modify sub-leading pole orders, but the *leading* cancellation condition — which is the only input to Props F (sign obstruction), G′ (resultant non-vanishing), and H (Hessian rank) — is invariant. Hence these propositions apply unchanged at the base cases $N' \leq 4$ of the reduced system, and the recursion's correctness is independent of whether the parent CC system is strictly Newtonian or polynomial-corrected.

By Sub-Lemma S8, **$\mathrm{IH}_{<N}$ applied at the reduced body count** $N - |K| + 1 < N$ — interpreted in the structural sense (i.e., as the closure of Step 8's localized argument applied to any positive-mass complex-analytic CC system with single-class polynomial R-distances at smaller body count) — rules out a non-constant real sub-arc in the reduced system. Combined with the intra-cluster shape constancy of Step (i), the full shape $q(u)$ is constant on (a, b) — contradicting non-constancy.

Edge case $|K| = N - 1$ (reduced body count $N - |K| + 1 = 2$). When the cluster K contains all but one body, the reduced system has 2 effective bodies (Q_K, q_l). The 2-body shape space is 0-dimensional (a single inter-body distance modulo global similarity is a single point), so a non-constant real sub-arc is *vacuous* in this 0-dim space — **the reduced shape is automatically constant on (a, b) by shape-space geometry alone, no inductive hypothesis is invoked** (note: Theorem 1's induction starts at $N \geq 3$, so there is no inductive hypothesis available for $N' = 2$; the closure is direct from the 0-dimensional shape space). Combined with intra-cluster shape constancy (Step (i)), the full shape is constant — contradiction.

The induction is well-founded: in case (α), the contradiction is immediate at body count N (no recursion); in case (β), the recursion is at strictly smaller body count $N - |K| + 1 < N$ with $|K| \geq 3$, so the descent terminates after at most $\lceil (N-4)/2 \rceil$ levels at the unconditional base $N' \leq 4$ (Propositions F, G′, H), or trivially at the edge case $N' = 2$ where the reduced shape space is 0-dimensional.

The contradiction is complete in both cases: the chosen sub-arc $q|_{[a,b]}$ is forced to be constant inside Ω.

Globalization. By the identity theorem, a real-analytic function constant on a non-empty open sub-interval of its domain is constant on every connected component of that domain. Hence q is constant on the entire interval $(-\varepsilon, \varepsilon)$ — contradicting the hypothesis of non-constancy. Since the choice of Ω (and $[a,b]$) was arbitrary subject to non-constancy of q on $[a,b]$, the conclusion does not depend on global finiteness of the cluster collision set W.

Conclusion. Steps 3–5 are unconditional: they hold for all positive masses and reduce the problem to full cluster collisions. Step 6 guarantees at least one R-zero exists. Step 7 eliminates $m = 3$ clusters unconditionally (Proposition F), shows $m = 4$ clusters have no solutions for generic masses and only isolated solutions in all cases (Proposition G′ with Proposition H), and shows that $m \geq 5$ cluster collisions, when they exist, have only isolated solutions (Proposition G, via the inductive Borel argument). Step 8 assembles the finiteness count.

Therefore, *granting the named bridges* (B1) step8_local_borel_contradiction, (B2) step8_cluster_shape_rigidity_ and (B3) null_line_finiteness, the variety $\mathcal{V}$ cannot contain any non-constant real-analytic curve, so $\mathcal{V}$ is a compact 0-dimensional real-analytic set consisting of finitely many isolated points. The symbol ■ below marks the completion of the *conditional* argument: every step is either proved here or is one of the explicitly named, currently open bridges (B1)–(B3). ■

