hexagon:2611.00563v1

OBLIGATORY TRIPLE SYSTEMS: ATOMS, SPECTRA AND DECOMPOSITIONS

SAMUIL PETKOV

24 October 2026

Abstract. Erdős Problem 593 asks which finite triple systems embed in every triple system of uncountable chromatic number. Such finite systems are called obligatory. We describe their structure and the ways of decomposing them while keeping their original hyperedges. Every obligatory system without isolated points has canonical pieces, called atoms: single triples and expansions of finite 2-connected bipartite graphs, obtained by adding one new private point to each graph edge. A forest records how the atoms meet and reconstructs the system. We determine exactly which numbers of vertices, hyperedges and components can occur, and which atom counts occur in a connected system, including the parity restriction when the vertex–hyperedge incidence graph has just one independent cycle. For a fixed system, we consider decompositions into connected pieces that meet in at most one point and whose shared-point incidences form a forest. Each such decomposition is specified by independently partitioning the atoms through each shared point. These choices, ordered by merging blocks, form a product of partition lattices. An exact piece-count formula identifies the single local mergers. Li gave the first public classification. The appendix includes our self-contained proof, developed independently in parallel with Li's work. We also record the precise coverage of the kernel-checked Lean 4 development. All arguments are in ZFC.

2020 Mathematics Subject Classification. Primary 05C65; Secondary 05C15, 05C35, 05C40, 05C63, 06A07, 03E05.

Keywords. obligatory triple systems, hypergraph blocks, one-point amalgamation, canonical decomposition, finite parameter spectra, partition lattices.

Introduction

A triple system is a collection of distinct three-element hyperedges on a set of vertices. A finite triple system is *obligatory* if it embeds in every triple system of uncountable chromatic number. Erdős asked which finite systems have this property [1, 2]. The answer is determined by finite incidence structure. We study which numerical parameters can occur and how each system can be assembled from its canonical pieces.

For a finite graph J, let J^+ be its private-vertex expansion: each graph edge xy is replaced by $\{x, y, p_{xy}\}$ with a new private point p_{xy}. The graph we started from is the core of this expansion. A one-point amalgamation joins two otherwise disjoint systems by identifying one vertex of each. Let $\mathcal{B}$ be the smallest class containing all such expansions of finite bipartite graphs and all finite edgeless systems, closed under finite disjoint unions and one-point amalgamations, up to isomorphism. Write F° for the result of deleting the isolated vertices of F.

A triple system is linear if distinct hyperedges meet in at most one point. Its Levi graph joins each hyperedge to its incident points. A bridge is an incidence edge whose deletion disconnects its component; a Berge cycle is a cycle of this graph, with length counted in hyperedges.

Theorem A (Li's classification)**.** *For every finite triple system F, the following are equivalent:*

(1) F is obligatory;
(2) $F \in \mathcal{B}$;
(3) F° is linear, every hyperedge-node of its Levi graph is incident with a bridge, and every Berge cycle of F° has even length.

Our classification argument was developed independently in parallel with Li's work. The first public resolution is due to Li [3, Theorem 1.1]. Theorem A allows us to identify the canonical pieces, called atoms, and study decompositions that keep the original hyperedges. Here *finite spectra* means the possible numbers of vertices, hyperedges, components and atoms. This differs from the exact chromatic-cardinal avoidance spectra studied in Li [3, Part II].

The canonical normal form records each atom and the points it shares with other atoms. Joining atoms to these shared points gives the atom–point forest of Theorem 2.3. The numerical spectra then determine which counts can occur (Theorems 3.3 and 3.8). Every admissible value is attained. For example, a connected reduced system, meaning one without isolated points, with surplus $|V| - |E| = 8$ and one independent cycle in its Levi graph has exactly one, three or five canonical atoms. The missing even counts come from the parity of its cyclic core.

For a fixed system, we join its atoms into connected pieces, retaining their original hyperedges. Pieces may meet in at most one point, and the incidence graph between pieces and their shared points must be a forest. These are the supported one-point forest decompositions. Write $\mathcal{D}(F)$ for their set, $S(F)$ for the set of shared atom points, and $\mu(p)$ for the number of atoms through p. The factor at each shared point is the partition lattice of its incident atoms, ordered by merging blocks. These local choices give the order isomorphism

$$\mathcal{D}(F) \cong \prod_{p \in S(F)} \Pi_{\mu(p)},$$

stated in Theorem 4.1. The factors use the actual incident atoms, and the pieces retain their original hyperedges. A cover is a single step in the order with no intermediate decomposition; here it changes the partition at one shared point and reduces the number of pieces by one (Theorem 4.3). The product therefore describes changes to a fixed system without reassembling an arbitrary prescribed list of atoms.

These descriptions let us check membership in the class, recover the atoms, test proposed numerical parameters and count local coarsenings. Sections 3 and 4 give examples. We make no claim about the computational complexity of these tasks.

The decomposition uses classical graph and hypergraph block theory [4, 5, 13]. Komjáth's reduction already relates obligatoriness to nonseparable blocks [9]. We identify the allowable atom types and describe their supported decompositions

as partitions of the original hyperedges, with local choices at shared points. The partition-lattice facts used here are classical [6]; the general block and partition-lattice identities are not claimed as new.

The main text proves the finite structural results. Appendix A gives a complete proof of Theorem A: bipartite expansions provide the basic examples, one-point closure joins them, and a sequence lift with one new apex produces the avoidance hosts. The lift and bridge arguments are also central to Li [3, Sections 3–5]; our proof uses expansion fibres and a forest that records their shared points. The preserved Lean implementation uses a shift-graph obstruction in the negative direction. Its closure argument uses a layering that also handles singular cardinals. Appendix B explains the relation between these arguments and lists the precise formal statements checked.

1. Preliminaries

A *triple system* is a simple 3-uniform hypergraph $H = (V(H), E(H))$, so $E(H) \subseteq [V(H)]^3$; the latter denotes the set of three-element subsets of the vertex set. A colouring $c : V(H) \to \lambda$ is *proper* if no hyperedge has all its vertices of one colour. The chromatic number $\chi(H)$ is the least cardinal λ for which such a colouring exists.

An embedding of a triple system F in H is an injective map $f : V(F) \to V(H)$ satisfying $f[e] \in E(H)$ for every $e \in E(F)$. It need not be induced: additional host hyperedges on the image vertices are allowed. A finite triple system F is *obligatory* if every triple system H with $\chi(H) > \aleph_0$ contains an embedding of F. A point is *isolated* if it belongs to no hyperedge, and F° is obtained by deleting all isolated points.

All ordinary graphs in this paper are simple. A triple system is *linear* if any two distinct hyperedges meet in at most one point. Its *Levi graph* $I(F)$ records incidence: its two vertex classes are $V(F)$ and $E(F)$, and a point-node p is adjacent to a hyperedge-node e exactly when $p \in e$. A *bridge* is an edge whose deletion increases the number of connected components.

A *Berge cycle of length* $\ell \geq 2$ consists of distinct points $p_0, \dots, p_{\ell-1}$ and distinct hyperedges $e_0, \dots, e_{\ell-1}$ such that $p_i \in e_i \cap e_{i+1}$, with indices modulo ℓ. Equivalently, it is a simple cycle of length 2ℓ in $I(F)$. For graph terminology and the incidence-graph interpretation of hypergraph connectivity, see Diestel [5, Chapters 1 and 3] and Bahmanian and Šajna [13].

For a finite simple graph J, its *private-vertex expansion* J^+ has vertex set

$$V(J) \dot\cup \{p_a : a \in E(J)\}$$

and hyperedges $\{u, v, p_{\{u,v\}}\}$ for $\{u, v\} \in E(J)$, where the points p_a are new and pairwise distinct. A *one-point amalgamation* of vertex-disjoint triple systems F_0, F_1 is obtained by choosing $x_i \in V(F_i)$, identifying x_0 with x_1, and making no other identifications or new hyperedges. Finite disjoint unions and finite edgeless systems include the empty cases.

Finite reductions. We call a finite system *reduced* when it has no isolated points. Connectivity and the number of components refer to the Levi graph. For a reduced

system with at least one hyperedge, write

$$n = |V(F)|, \qquad m = |E(F)|, \qquad c = \#\pi_0(I(F)),$$
$$s = n - m, \qquad \beta = 2m - n + c.$$

The last quantity counts the independent cycles in the finite Levi graph. Every decomposition below partitions the original hyperedge set. A piece contains exactly the points incident with its hyperedges.