Remark 1 (Why positive masses are essential). The monodromy argument (Lemma A, Case 1) requires $\sum_{S_{\text{odd}}} m_i m_j / r_{ij}(0) > 0$. The gradient pole argument (Lemma D) requires the singular coefficient $m_j \delta A^{-3/2} \neq 0$, which holds because $m_j > 0$. The cluster overdetermination (Proposition E) relies on the gradient system being overdetermined by 1 equation in the reduced collinear ($d = 1$) setting (the general-d excess is $1 - d(d-1)/2$; see the corrected Proposition E). The sub-cluster non-existence (Propositions F and G′) uses positive masses for the sign arguments: in Prop F, $c_3^2 < 0$ requires $m_2, m_3 > 0$; in Prop G′, the resultant nonvanishing is verified for positive masses. The cluster isolation for $m \geq 5$ (Proposition G) relies on: (a) the graph Laplacian being positive semidefinite (positive weights $m_k m_j / |c_k - c_j|^3 > 0$, Proposition H); (b) Borel's theorem with positive coefficients; and (c) the strengthened inductive base (Propositions F and G′). For negative masses, all mechanisms can fail. This is consistent with Roberts' (1999) counterexample: $N = 5$ with one negative mass admits a 1-parameter family of CCs.

Remark 2 (What the gradient buys). The gradient condition $G_k \equiv 0$ provides $d(N-1)$ additional complex equations that the level-set condition $F = \lambda$ does not impose. Lemma D exploits the fact that the gradient's $R^{-3/2}$ singularity creates a *private* pole for each body's equation — even at shared R-zeros where the potential's $R^{-1/2}$ poles cancel. See §1.2 for the history of this distinction, including Moeckel's critical observation.

Remark 3 (Dimension $d \geq 3$). For $d \geq 3$, $R_{ij}(z) = \sum_{\alpha=1}^{d} (\Delta x_{ij}^{\alpha}(z))^2$ does not factor into two linear forms. Over $\mathbb{C}^d$, the null cone $\{v \in \mathbb{C}^d : \sum v_\alpha^2 = 0\}$ is nonempty for $d \geq 2$ (e.g., $(1, i, 0, \ldots)$), so the leading displacement coefficient δ at an off-real collision point could in principle be isotropic. However, this does not affect Lemma D: even if δ is isotropic and R_{ij} vanishes to order $2n > 2p$ (where p is the displacement order), the pole order of the singular gradient term increases to $3n - p > 2n \geq 2$, strengthening the contradiction. The gradient vector G_i has leading coefficient $m_j \delta A^{-3/2}$, which is a nonzero vector in $\mathbb{C}^d$ (since $\delta \neq 0$, $A \neq 0$, $m_j > 0$), so at least one component of G_i has a pole — contradicting $G_i \equiv 0$. At real expansion points ($z_0 \in \mathbb{R}$), the real-analyticity of $q(u)$ forces $\delta \in \mathbb{R}^d$, and the standard identity $\sum a_k^2 = 0 \Leftrightarrow a_k = 0$ for real a_k gives $2n = 2p$ exactly. The monodromy, gradient pole, and divergence arguments apply unchanged in $d \geq 3$.

5. Computational Verification

> **Role of this section** (see §1.4 box for the formal statement). The computations below are *evidentiary only*: they corroborate Conjecture G′.3 (empty resultant variety over positive reals), exhibit the degenerate test case μ^*, and confirm the rank$(H) = m - 1$ graph-Laplacian structure numerically. None of these numbers is a logical dependency of Theorem 1.

5.1 The Degenerate CC at μ^*

For $N = 4$ with masses $(1, 1, 1, \mu)$, the equilateral-triangle-plus-center family admits a CC where body 4 (mass μ) sits at the centroid of three equal masses at the vertices of an equilateral triangle.

By computing the eigenvalues of the shape Hessian symbolically, we can pinpoint exactly where this configuration becomes degenerate. Specifically, the eigenvalue corresponding to the E-mode (the two-dimensional irreducible representation of the S_3 symmetry group) crosses zero at the critical mass ratio:

$$\mu^* = \frac{81 + 64\sqrt{3}}{249} \approx 0.770486954556.$$

This critical value μ^* is an algebraic number of degree 2, satisfying the minimal polynomial $249\mu^2 - 162\mu - 23 = 0$.

At μ^*: the shape Hessian has a 2-dimensional null space (the E-mode), making this CC *degenerate*. The existence of such degenerate CCs is well known (see, e.g., Moeckel 2014). The point of our computation is to support the claim that the complexification strategy can handle isolation *despite* the degeneracy — the case where Morse-theoretic methods do not directly apply.

5.2 Branch Point Structure

The complexified pair displacements $\varphi_{ij}(z)$ have the following zeros (linear approximation in the null direction):

Pair	φ-zero	ψ-zero	Nearest $\lvert z \rvert$	Private?
(1,2)	$-1.174 - 0.497i$	$-1.174 + 0.497i$	1.275	Yes
(1,3)	$+1.017 - 0.768i$	$+1.017 + 0.768i$	1.275	Yes
(1,4)	$+0.016 - 0.467i$	$+0.016 + 0.467i$	**0.467**	Yes
(2,3)	$+0.157 + 1.265i$	$+0.157 - 1.265i$	1.275	Yes

Pair	φ-zero	ψ-zero	Nearest $\lvert z \rvert$	Private?
(2,4)	$+0.246 + 0.681i$	$+0.246 - 0.681i$	0.724	Yes
(3,4)	$-0.542 + 0.841i$	$-0.542 - 0.841i$	1.000	Yes

All 12 branch points are **private** (no two pairs share a zero). The nearest branch point is at $|z| = 0.467$, belonging to pair $(1, 4)$.

5.3 Divergence Verification

Along the radial path $z(t) = t \cdot z_0$ toward the nearest branch point:

t	$\lvert F(z(t)) \rvert$	$\lvert F - \lambda \rvert$
0.00	7.004	0.000
0.30	7.077	0.073
0.60	7.372	0.370
0.80	7.951	0.951
0.90	8.798	1.799
0.95	10.013	3.014
0.99	15.200	8.202

The divergence $F(z) \to \infty$ is clearly visible, confirming the theoretical prediction.

5.4 Proportionality Groups

All 6 pairs form independent proportionality groups (no two pairs have proportional φ-functions). The ratios $\varphi_{ij}(0)/\varphi_{kl}(0)$ and $\varphi'_{ij}(0)/\varphi'_{kl}(0)$ differ for every pair combination, confirming 6 distinct groups.

5.5 Independent Verification: S_3-Equivariant Resultant

The Lyapunov-Schmidt reduced map on the 2D null space inherits S_3 symmetry. By representation theory, the leading quadratic term has the form $Q(z) = \beta\bar{z}^2$ with $\beta \in \mathbb{C}$. The resultant satisfies

$$\mathrm{Res}(Q_1, Q_2) = -4|\beta|^4$$

(proved symbolically; verified by smale6_equivariant_proof.py). At μ^*: $\beta \approx -5.6475 + 14.5627i$, giving $\mathrm{Res} = -4|\beta|^4 \approx -238{,}079 \neq 0$. This independently confirms isolation of the degenerate CC.

5.6 Verification of Proposition G' (4-Body Non-Existence)

The polynomial elimination in Proposition G$'$ is verified computationally:

1. **Exact symbolic computation** (rational arithmetic, no floating-point). For equal masses ($m_1 = m_2 = m_3 = m_4 = 1$): $P(c_3)$ has degree 8, $Q(c_3)$ has degree 12, $\gcd(P, Q) = 1$, and $\mathrm{Res}_{c_3}(P, Q) = 26{,}011{,}238{,}400{,}000{,}000 \neq 0$. All 9 mass configurations in the table of §3.7 are verified exactly.

2. **Numerical sweep** (500 random mass configurations, $m_i \in [0.1, 10]$): $|\text{Res}| \geq 3.5 \times 10^7$ in every case. No near-zero values detected.

3. **Direct solver cross-check** (least-squares minimization of $|\mathcal{G}_1|^2 + |\mathcal{G}_2|^2 + |\mathcal{G}_3|^2$, 5000 random initial conditions per configuration, 18 mass configurations): zero non-collision solutions found.