2. The intrinsic decomposition

We prove that the construction rule and the incidence conditions in Theorem A are equivalent. Deleting the bridges of the Levi graph gives expansion pieces; its remaining forest tells us how to join them one point at a time. This is the explicit argument used by the Lean development. For a different structural formulation, compare Li [3, Sections 4–5]. Lemma A.1 handles isolated points.
Forward preservation.

Proposition 2.1 (The intrinsic conditions are preserved by the generators and operations)**.** *Every generator of $\mathcal{B}$ satisfies the intrinsic conditions, and the conditions are preserved under finite disjoint unions and one-point amalgamations.*

Proof. For a bipartite graph J, the expansion J^+ is linear, and the incidence from each hyperedge-node to its private point is a bridge. A Berge cycle in J^+ cannot pass through a private point. It therefore follows a graph cycle of J, with the same length, which is even. Edgeless systems satisfy the conditions vacuously, and disjoint union preserves them component by component.

In a one-point amalgamation, hyperedges from different factors meet only at the identified point, so linearity is preserved. The Levi graph is the vertex-sum of the two Levi graphs at that point-node. A simple cycle stays in one factor, since moving between factors would require a second visit to the cut vertex. Thus the old bridges remain bridges, and every Berge cycle keeps its parity. This is the usual cut-vertex argument for graph blocks [5, Chapter 3]; the hypergraph terminology is discussed in Bahmanian and Šajna [13]. □

Thus

$$(2) \Longrightarrow (3). \tag{2.1}$$

Bridge-block reconstruction.

Proposition 2.2 (Bridge-block decomposition)**.** *Suppose F is finite, has no isolated vertices, and satisfies the three intrinsic conditions. Then $F \in \mathcal{B}$.*

Proof. If $E(F) = \varnothing$, then the assumption that F has no isolated vertices gives $V(F) = \varnothing$, and F is an edgeless generator. Hence assume $E(F) \neq \varnothing$. Put $L = I(F)$, let $B(L)$ be its set of bridges, and let $\mathscr{C}$ be the set of connected components of $L - B(L)$.

Consider a hyperedge-node e. The node e has degree three and is incident with a bridge, so at most two of its incidences are nonbridges. Exactly one is impossible: a

cycle containing that incidence must leave e through another nonbridge incidence. Each hyperedge-node therefore has either zero or two nonbridge incidences.

The nontrivial components. Let $C \in \mathscr{C}$ contain a nonbridge incidence. Every hyperedge-node $e \in C$ has exactly two point-neighbours in C. Contract each such e to an ordinary graph edge between those two point-neighbours. This produces a finite graph J_C.

The graph J_C is simple: parallel edges would correspond to two hyperedges of F sharing two points, contrary to linearity. Every cycle in J_C gives a Berge cycle of the same length in F. Hence J_C has no odd cycle and is bipartite.

Each hyperedge-node $e \in C$ has a third point-neighbour p_e across a bridge. These third points satisfy:

1. $p_e \notin C$; otherwise a path in C from p_e to e, together with the incidence ep_e, would place that incidence on a cycle.
2. If $e \neq f$ lie in C, then $p_e \neq p_f$; otherwise a path in C from e to f, together with ep_e and $p_f f$, would place those bridge incidences on a cycle.

We now identify the expansion J_C^+ with the subsystem P_C formed by these hyperedges. Map a graph vertex of J_C to its point-node in C, the private vertex corresponding to e to p_e, and the expanded edge to the hyperedge e of F. The two observations above make the vertex map injective: the private points are distinct and lie outside C. Every e consists of its two contracted endpoints and p_e, so the map preserves incidence and covers every vertex and hyperedge of P_C. It is therefore an isomorphism $J_C^+ \cong P_C$.

Components consisting of one hyperedge-node. If a hyperedge-node e has no nonbridge incidence, it is an isolated component after the bridges are deleted. Its single triple is the expansion of a one-edge bipartite graph. Thus every hyperedge of F belongs to a piece P_C isomorphic to J^+ for a finite bipartite graph J. The hyperedge sets of these pieces partition $E(F)$. Their vertex sets cover $V(F)$. Indeed, every point is incident with a hyperedge. The point either lies in that hyperedge-node's bridge-deleted component or is the point endpoint of a bridge leaving it.

Quotient forest. Contract every component of $L - B(L)$. The original bridges become edges of a quotient multigraph T. Two distinct bridges cannot join the same pair of bridge-deleted components: paths inside those two components, together with the two bridges, would form a cycle containing both, contradicting that they are bridges. Thus T is a simple graph.

To see that T is a forest, suppose that it contained a cycle. Within each contracted component on this cycle, choose a path joining the endpoints of the two incident quotient edges. Concatenating these paths with the corresponding bridges gives a closed walk in L that traverses each of those bridges exactly once. This is impossible: deleting a bridge separates its endpoints, so every closed walk crosses the resulting cut an even number of times. Hence T is a forest.

A vertex $C \in V(T) = \mathscr{C}$ is called *active* if the bridge-deleted component C contains a hyperedge-node. For an active C, the piece P_C consists of the hyperedges represented in C, all point-nodes of C, and the point endpoints of bridge incidences

leaving hyperedge-nodes in C. An inactive component contains no hyperedge-node and hence no incidence of $L - B(L)$, so it is a singleton point-node.

Every hyperedge now lies in a bipartite expansion piece. To reconstruct the system by one-point amalgamations, we need an order in which each new piece meets the preceding union in at most one point. The quotient forest provides this order.

Claim 2.2.1 (Running intersection)**.** *Every component of T contains an active vertex, because F has no isolated vertices. Choose an active root in each component, and order the active vertices by nondecreasing depth, with arbitrary order among equal-depth vertices. When an active component C is added, its piece P_C meets the union of the earlier pieces in at most one point.*

Proof of the claim. For a point $p \in V(F)$, let $X(p) \in \mathscr{C}$ be the component of $L - B(L)$ containing the point-node p. The active pieces containing p are precisely $P_{X(p)}$, when $X(p)$ is active, together with the pieces P_D at active neighbours D of $X(p)$ whose corresponding original bridge has point endpoint p. Thus the vertices of T indexing pieces that contain p are contained in the closed star centred at $X(p)$.

Suppose that $p \in P_C$ also lies in an earlier piece. If $X(p) = C$, every other piece containing p is indexed by a neighbour of C. A tree has only one neighbour of C at smaller depth, namely its parent, so only that piece can be earlier. If $X(p) \neq C$, then C is a neighbour of $X(p)$. The vertex $X(p)$ cannot be a child of C, because then $X(p)$ and all its children would be deeper than C, leaving no earlier piece containing the point. Hence $X(p)$ is the parent of C. In either case, p is the point endpoint of the unique parent edge of C. The same reasoning applies to every point shared with an earlier piece, so all such points coincide with this parent point p_C, and

$$P_C \cap \bigcup_{D \text{ earlier}} P_D \subseteq \{p_C\}.$$

For a root active component the intersection is empty. □

Start with the root piece of each component and add the active pieces in the chosen order. An empty intersection calls for a disjoint union; an intersection equal to $\{p_C\}$ calls for a one-point amalgamation. Several earlier pieces may already contain the same inactive attachment point. This causes no extra identification: the new step still uses only that one existing vertex.

The identity maps from the pieces into F assemble inductively to an incidence isomorphism onto the union of the pieces already added. By the running-intersection claim, the next piece and the preceding union are either disjoint or meet exactly at the point used in the amalgamation, so no unintended identification is introduced. The piece hyperedge sets are disjoint and partition $E(F)$, and their vertex sets cover $V(F)$. After all active pieces have been added, the resulting subsystem is F. Different components of T are combined by disjoint union, proving $F \in \mathcal{B}$. □

This proves

$$(3) \Longrightarrow (2). \tag{2.2}$$

Combining (2.1) and (2.2),

$$(2) \iff (3). \tag{2.3}$$

Theorem 2.3 (Canonical atom normal form)**.** *Let F be a finite triple system without isolated points satisfying the three intrinsic conditions in Theorem A. Its hyperedges admit a canonical partition into* atoms *with the following properties.*

(1) Every atom is either one triple or J^+ for a finite 2-connected simple bipartite graph J.