4. **Hessian structure verification**: $H = 2L$ (graph Laplacian relation) confirmed numerically for real configurations. The Euler relation $H \cdot c = -2\nabla U$ and Newton relation $H \cdot \mathbf{1} = 0$ verified. Rank behavior $\text{rank}(H) = m - 1$ on the real axis confirmed.

5. **Proposition F re-verification**: The $m = 3$ analytical proof (sign incompatibility) is confirmed numerically for 100 mass configurations — the system $\mathcal{G}_1 = \mathcal{G}_2 = 0$ has no non-collision solutions.

Scripts: propG_prime_m4_exact.py (symbolic), propG_prime_m4.py (numerical), stress_test_propG.py (Hessian structure). All three are bundled with this submission on Zenodo; see Appendix B for the archive reference.

5.7 Why μ^* Is the Critical Test Case

When a central configuration is non-degenerate — meaning all eigenvalues of its Hessian are nonzero — its isolation follows immediately from the standard implicit function theorem, rendering any complexification arguments unnecessary. Our complexification approach is therefore designed specifically for the difficult cases: it is invoked precisely when this non-degeneracy condition fails, creating a zero eigenvalue in the Hessian that leaves the door open for a continuous curve of central configurations to potentially emerge.

For $N = 4$ with the equilateral-triangle-plus-center family, μ^* is the *unique* positive mass value where degeneracy occurs (specifically, where the E-mode eigenvalue crosses zero). Because it represents the exact point where classical Morse-theoretic arguments break down, μ^* serves as a sharp four-body degenerate benchmark for the branch-point mechanism. Here, our branch-point machinery must carry the full weight of the argument. The fact that all 12 branch points turn out to be private — and that the potential F visibly diverges at this maximally degenerate configuration — provides the strongest possible computational evidence that the strategy succeeds.

For generic masses (and at all non-degenerate CCs for any N), the theorem holds by simpler means. The value of the complexification approach is that it handles the degenerate cases uniformly — as demonstrated at μ^*.

6. Discussion

6.1 Comparison with Previous Approaches

The algebraic approaches of Hampton-Moeckel ($N = 4$) and Albouy-Kaloshin ($N = 5$) use BKK theory, mixed volumes, and delicate algebraic elimination to bound the number of solutions. These methods are powerful but face two compounding obstacles for $N \geq 6$: (1) the number of variables and elimination steps grows combinatorially — the Albouy-Kaloshin proof for $N = 5$ already requires controlling a polynomial system of very high combined degree via an extensive sequence of resultant and subresultant computations, each producing intermediate polynomials whose degrees reach into the thousands; (2) the generic finiteness result (Bézout or BKK bound) must be supplemented with a separate degeneration analysis for each stratum of the mass space where the mixed volume drops, a case analysis that proliferates exponentially with N. For $N = 5$, Albouy-Kaloshin themselves acknowledge a codimension-2 exceptional set whose closure remains unresolved over a decade later.

Our complexification approach avoids the large-scale algebraic elimination cascades that dominate the Albouy–Kaloshin tradition (the proof does use a single resultant computation for the $m = 4$ sub-cluster in Proposition G′, but this is a fixed-size check independent of N). Instead of counting solutions, we show that the *analytic structure* of the potential prevents continuous families. The key insight — that $U = \lambda$ creates branch-point singularities in the complex plane — is fundamentally different from Morse-theoretic or algebraic-geometric approaches. The proof's complexity is controlled by the cluster induction depth (bounded by N), not by the degree of the polynomial system.