(2) Two distinct atoms meet in at most one point. The bipartite incidence graph between atoms and points belonging to at least two atoms is a forest.

(3) The atom partition, the nontrivial cores J up to graph isomorphism, and the atom–point forest are determined by F.

Conversely, every forest assembly of these atoms by disjoint unions and one-point amalgamations satisfies the intrinsic conditions.

Proof. Let $L = I(F)$. A *cyclic block* is a maximal 2-connected subgraph of L containing a cycle. Since every hyperedge-node has degree three and is incident with a bridge, its degree inside a cyclic block is two. Suppressing the hyperedge-nodes in such a block C gives a finite graph J_C on the point-nodes of C. It is simple by linearity and 2-connected by standard block theory [4, 5]. Its cycles correspond to Berge cycles of the same lengths, so J_C is bipartite.

For a hyperedge-node $e \in C$, the third incidence is a bridge. Its point endpoint lies outside C, since otherwise that incidence would lie on a cycle. These third points are pairwise distinct as e varies in C; otherwise a path inside C between two hyperedge-nodes, completed through the common third point, would put both bridge incidences on a cycle. The hyperedges represented in C therefore form exactly J_C^+.

A hyperedge-node outside every cyclic block has no nonbridge incidence and contributes one single-triple atom. Thus the atoms partition the hyperedges. Two atoms cannot share two points, since paths inside the two atoms would produce a Levi cycle crossing two cyclic blocks. Similarly, a cycle in the atom–shared-point incidence graph expands to a closed Levi walk and hence to a Levi cycle crossing atom boundaries. The incidence graph is therefore a forest. Cyclic blocks and the block–cut forest are canonical, proving the forward statement. The converse follows because an atomic expansion is linear, has a private bridge at each hyperedge-node, and has only even Berge cycles, while a forest of one-point sums preserves these properties. □

The underlying graph operations are classical. Tarjan's algorithm finds bridges and bridge-connected components [14], and depth-first search finds biconnected components [15]. These algorithms are separate from the structural identification proved above. We do not assert a verified running-time bound for the Lean development.

Call a connected triple system with at least one hyperedge *one-point indecomposable* if it cannot be written as a one-point amalgamation of two systems that both contain a hyperedge.

Corollary 2.4 (Minimal generators and one-point indecomposables)**.** *Up to isomorphism, the class $\mathcal{B}$ is generated by finite edgeless systems, one triple, and J^+ for finite 2-connected simple bipartite graphs J, under finite disjoint unions and one-point amalgamations. A connected reduced obligatory system with at least one hyperedge is one-point indecomposable if and only if it is one triple or J^+ for a finite 2-connected simple bipartite graph J.*

Proof. The theorem decomposes every reduced member of $\mathcal{B}$ into the displayed atoms; isolated points are restored by an edgeless factor. For the indecomposability assertion, the system is connected. If the atom forest has at least two atom-nodes, a leaf atom separates from the rest at its unique shared point. Conversely, one triple is indecomposable, and if J is 2-connected then deleting any point of J^+ cannot separate its hyperedge-nodes into two nonempty classes: a private point is a leaf, while a core point leaves J connected. Hence J^+ is indecomposable. □

Remark 2.5 (Blocks and the earlier reduction). Komjáth proved that a finite triple system is obligatory if and only if all of its 2-connected components are obligatory [9]. General hypergraph block decomposition and its relation to the incidence graph are classical; see Whitney [4], Diestel [5, Chapter 3], and, for the hypergraph formulation, Bahmanian and Šajna [13]. For the reduced 3-uniform systems considered here, the canonical atoms in Theorem 2.3 are precisely these maximal nonseparable edge-supported blocks. The canonical atom theorem identifies the allowable block types, and the finite spectra below quantify the systems assembled from those types.

Corollary 2.6 (Block formulation of the classification)**.** *Let F be a finite triple system. After isolated points are deleted, F is obligatory if and only if every nonempty block of F° is either one triple or J^+ for a finite 2-connected simple bipartite graph J.*

Proof. If $E(F^\circ) = \varnothing$, then F° is empty and the assertion follows from Lemma A.1. Otherwise suppose first that F is obligatory. Lemma A.1 gives that F° is obligatory, and Komjáth's reduction gives obligatoriness of each of its blocks. A block is one-point indecomposable, so Corollary 2.4 gives exactly the displayed list.

Conversely, each displayed block is obligatory by the positive expansion result (the one-triple case included). Komjáth's reduction therefore makes F° obligatory, and Lemma A.1 restores the isolated points. □

3. Finite parameter consequences

We now determine the possible numbers of vertices, hyperedges and components. A triple system is *connected* when its Levi graph is connected. We first work without isolated vertices. For an integer $r \geq 1$, set

$$q(r) = \lceil 2\sqrt{r} \rceil .$$

Proposition 3.1 (Edge-deletion form of the bridge condition)**.** *Let F be a finite triple system and let $e = \{x, y, z\} \in E(F)$. The hyperedge-node e has no incident bridge in $I(F)$ if and only if x, y, z lie in one connected component after the hyperedge e is deleted and all points are retained. Consequently the bridge condition in Theorem A is equivalent to the following hypergraph-native condition:*

$$\textit{for every } e \in E(F^{\circ}),\ \textit{deleting } e$$

$$\textit{separates its three points into at least two components.}$$

Proof. Fix $x \in e$. The Levi incidence xe is not a bridge precisely when there is an alternative path from e to x after that incidence is deleted. A simple such path leaves e through y or z and thereafter lies entirely in the Levi graph of $F - e$. Thus xe is non-bridging exactly when x is connected in $F - e$ to at least one of the other two points of e.

If all three incidences at e are non-bridging, each of x, y, z is connected to another of these three points after the hyperedge is removed. Their component partition therefore has no singleton block. For three points this forces one block, so all three are connected in $F - e$. The converse follows by reversing the same paths. □

Lemma 3.2 (Bipartite shadow)**.** *Let F be an obligatory finite triple system with no isolated vertices and at least one hyperedge. There is a finite bipartite graph J, without isolated vertices, such that*

$$|E(J)| = |E(F)|, \qquad |V(J)| = |V(F)| - |E(F)|,$$

and J has the same number of connected components as F. Conversely, for every finite bipartite graph J without isolated vertices, the expansion J^{+} is obligatory and has these parameter relations.

Proof. By Theorem A and Proposition 2.2, each connected component H of F is obtained by one-point amalgamations from connected expansion pieces

$$J_1^{+}, \ldots, J_k^{+},$$

where each J_i is a connected bipartite graph. Disconnected graph pieces may simply be split into their connected components. Write $m_i = |E(J_i)|$ and $s_i = |V(J_i)|$. Since H is connected, its k pieces are joined by exactly $k - 1$ one-point identifications. Hence

$$|E(H)| = \sum_{i=1}^{k} m_i, \qquad |V(H)| = \sum_{i=1}^{k} (m_i + s_i) - (k - 1). \tag{3.1}$$

Take one-point sums of the ordinary graphs $J_1, \ldots, J_k$ along any tree on the k indices. Before each identification, interchange the two bipartition classes of the new

factor if necessary. The resulting graph J_H is connected and bipartite, and (3.1) gives

$$|E(J_H)| = |E(H)|, \qquad |V(J_H)| = |V(H)| - |E(H)|.$$

Take the disjoint union of these shadows over the components of F. This proves the forward assertion. Conversely, expanding a bipartite graph gives an obligatory system by Theorem A, and the definition of J^+ gives the stated counts. □

Theorem 3.3 (Exact order–size–component spectrum)**.** *Let $m \geq 1$, $1 \leq c \leq m$, and n be integers. There exists an obligatory triple system F with no isolated vertices, exactly m hyperedges, exactly n vertices, and exactly c connected components if and only if*

$$\boxed{m + 2(c-1) + \left\lceil 2\sqrt{m-c+1} \right\rceil \leq n \leq 2m + c.} \tag{3.2}$$

Every integer in this interval occurs.

Proof. We begin with the corresponding elementary graph fact. A connected simple bipartite graph with $r \geq 1$ edges and s vertices exists if and only if

$$q(r) \leq s \leq r + 1. \tag{3.3}$$

Indeed, connectedness gives $r \geq s - 1$. If the bipartition has sizes a, b, then

$$r \leq ab \leq \left\lfloor \frac{s^2}{4} \right\rfloor,$$

which is equivalent to $s \geq q(r)$. Conversely, $K_{\lfloor s/2 \rfloor, \lceil s/2 \rceil}$ contains a spanning tree, and starting with that tree and adding arbitrary unused edges realises every edge count between $s - 1$ and $\lfloor s^2/4 \rfloor$.

Apply Lemma 3.2. Let the c connected components of the shadow have positive edge counts $m_1, \ldots, m_c$, where $\sum_i m_i = m$. If their orders are s_i, then (3.3) gives

$$\sum_{i=1}^{c} q(m_i) \leq \sum_{i=1}^{c} s_i \leq m + c. \tag{3.4}$$

For positive integers a, b,

$$q(a) + q(b) \geq q(a + b - 1) + 2. \tag{3.5}$$

To see this, note that

$$\sqrt{a} + \sqrt{b} \geq \sqrt{a + b - 1} + 1,$$

because, after squaring, the assertion reduces to $(a-1)(b-1) \geq 0$; taking ceilings gives (3.5). Repeatedly applying (3.5) to the left side of (3.4) yields

$$\sum_{i=1}^{c} q(m_i) \geq q(m - c + 1) + 2(c - 1).$$

The shadow has $n - m$ vertices, so these inequalities give (3.2).

To attain the full interval, put $M = m - c + 1$ and take $c - 1$ components equal to K_2. By (3.3), a final connected bipartite component with M edges can have any order from $q(M)$ through $M + 1$. The disjoint union therefore has every shadow order from $q(M) + 2(c-1)$ through $m + c$. Its expansion adds exactly m private vertices, attaining every integer in (3.2). □

Corollary 3.4 (Connected and unrestricted spectra)**.** *For $m \geq 1$:*

(1) a connected obligatory triple system without isolated vertices has m hyperedges and n vertices if and only if

$$m + \lceil 2\sqrt{m} \rceil \leq n \leq 2m + 1;$$

(2) without prescribing the number of components, a reduced obligatory triple system has m hyperedges and n vertices if and only if

$$m + \lceil 2\sqrt{m} \rceil \leq n \leq 3m;$$

(3) if isolated vertices are permitted, an obligatory triple system with m hyperedges and n vertices exists if and only if

$$n \geq m + \lceil 2\sqrt{m} \rceil .$$

Proof. The first assertion is Theorem 3.3 with $c = 1$. For the second, the connected interval covers through $2m + 1$; for $n > 2m + 1$, take $c = n - 2m$ and use the upper endpoint $n = 2m + c$ of (3.2). The largest possible value is obtained from m disjoint triples. The final assertion follows by adjoining arbitrary isolated vertices. □

Corollary 3.5 (Size spectrum at fixed order)**.** *Let $c \geq 1$ and $n \geq 3c$. A reduced obligatory triple system with exactly n vertices, c connected components, and m hyperedges exists if and only if*

$$\boxed{\left\lceil \frac{n-c}{2} \right\rceil \leq m \leq n - 2c + 4 - \left\lceil 2\sqrt{n - 3c + 4} \right\rceil .} \tag{3.6}$$

Consequently, a reduced obligatory triple system with exactly c connected components and n vertices exists if and only if

$$n = 3c, \qquad n = 3c + 2, \qquad \text{or} \qquad n \geq 3c + 4. \tag{3.7}$$

Proof. The upper-order inequality in (3.2) is equivalent to $m \geq \lceil (n - c)/2 \rceil$. For the other inequality put

$$M = m - c + 1, \qquad N = n - 3c + 3.$$

It becomes

$$M + \left\lceil 2\sqrt{M} \right\rceil \leq N. \tag{3.8}$$

For integers $N \geq 3$ and $M \geq 1$, the largest M satisfying (3.8) is

$$U = N + 2 - \left\lceil 2\sqrt{N + 1} \right\rceil . \tag{3.9}$$

Indeed, write $k = \lceil 2\sqrt{N + 1} \rceil$. Since $k^2 \geq 4(N + 1)$,

$$4U = 4N + 8 - 4k \leq (k - 2)^2,$$

so $\lceil 2\sqrt{U} \rceil \leq k - 2$ and $U + \lceil 2\sqrt{U} \rceil \leq N$. On the other hand, $(k - 1)^2 < 4(N + 1)$ gives

$$(k - 3)^2 < 4(N + 3 - k) = 4(U + 1),$$

so $\lceil 2\sqrt{U + 1} \rceil \geq k - 2$ and $(U + 1) + \lceil 2\sqrt{U + 1} \rceil \geq N + 1$. The left side of (3.8) is strictly increasing in M, proving (3.9). Translating back to m gives (3.6).

For (3.7), Corollary 3.4 shows that a connected reduced obligatory system has order 3, order 5, or any order at least 7. The first two statements come from $m = 1, 2$. For $m \geq 3$, the connected intervals have no gaps because $\lceil 2\sqrt{m+1} \rceil \leq m+1$ for $m+1 \geq 4$. Thus every component has order

$$3 + r, \qquad r \in \{0, 2\} \cup [4, \infty).$$

The sum of c such increments is again 0, 2, or any integer at least 4: use one nonzero increment and take all remaining increments to be zero. The increments 1 and 3 are impossible. This proves (3.7). □

Corollary 3.6 (Exact Levi cycle-rank spectrum)**.** *Let F be reduced and obligatory, with m hyperedges, n vertices, and c connected components. The cyclomatic number of its Levi graph is*

$$\beta(I(F)) = |E(I(F))| - |V(I(F))| + c = 2m - n + c. \tag{3.10}$$

For fixed integers $m \geq 1$ and $1 \leq c \leq m$, every integer in the interval

$$0 \leq \beta(I(F)) \leq m - c + 2 - \left\lceil 2\sqrt{m-c+1} \right\rceil \tag{3.11}$$

occurs, and no other value occurs. In particular,

$$|V(F)| = 2|E(F)| + c \quad \Longleftrightarrow \quad I(F) \text{ is a forest.}$$

Proof. The Levi graph has $n + m$ vertices, $3m$ incidence edges, and c components, which gives (3.10). Substitution of the two endpoints in (3.2) gives (3.11), and every intermediate value occurs because every intermediate order does. A finite graph has cyclomatic number zero exactly when it is a forest. □

3.1. **Canonical atoms and structural spectra.** For a connected reduced obligatory system F, put

$$s = |V(F)| - |E(F)|, \qquad \beta = 2|E(F)| - |V(F)| + 1.$$

Thus s is the order of a connected bipartite shadow and β is its cycle rank. Let $a(F)$ denote the number of canonical atoms from Theorem 2.3.

Lemma 3.7 (Minimum order at fixed cyclic rank)**.** *Let J be a finite 2-connected simple bipartite graph of cycle rank $r = |E(J)| - |V(J)| + 1 \geq 1$. Then*

$$|V(J)| \geq 2 + q(r). \tag{3.12}$$

For $r \geq 2$, every order $v \geq 2 + q(r)$ occurs. For $r = 1$, the possible orders are exactly the even integers $v \geq 4$.

Proof. Writing $v = |V(J)|$, the bipartite extremal bound gives

$$r \leq \left\lfloor \frac{v^2}{4} \right\rfloor - v + 1 = \left\lfloor \frac{(v-2)^2}{4} \right\rfloor,$$

which is equivalent to (3.12). If v is even, start with the spanning cycle C_v in the balanced complete bipartite graph and add edges until the desired rank is reached. If v is odd, start with C_{v-1} and add the remaining vertex with two neighbours in the opposite bipartition class; this is 2-connected and has rank two. Add arbitrary

missing cross-edges. These constructions realise every allowed $r \geq 2$. A 2-connected graph of rank one is a cycle, which is bipartite exactly when its order is even. □

Theorem 3.8 (Exact canonical atom-count spectrum)**.** *Fix nonnegative integers s, β satisfying either $\beta = 0$ and $s \geq 2$, or $\beta \geq 1$ and $s \geq 2 + q(\beta)$. Let F range over connected reduced obligatory systems with at least one hyperedge, shadow order s and cycle rank β. The possible values $k = a(F)$ are exactly*

$$\begin{aligned} &\beta = 0: \quad k = s - 1, \\ &\beta = 1: \quad 1 \leq k \leq s - 3, \quad k \equiv s + 1 \pmod 2, \\ &\beta \geq 2: \quad 1 \leq k \leq s - 1 - q(\beta). \end{aligned} \tag{3.13}$$

For $\beta \geq 1$, the maximum is

$$a_{\max}(s, \beta) = s - 1 - q(\beta). \tag{3.14}$$

Every maximiser has one minimum-order cyclic atom carrying the full rank β, and every other atom is one triple. The parity obstruction in the rank-one line is the only gap in any atom-count spectrum.