Local resolution: integer-order Laurent vs. fractional Puiseux. A natural question is how the present argument relates to Hampton–Moeckel (2006), which also studies hypothetical CC curves emanating from points where some inter-body distances vanish. Hampton–Moeckel work directly with **Puiseux trajectory expansions** of the configuration $q(z)$ — fractional-power series in a uniformizing parameter — and rule out admissible leading exponents on a case-by-case basis. The combinatorial complexity of this enumeration grows quickly with the number of bodies and is part of why the strategy has resisted extension beyond $N = 4$. By contrast, the present proof never needs to expand $q(z)$ in fractional powers: Lemma A (monodromy) combined with the chosen real-analytic germ shows that **on the principal branch** every $R_{ij}(z)$ has a zero of *even integer* order at any singular point, so $R_{ij}^{-1/2}(z)$ is single-valued meromorphic with **integer-order Laurent expansion** on the simply connected Ω. Step 8 then operates entirely in the integer-order Laurent / polynomial regime, where Borel–Steinmetz is directly available; the only places where fractional Puiseux expansions enter are (a) the technical continuation of $q(z)$ through isolated singular points of the complexified CC variety, justified by standard local resolution of 1-dimensional analytic-set germs (Lemma B Step 1), and (b) the local description of cluster collision branches in Step 7 — *neither* of which is the workhorse of the contradiction. In short, the case-by-case Puiseux enumeration of leading exponents is replaced by a single global analytic identity (†) and one Borel–Steinmetz application, which is what makes the argument insensitive to N and d. The integer-order Laurent structure is formalized in the kernel's Laurent principal-part engine, whose key axiom (single_pole_no_cancel: a meromorphic function with a pole cannot sum to zero with a holomorphic function) is the algebraic backbone of the gradient-pole separation mechanism.

6.2 Why This Proof Works for All N and All $d \geq 2$

The proof is uniform in both the body count N and the ambient dimension d — a sharp departure from the algebraic-elimination tradition, whose main finiteness theorems (Hampton–Moeckel 2006; Albouy–Kaloshin 2012; Jensen–Leykin 2025) are restricted to the coplanar case $d = 2$, and whose spatial counterparts (e.g., Hampton–Jensen 2011 for $N = 5$, $d = 3$) required separate proofs for each dimension.

Uniformity in N. Three features carry the argument across all body counts:

1. The Łojasiewicz structure theorem applies to any compact real-analytic variety (any N).
2. The branch-point divergence argument depends only on the *pairwise* structure of $U = \sum m_i m_j / r_{ij}$ — which is the same for all N.

3. The positive-mass condition $m_i m_j > 0$ ensures that no cancellations occur in the monodromy argument — regardless of the configuration's symmetry or degeneracy structure.

Uniformity in d. The proof nowhere uses the ambient dimension except through the *scalar* complexified squared distance $R_{ij}(z) = \sum_{\alpha=1}^{d} (\Delta x_{ij}^{(\alpha)}(z))^2$ (equation (R)) and the *number of components* in each body's gradient vector. Concretely:

1. **Setup** (§2.0–2.2): the complexified squared distance $R_{ij}(z) = \sum_{\alpha} (\Delta x_{ij}^{(\alpha)}(z))^2$ is dimensionally generic; the planar factorization $R_{ij} = \varphi_{ij}\psi_{ij}$ is a convenience of $d = 2$ but plays no role in the proof.
2. **Monodromy (Lemma A)** and **divergence (Lemma C)**: both depend only on the branch structure of $R_{ij}(z)^{-1/2}$ around its zeros — unchanged across d.
3. **Gradient pole obstruction (Lemma D)**: body i's gradient in d dimensions is $\nabla_i U \in \mathbb{R}^d$ with d components; at a collision where body i has a unique partner j, the private $R^{-3/2}$ pole appears in *at least one* component of $\nabla_i U$ and cannot be cancelled by any other pair term, regardless of how many components d contains. Since the full CC condition requires every component to vanish, Lemma D fires in every dimension.
4. **Cluster analysis (Propositions F, G′, G)**: Remark F.2 shows that the transverse gradient components ($\alpha \geq 2$) vanish identically at leading order at any cluster collision, forcing the cluster into a common line (the collinear case $d = 1$); the same reduction is documented in Remark G′.1 for $m = 4$ and carries through the inductive argument of Proposition G for $m \geq 5$.
5. **Finiteness assembly (§4 Step 8)**: the assembly uses compactness on bounded simply connected regions, local pole-clearing/Borel–Steinmetz, Łojasiewicz, and the identity theorem; all of these are dimensionally generic.

The single place where d could in principle matter is the *shape space* $\mathcal{S}$, which has dimension $dN - d - \binom{d}{2} - 1$ and varies with d; but we use $\mathcal{S}$ only to invoke compactness and the Łojasiewicz decomposition, and both properties hold for every $d \geq 2$. The net effect is that the *reduction strategy* is dimension-uniform: the local obstructions run by a single argument, with no spatial-vs-planar split of the kind required by the algebraic tradition.

Honest scope of the dimension-uniformity claim. This uniformity is a property of the *strategy*, not a completed unconditional proof, and it is not literally "every $d \geq 2$ without modification." Two dimension-sensitive points must be stated plainly. (i) For $d \leq 3$ the collinearity/tight-cluster reduction (Cor. D.1′) is argued directly, but for $d \geq 4$ the Witt index is ≥ 2, the null cone supports non-proportional isotropic displacements, and the exclusion of positive-dimensional null-line residual families is delegated to the bridge null_line_finiteness (Prop. D.4), which is *not* discharged and delivers only generic absence / isolated-at-special-mass. So the $d \geq 4$ case is *conditional*, not modification-free. (ii) The general-d leading-order excess count of Proposition E is $1 - d(d-1)/2$ (not a universal 1); the genuine overdetermination is recovered only after reduction to the collinear $d = 1$ sub-cluster. Both points leave the elementary machinery intact but mean the uniform-in-d finiteness *conclusion* inherits the same conditional status (via the Step-8 bridges) as every other dimension.

6.3 The Role of Positive Masses

The positivity of masses enters at five distinct points in the proof:

1. **Monodromy** (Lemma A): the sum $\sum_{S_{\mathrm{odd}}} m_i m_j / r_{ij}(0) = 0$ forces $S_{\mathrm{odd}} = \emptyset$ because each

term $m_i m_j > 0$.

2. **Gradient pole** (Lemma D): the singular coefficient $m_j \neq 0$ creates an uncancellable pole.
3. **Graph Laplacian** (Proposition H(iii)): the edge weights $m_k m_j / |c_k - c_j|^3 > 0$ make the Laplacian positive semidefinite, giving full column rank on the real axis.
4. **Borel at sub-cluster level** (Proposition G): the coefficients $m_i m_j > 0$ guarantee non-proportional exponents, so Borel's theorem applies.
5. **Inductive base** (Proposition F): the 3-body gradient system is inconsistent for positive masses.

For negative masses, all five mechanisms can fail simultaneously. This is consistent with Roberts' (1999) counterexample: $N = 5$ with one negative mass admits a 1-parameter family of CCs.

6.4 Formal Companion Layer and Verification Boundary

The proof is tracked by a machine-checkable companion layer rather than by prose alone. Each paper-facing claim is represented in one of two forms: (1) a *semantically typed kernel statement* under a canonical name, encoding the full mathematical content in the kernel language, or (2) a *named bridge label* with prose description, citation, dependency metadata, and a canonical identifier — serving as a tracked placeholder whose mathematical justification lives in the paper's prose sections (§§3–4) or in cited references. The deterministic projection formal_statements.md renders both forms — typed statements and bridge labels — together with their anchors and trust metadata, into a human-readable companion file. The manuscript's <!-- ssot: ... formal_ref=... --> annotations link the narrative theorem/lemma/proposition statements back to that projection.

This creates three distinct layers:

Layer	Role	Audit status
Manuscript prose	Human-readable proof narrative, motivation, local explanations	Checked against SSOT annotations
formal_statements.md	Deterministic English projection of kernel statements and bridge assumptions	Render-fidelity gate PASS
formal proof kernel + Lean export	Typed statements, dependency graph, proof skeleton, Lean-verifiable exports	Statement graph tracked; Lean export compiles with zero sorry

The current audit state is:

Gate	What it checks	Current result
G10d	formal_statements.md exists and is the rendered companion layer	PASS