Proof. Represent each single-triple atom as the expansion of a one-edge graph. Let the atom cores be J_i, with $v_i = |V(J_i)|$, $e_i = |E(J_i)|$, and $r_i = e_i - v_i + 1$. Since a tree of k atoms makes $k - 1$ one-point identifications,

$$s = 1 + \sum_i (v_i - 1), \qquad \beta = \sum_i r_i. \tag{3.15}$$

If $\beta = 0$, every atom is one triple and (3.15) gives $k = s - 1$. If $\beta = 1$, there is one cyclic atom, whose core is an even cycle of order $v \geq 4$; the remaining atoms are single triples, so $s = k + v - 1$. This is exactly the middle line of (3.13).

Suppose $\beta \geq 2$. For positive integers a, b, the inequality $\sqrt{a} + \sqrt{b} \geq \sqrt{a + b - 1} + 1$ follows after squaring from $(a - 1)(b - 1) \geq 0$. Hence $q(a) + q(b) \geq q(a + b - 1) + 2 \geq q(a + b) + 1$, where the second inequality uses that q increases by at most one between consecutive integers. Apply this inequality to the positive-rank atom cores. Together with Lemma 3.7, it gives

$$s \geq 1 + k + q(\beta),$$

so $k \leq s - 1 - q(\beta)$. Conversely, for every integer in the displayed interval choose one cyclic core of rank β and order $s - k + 1$, supplied by the lemma, and attach $k - 1$ single-triple atoms. Combining two positive-rank cores incurs a strict surplus in the preceding bound. Equality in (3.14) therefore forces one positive-rank atom of minimum order, proving the rigidity statement. □

The equality case determines how cyclic rank is distributed: all positive rank lies in one atom of minimum order. It does not determine the graph-isomorphism type of that atom's core from the numerical parameters alone.

Example 3.9 (A parity-sensitive feasibility test). Take a connected reduced obligatory system with $s = 8$ and $\beta = 1$. Theorem 3.8 gives $k \in \{1, 3, 5\}$, and each value occurs. The corresponding cyclic cores can be taken to be C_8, C_6, C_4, respectively, with zero, two or four single-triple atoms attached by one-point amalgamations. Each system

has eight hyperedges and sixteen vertices. The even values of k are excluded by the rank-one parity condition, rather than by a failure of the size bound.

4. Supported decompositions and exact piece counts

We now keep a system fixed and ask which of its atoms can be joined into larger pieces. At each shared point, we record which incident atoms belong together. Throughout this section F is reduced, obligatory and has at least one hyperedge. A *supported piece* is a nonempty set of original hyperedges together with all their incident points. Thus neither the hyperedges nor their vertex incidences are changed.

Theorem 4.1 (Local product for supported decompositions)**.** *Let F be reduced and obligatory with at least one hyperedge. Let $\mathcal{A}(F)$ be its canonical atoms, let $S(F)$ be the points belonging to at least two atoms, and put*

$$\mu(p) = |\{A \in \mathcal{A}(F) : p \in V(A)\}|.$$

Let $\mathcal{D}(F)$ be the set of partitions of $E(F)$ into nonempty connected supported pieces such that two pieces meet in at most one point and the piece–shared-point incidence graph is a forest, ordered by refinement. Then $\mathcal{D}(F)$ is a lattice and

$$\mathcal{D}(F) \cong \prod_{p \in S(F)} \Pi_{\mu(p)}, \tag{4.1}$$

where Π_r denotes the partition lattice of an r-element set. The empty product in (4.1) *is the one-element lattice.*

Proof. A supported one-point forest decomposition cannot split a block. Suppose that two hyperedges of one block lie in different pieces. The unique path between those pieces in the piece–shared-point forest contains a point whose deletion separates the two hyperedges. This contradicts nonseparability of the block. Here deletion removes a point-node of the Levi graph and retains all hyperedge-nodes. Such a deletion cannot separate the hyperedge-nodes of an atom: this is immediate for a single triple and follows from 2-connectivity of the core for an expansion J^+. Thus every element of $\mathcal{D}(F)$ is a coarsening of the canonical atom decomposition.

For a shared point p, put two incident atoms in the same local block when they lie in the same supported piece. This partitions the $\mu(p)$ atoms through p and gives an element of $\Pi_{\mu(p)}$. Conversely, choose a partition at each shared point and generate an equivalence relation on all atoms from these local identifications. We describe its classes with an auxiliary forest. At each p, retain a central node p, insert a node (p, B) for every local block B, join p to (p, B), and join (p, B) to each atom in B. Replacing a star of a forest by this tree preserves acyclicity, so the auxiliary graph is again a forest.

Delete the edges from the central nodes p to their nodes (p, B). The components containing atoms are vertex-disjoint in the auxiliary graph and correspond exactly to the generated equivalence classes. The union of the original hyperedges in each class, with their unchanged incident points, is a connected supported piece. Two different blocks at one p cannot lie in the same component, since a path between them, together with their two edges to p, would give a cycle in the auxiliary graph. Restore

the deleted edges and contract these components in the full auxiliary graph, retaining the central nodes p. The preceding acyclicity argument prevents repeated incidences between a contracted piece and p. Delete each degree-one central node, which is now internal to one piece. The resulting graph is precisely the piece–shared-point incidence graph and is still a forest. In particular, two pieces cannot meet at two shared points. Restriction recovers the prescribed local partitions. Conversely, a connected supported piece is recovered from its local identifications, since its atom nodes are connected through the shared points it contains. Thus the constructions are inverse and preserve refinement, proving (4.1). For the related theory of connected set partitions and bond lattices, see Simon et al. [16] and Stanley [17, Lecture 2]. The factors here are full partition lattices on the atoms through each shared point, not the bond lattice of the incidence forest itself. A point shared by two atoms contributes a factor that is a two-element chain. Three atoms through one point give the five-element diamond $\Pi_3 \cong M_3$, which is not distributive [18, Chapter II, Theorem 101].

□

Theorem 4.2 (Exact piece accounting)**.** *Let $k = |\mathcal{A}(F)|$. For $D \in \mathcal{D}(F)$, let $p(D)$ be its number of actual pieces, and let $b_p(D)$ be the number of blocks of its local partition at $p \in S(F)$. Define*

$$\delta(D) = \sum_{p \in S(F)} (\mu(p) - b_p(D)).$$

Then every summand is nonnegative, and

$$p(D) + \delta(D) = k. \tag{4.2}$$

In particular, $\delta(D) = k - p(D)$ is the actual piece deficit.

Proof. A partition of a $\mu(p)$-element set has at most $\mu(p)$ blocks. Use the auxiliary forest from the proof of Theorem 4.1. The canonical atom–point forest has $k + |S(F)|$ vertices and c components, so $\sum_p \mu(p) = k + |S(F)| - c$. Subdividing each shared-point star by the local block nodes preserves its component count. Deleting the $\sum_p b_p(D)$ edges from the central point nodes to their block nodes increases the number of components by that amount. Exactly $|S(F)|$ components are now isolated central point nodes; the others contain the atoms of the actual pieces. Hence

$$p(D) = c + \sum_p b_p(D) - |S(F)| = k - \sum_p (\mu(p) - b_p(D)),$$

which proves (4.2). □

Theorem 4.3 (Covers and local mergers)**.** *For $D, D' \in \mathcal{D}(F)$, the following are equivalent:*

(1) D' covers D in refinement order;
(2) exactly one shared point has a changed local partition, and there the change merges two blocks;
(3) $D < D'$ and $p(D) = p(D') + 1$;
(4) $D < D'$ and $\delta(D') = \delta(D) + 1$.

The shared point in the second condition is unique. Thus an upward cover merges exactly two actual original-edge pieces.

Proof. In a finite product, a cover changes exactly one coordinate, by a cover in that factor. In a partition lattice this means merging two blocks: merging more leaves an intermediate partition, while merging two leaves none. The order isomorphism in Theorem 4.1 proves the first two conditions equivalent and makes the shared point unique. Theorem 4.2 then shows that a cover reduces $p(D)$ by one and increases $\delta(D)$ by one.

Any strict coarsening reduces the number of pieces: every local block count weakly decreases, and at least one decreases strictly. Thus, if $D < D'$ and the piece count drops by one, there can be no intermediate element, since that would require two strict drops. This proves the converse from the third condition. The third and fourth conditions are equivalent by (4.2). □

Proposition 4.4 (Least supported completion of merger requests)**.** *Fix $D_0 \in \mathcal{D}(F)$ and sets $R^+, R^- \subseteq E(F) \times E(F)$. We seek a coarsening $D \geq D_0$ that puts each pair in R^+ in one piece and each pair in R^- in different pieces. Splitting a piece of D_0 is not allowed. Let Q_0 be the actual piece–shared-point forest of D_0, and let $\rho(e)$ be the unique D_0-piece containing the original hyperedge e.*

A positive request joining different components of F makes the problem infeasible. Otherwise, for every point x shared by D_0-pieces, form a graph H_x on the entire set $\mathcal{C}_x = \{C \in D_0 : x \in V(C)\}$, including isolated vertices. For each $(e, f) \in R^+$, follow the unique path in Q_0 from $\rho(e)$ to $\rho(f)$ and insert an edge CC' in H_x at each turn $C - x - C'$. Write $\nu_x = |\mathcal{C}_x|$, let h_x be the number of components of H_x, and put $M = \sum_x (\nu_x - h_x)$.

There is a unique finest supported coarsening D_* satisfying R^+: for every $D \geq D_0$,

$$D \text{ satisfies } R^+ \quad \Longleftrightarrow \quad D_* \leq D.$$

The joint requests are feasible exactly when D_* separates every pair in R^-. In that case the minimum number of upward local covers is M, and

$$p(D_*) = p(D_0) - M, \qquad \delta(D_*) - \delta(D_0) = M, \qquad 0 \leq M \leq p(D_0) - c.$$

Here c is the number of components of F. All pieces are unions of unchanged original hyperedges.

Proof. A connected union of D_0-pieces contains the whole Q_0-path between any two of its pieces. Indeed, a Levi path inside the union gives a walk in Q_0 using only its piece-nodes, and a walk in a forest contains the unique path between its endpoints.

Therefore two starting pieces lie in the same coarsened piece exactly when every turn on their Q_0-path is identified locally. A path of length zero requires no identification.

At each x, combine the local blocks of D_0 according to the components of H_x. The starting block for $C \in \mathcal{C}_x$ consists of its canonical atoms through x; it is nonempty because C has an original edge through x. Other coordinates already have one block and remain unchanged. Theorem 4.1 reconstructs a supported coarsening D_*. The path observation shows that it satisfies every positive request. Conversely, any coarsening satisfying a request must identify every turn of its path, hence every component of every H_x. Coordinatewise refinement gives $D_* \leq D$.

If D_* merges a forbidden pair, every $D \geq D_*$ also merges it. Otherwise D_* satisfies both sets of requests. For a negative pair in one component, separation means that some turn on its path has endpoints in different components of H_x. A negative pair already inside one D_0-piece is therefore impossible to satisfy. A pair in different components is always separated.

The local block count drops from ν_x to h_x at x. Theorem 4.2 gives the two equalities. The forest Q_0 has $p(D_0) + |S_0|$ vertices and c components, where S_0 is its shared-point set. Hence $\sum_x(\nu_x - 1) = p(D_0) - c$, proving the bound on M. Choose a spanning tree in each nontrivial component of every H_x. Moves at other points leave this coordinate unchanged. The selected links, in any order, merge distinct current local blocks: a redundant link would create a cycle in that selected tree. Theorem 4.3 therefore gives exactly M covers ending at D_*. Every successful upward sequence ends at a coarsening of D_*, so must reduce the piece count by at least M. □

The finest completion depends on where the pieces meet. Start with three single-triple pieces, the sole positive request (e_1, e_3) and no negative requests. If all three triples share one point, the finest completion leaves e_2 separate. If instead consecutive triples meet at distinct points, forming the loose chain $e_1 - a - e_2 - b - e_3$, the same request forces both intervening turns and joins all three pieces. Counts of vertices, edges and atoms alone do not record this difference.

Corollary 4.5 (Minimum number of local changes)**.** *Give each cover or reverse cover of $\mathcal{D}(F)$ unit cost, and let $d(D, E)$ be the minimum number of these local changes from D to E. Write $D \vee E$ for the least common coarsening in $\mathcal{D}(F)$. Then*

$$d(D, E) = p(D) + p(E) - 2p(D \vee E).$$

A shortest path first merges from D to $D \vee E$, then splits to E. If F has k atoms and c components, the diameter of this graph of decompositions is $k - c$.

Proof. The distance formula for a partition lattice is classical; see the rank case of Rossi [24, Section 5, Proposition 11]. Here is the argument in the present coordinates. For $r \geq 1$, let P, Q be partitions of an r-element set, let $b(P)$ count blocks and set

$$\psi_Q(P) = b(P) + b(Q) - 2b(P \vee Q).$$

Merging two blocks of P decreases $b(P)$ by one. It either leaves $b(P \vee Q)$ unchanged or decreases it by one. To see this, form the bipartite block-intersection graph of P and Q: its vertices are the blocks of the two partitions, and two blocks from opposite sides are adjacent when they intersect. Each connected component gives one block

of their least common coarsening. The merge identifies two vertices in the same or different components. Thus each cover or reverse cover changes ψ_Q by exactly one in absolute value. Since $\psi_Q(Q) = 0$, every path from P to Q has length at least $\psi_Q(P)$. The path through $P \vee Q$ has exactly that length.

In the product (4.1), one move changes exactly one coordinate. Distances therefore add. Substituting $p(D) = c + \sum_p b_p(D) - |S(F)|$ gives the displayed formula. For the diameter bound, the bipartite block-intersection multigraph of P, Q, with one edge per underlying element, has r edges, $b(P) + b(Q)$ vertices and $b(P \vee Q)$ components. Hence

$$b(P) + b(Q) - 2b(P \vee Q) \leq r - b(P \vee Q) \leq r - 1.$$

The discrete and one-block partitions attain $r-1$. Consequently the product diameter is $\sum_{p \in S(F)} (\mu(p) - 1) = k - c$, by the atom–point forest identity. The empty product has diameter zero, as does the formula when $k = c$. □

Example 4.6 (Three triples through one point). Let the hyperedges be

$$e_i = \{o, a_i, b_i\}, \qquad 1 \leq i \leq 3,$$

with all six noncentral points distinct. The Levi graph is a tree, and the three single triples are the canonical atoms. There is one shared point, o, and its incident-atom star has size three. Thus $\mathcal{D}(F) \cong \Pi_3$. The five supported decompositions are

partition of the hyperedges	$p(D)$	$\delta(D)$
$e_1 \mid e_2 \mid e_3$	3	0
$e_1e_2 \mid e_3$, $e_1e_3 \mid e_2$, $e_2e_3 \mid e_1$	2	1
$e_1e_2e_3$	1	2

A local merger at o changes the global count by exactly one. The example uses a specified actual system; no realization theorem for an arbitrary prescribed atom list is needed.

For a fixed system, the local partitions therefore determine the finest completion of merger requests and the minimum number of mergers needed. The parameter spectra can rule out impossible numerical data before a construction is attempted. We make no running-time claim for enumeration or optimization.

Scope of formal verification

The classification, canonical reconstruction, finite spectra and supported decomposition interfaces used here belong to the preserved kernel-checked Lean 4 development [19, 20]. Appendix B lists the relevant declarations and preserved validation records. Every use retains the exact assumptions on finiteness, isolated points and connectedness, as well as any requirement that the vertex and hyperedge sets remain fixed. The small examples are elementary consequences of those statements; they are not presented as separate Lean results. The merger completion in Proposition 4.4 and the distance and diameter in Corollary 4.5 are proved in the text and are not claimed as separately formalised.

Formal verification checks a precise statement. Matching that statement to the prose remains a mathematical responsibility [21]. As explained in the appendix, parts

of the Lean classification proof use different intermediate arguments. This paper neither needs nor claims verification of density inequalities, realization of arbitrary prescribed atom profiles, results on lattice height or general counting-coefficient bounds.

Appendix A. A proof of the classification

This appendix completes the proof of Theorem A by giving the embedding arguments for hosts of uncountable chromatic number. The equivalence of the two finite descriptions was proved in Section 2. We first show that bipartite expansions are obligatory and that the allowed operations preserve obligatoriness. We then construct a host avoiding each finite system that fails the intrinsic criterion. All embeddings are injective and non-induced, and all arguments are in ZFC.

A.1. **Finite reductions and colouring tools.**

Lemma A.1 (Isolated-vertex reduction)**.** *Let F be a finite triple system and let F° be obtained by deleting all isolated vertices. Then*

$$F \text{ is obligatory} \iff F^\circ \text{ is obligatory},$$

and

$$F \in \mathcal{B} \iff F^\circ \in \mathcal{B}.$$

Proof. Suppose F° is obligatory and H has uncountable chromatic number. Then H has infinitely many vertices. Embed the finite system F°, then send the isolated vertices of F to distinct unused host vertices. Because containment is non-induced, additional host edges do not affect the embedding. The converse holds because F° is a subhypergraph of F.

For membership in $\mathcal{B}$, the implication $F^\circ \in \mathcal{B} \Rightarrow F \in \mathcal{B}$ follows by taking the disjoint union with a finite edgeless system.

For the converse, induct on a construction of F in $\mathcal{B}$. For a generator $F = J^+$, delete the isolated vertices of J, leaving J_0; then $(J^+)^\circ = J_0^+ \in \mathcal{B}$. If F is edgeless, its reduction is the empty edgeless system. Disjoint unions cause no further difficulty, because reduction commutes with them.

Suppose finally that F is a one-point amalgamation of F_0 and F_1, identifying $x_i \in V(F_i)$. By induction, $F_0^\circ, F_1^\circ \in \mathcal{B}$. If both selected vertices are nonisolated in their factors, then F° is the one-point amalgamation of the two reductions at those vertices. If exactly one selected vertex is nonisolated, the isolated side disappears at the identified point and F° is the disjoint union of the two reductions. The same description applies when both selected vertices are isolated, because the identified point is then deleted. Hence $F^\circ \in \mathcal{B}$ in every case. □

We may therefore suppress isolated vertices in the substantive argument and restore them at the end.

Lemma A.2 (Finite deletion)**.** *If $\chi(H) > \aleph_0$ and $S \subseteq V(H)$ is finite, then*

$$\chi(H - S) > \aleph_0.$$

Proof. Otherwise a countable colouring of $H - S$, together with finitely many fresh colours for S, would colour H countably. □

Colouring and closure tools.

Lemma A.3 (Two elementary colouring facts)**.** *We use the following two facts repeatedly.*

1. *If $E(H) = E_1 \cup \cdots \cup E_r$ and every $(V(H), E_i)$ is countably chromatic, then H is countably chromatic.*
2. *If a graph G admits an orientation in which every vertex has outdegree at most $d < \omega$, then $\chi(G) \leq 2d + 1$.*

Proof. For the first assertion, take the product of the r countable colourings.

For the second, every finite subgraph has at most $d|V|$ edges and therefore a vertex of degree at most $2d$. It is $(2d + 1)$-colourable by greedy deletion. The de Bruijn–Erdős compactness theorem [11, Theorem 1] then gives a $(2d + 1)$-colouring of all of G. □

We will use chains of sets closed under an operation that takes finitely many inputs and returns finitely many outputs. A chain is *continuous* when its set at each limit stage is the union of the earlier sets.

Lemma A.4 (Closure-chain lemma)**.** *Let κ be an uncountable cardinal, let $1 \leq r < \omega$, and let*

$$\Phi : [\kappa]^r \longrightarrow [\kappa]^{<\omega}$$

have uniformly finite values. There is an increasing continuous chain

$$\varnothing = M_0 \subseteq M_1 \subseteq \cdots \subseteq M_i \subseteq \cdots \qquad (i < \operatorname{cf} \kappa)$$

such that

$$\bigcup_{i<\operatorname{cf}\kappa} M_i = \kappa, \qquad |M_i| < \kappa,$$

and every M_i is Φ-closed:

$$a \in [M_i]^r \quad \Longrightarrow \quad \Phi(a) \subseteq M_i.$$

Proof. Choose sets $A_i \subseteq \kappa$, each of cardinality below κ, that together cover κ. At a successor stage, close $M_i \cup A_i$ under Φ by iterating the operation ω times. The operation Φ uses only finitely many inputs and returns a finite set, so this process does not increase an infinite cardinal below κ. At a limit stage, take the union. Such a union is still closed because every finite set of inputs occurs at an earlier stage. Because the stage is below $\operatorname{cf} \kappa$, the union still has cardinality below κ. □

The new vertices added at each stage form a layer. Write

$$I_i = M_{i+1} \setminus M_i.$$

Then $(I_i)_{i<\operatorname{cf}\kappa}$ partitions κ, and $M_i = \bigcup_{j<i} I_j$.

A.2. **A bipartite subgraph lemma.** The next lemma guarantees any prescribed finite complete bipartite graph. It follows from Erdős and Hajnal [7, Corollary 5.6, p. 72]. We give a closure-chain proof because we will use the same construction for triple systems.

Lemma A.5 (Uncountable chromatic number forces complete bipartite graphs)**.** *For every positive integer n, every graph of uncountable chromatic number contains $K_{n,n}$.*

Proof. Suppose the conclusion fails. Choose a $K_{n,n}$-free graph G of uncountable chromatic number with the smallest possible vertex cardinality κ.

For $A \in [\kappa]^n$, let

$$N(A) = \{v \in \kappa : v \text{ is adjacent to every member of } A\}.$$

Because G is $K_{n,n}$-free, $|N(A)| < n$. Apply Lemma A.4 to $A \mapsto N(A)$, and let M_i, I_i be the resulting chain and layers.

By minimality of κ, every induced graph $G[I_i]$ is countably chromatic. Let $G^{\times}$ consist of the edges joining distinct layers, and orient each such edge from the later endpoint toward the earlier endpoint. If $v \in I_i$ had n earlier neighbours, those neighbours would form an n-set $A \subseteq M_i$, and closure would give

$$v \in N(A) \subseteq M_i,$$

contrary to $v \in I_i$. Thus the orientation of $G^{\times}$ has outdegree at most $n-1$. By Lemma A.3, $G^{\times}$ is finitely colourable.

Colour each $G[I_i]$ with colours in ω, reusing the same palette on all layers, and combine this with a finite proper colouring of $G^{\times}$. The resulting countable product colouring is proper on all of G, a contradiction. □

A.3. **Bipartite expansions.**
Rainbow bipartite submatrices.

Lemma A.6 (Rainbow bipartite submatrices)**.** *For all positive integers n, t, there exists $q = q(n,t)$ with the following property. Suppose the edges of $K_{q,q}$ are coloured, with an arbitrary colour set, so that at every vertex each individual colour occurs on fewer than t incident edges. Then there are n-element subsets X' and Y' of the two vertex classes such that all n^2 edges between X' and Y' have distinct colours.*

Proof. Choose independent uniformly random n-subsets X' and Y'. For every unordered pair P of distinct same-coloured edges, let I_P be the indicator that both edges of P lie in the selected $K_{n,n}$, and put

$$Z = \sum_P I_P.$$

Thus $Z = 0$ exactly when all selected edges have distinct colours.

Let m_γ be the number of edges of colour γ. The local bound gives

$$m_\gamma \le (t-1)q.$$

At a fixed vertex and for a fixed colour γ, fewer than t incident edges have colour γ. In $\binom{d_{v,\gamma}}{2} \le (d_{v,\gamma})(t-1)/2$, the subscripted degree counts these incident edges. Summing this bound over the $2q$ vertices and all colours shows that the number of same-coloured pairs sharing a vertex is at most $(t-1)q^2$. The total number of same-coloured pairs is at most

$$\sum_{\gamma} \binom{m_\gamma}{2} \le \frac{1}{2}(\max_{\gamma} m_\gamma) \sum_{\gamma} m_\gamma \le \frac{1}{2}(t-1)q^3.$$

In the next two probabilities, a parenthesised integer with subscript two denotes its product with the preceding integer. If the two edges share a vertex in the left class, their exact selection probability is

$$\frac{n}{q}\frac{(n)_2}{(q)_2};$$

the same formula holds when the shared vertex is on the right. If the two edges are disjoint, the exact probability is

$$\left(\frac{(n)_2}{(q)_2}\right)^2.$$

For $q \ge 2n$, these are at most $2n^3/q^3$ and $4n^4/q^4$, respectively. By linearity of expectation,

$$\mathbb{E}[Z] \le \frac{2(t-1)n^3}{q} + \frac{2(t-1)n^4}{q}.$$

This is below 1 when q is sufficiently large. The variable Z is integer-valued and nonnegative. If $Z \ge 1$ for every choice, we could not have $\mathbb{E}[Z] < 1$. Thus some choice has $Z = 0$, so every edge of the selected $K_{n,n}$ has a different colour. □

The complete bipartite expansion atom.

Proposition A.7 (The complete bipartite expansion atom)**.** *For every positive integer n, the triple system $K^+_{n,n}$ is obligatory.*

Proof. Suppose not. Choose a $K^+_{n,n}$-free triple system

$$H = (\kappa, E)$$

with $\chi(H) > \aleph_0$ and with κ minimal.

For a pair xy, its codegree counts the vertices that complete it to a hyperedge:

$$d_H(xy) = |\{z : \{x, y, z\} \in E\}|.$$

Set $t = 3n^2 + 1$ and define a graph R on κ by

$$xy \in E(R) \quad \Longleftrightarrow \quad d_H(xy) \ge t.$$

Claim A.7.1 (the graph R is countably chromatic).

Otherwise Lemma A.5 gives a $K_{n,n}$ in R, with core vertex set C, $|C| = 2n$. For each of its n^2 graph edges xy, choose a third vertex p_{xy} with

$$\{x, y, p_{xy}\} \in E,$$

outside C and different from all previously chosen private vertices. At each step, fewer than $2n + n^2 < t$ vertices are forbidden, while xy has at least t possible third vertices. We can therefore choose the p_{xy} to be distinct and outside C. They give a copy of $K^+_{n,n}$, a contradiction. □

Let

$$E_1 = \{e \in E : \text{some pair contained in } e \text{ is an edge of } R\}, \qquad E_2 = E \setminus E_1.$$

A proper colouring of R is a proper colouring of (κ, E_1), since every E_1-edge contains an R-edge. Hence $\chi(\kappa, E_1) \leq \aleph_0$, and therefore

$$\chi(\kappa, E_2) > \aleph_0. \tag{A.1}$$

For a pair xy, define

$$\Phi(xy) = \begin{cases} \{z : \{x, y, z\} \in E\}, & d_H(xy) < t, \\ \varnothing, & d_H(xy) \geq t. \end{cases}$$

Apply Lemma A.4 to Φ, obtaining M_i, I_i.

Every $e \in E_2$ has at least two vertices in its highest layer. Indeed, suppose the highest layer met e only in $z \in I_i$. Then the other pair $xy = e \setminus \{z\}$ lies in M_i. Since $e \in E_2$, $xy \notin R$, so $d_H(xy) < t$. Closure gives $z \in \Phi(xy) \subseteq M_i$, a contradiction.

Thus every E_2-edge is either contained in one layer, or has exactly two vertices in its highest layer I_i and its third vertex in M_i. Let E^*_i be the E_2-edges contained in I_i. Since $|I_i| < \kappa$ and H is $K^+_{n,n}$-free, minimality of κ gives $\chi(I_i, E^*_i) \leq \aleph_0$. Consequently

$$E^* = \bigcup_i E^*_i$$

is countably chromatic. By (A.1), the remaining crossing-edge system

$$E^{**} = E_2 \setminus E^*$$

has uncountable chromatic number.

For each i, define a graph G_i on I_i by declaring $xy \in E(G_i)$ when there is a $\beta \in M_i$ such that $\{x, y, \beta\} \in E^{**}$. If every G_i were countably chromatic, colour each I_i properly in G_i, using the same countable palette on every layer. Every edge of E^{**} would then have differently coloured top-layer vertices, giving a countable colouring of (κ, E^{**}), a contradiction. Hence some G_i is uncountably chromatic.

Choose $q = q(n, t)$ from Lemma A.6. Lemma A.5 gives a $K_{q,q}$ in G_i, with sides X, Y. For every $xy \in X \times Y$, choose

$$\beta_{xy} \in M_i$$

such that $\{x, y, \beta_{xy}\} \in E^{**}$, and regard β_{xy} as the colour of the edge xy.

This edge-colouring is locally $(t - 1)$-bounded: at any fixed vertex, a colour appears on at most that many incident edges. Indeed, if fixed x and β admitted

t different vertices y with $\beta_{xy} = \beta$, then $\{x, \beta\}$ would have codegree at least t. It would be an edge of R, placing the triples $\{x, y, \beta\}$ in E_1, contrary to $E^{**} \subseteq E_2$. The same argument applies at vertices of Y.

Lemma A.6 gives n-sets $X' \subseteq X$, $Y' \subseteq Y$ for which the β_{xy}, $x \in X'$, $y \in Y'$, are all distinct. They lie in M_i, whereas $X' \cup Y' \subseteq I_i$, so none is a core vertex. The triples

$$\{x, y, \beta_{xy}\}, \qquad x \in X',\ y \in Y',$$

form $K^+_{n,n}$, the final contradiction. □

Corollary A.8 (Every finite bipartite expansion is obligatory)**.** *If J is any finite bipartite graph, then J^+ is obligatory.*

Proof. Embed the two vertex classes of J in those of a sufficiently large $K_{n,n}$. This embeds J^+ as a subhypergraph of $K^+_{n,n}$. A copy of the larger system contains a copy of the smaller one, so obligatoriness passes to subhypergraphs under non-induced containment. The result now follows from Proposition A.7. □

Corollary A.8 also follows from Reiher [10, Theorem 1.2]. The argument above gives a direct proof in the notation used here.

A.4. **Closure properties.** Obligatoriness is preserved by disjoint unions and one-point amalgamations; see Komjáth [9, pp. 233–238] and the summary in Reiher [10, p. 1]. We include the proofs. For a one-point amalgamation, we will find several copies meeting at a chosen root and select one that avoids the other factor away from that root.

Disjoint-union closure.

Lemma A.9 (Disjoint-union closure)**.** *The class of finite obligatory triple systems is closed under finite disjoint unions.*

Proof. Let F_1, F_2 be obligatory and let H be uncountably chromatic. First find a copy of F_1. Delete its finite vertex set. By Lemma A.2 the remaining triple system is still uncountably chromatic, so it contains F_2. The two copies are disjoint. Iteration proves the finite case. □

Rooted abundance and one-point amalgamation.

Definition A.10 (Rooted abundance)**.** Let (F, r) be a finite rooted triple system with $|V(F)| > 1$. For a triple system H and $m \geq 1$, let $R_m(F, r; H)$ be the set of vertices v for which there are m rooted copies of F, all sending r to v, whose off-root vertex sets are pairwise disjoint.

Lemma A.11 (Rooted-abundance lemma)**.** *If (F, r) is finite, $|V(F)| > 1$, and F is obligatory, and if $\chi(H) > \aleph_0$, then*

$$\chi(H[R_m(F, r; H)]) > \aleph_0.$$

Proof. Put

$$B = V(H) \setminus R_m(F, r; H).$$

For each $v \in B$, consider the off-root vertex sets of all copies of F rooted at v. Since $v \notin R_m(F, r; H)$, at most $m-1$ of these sets can be pairwise disjoint. Choose a maximal disjoint subfamily. It has at most $m-1$ members; let S_v be their union. Then

$$|S_v| \leq (m-1)(|V(F)|-1) =: D.$$

The set S_v meets the off-root part of every copy rooted at v. Otherwise we could add that off-root set to the chosen disjoint subfamily, contradicting maximality.

Form a graph D_B on B by joining each v to the vertices of $S_v \cap B$. Orient each edge away from a vertex whose chosen set generated it, choosing one endpoint if both did. Every vertex then has outdegree at most D. Lemma A.3 therefore gives a finite proper colouring of D_B.

Let A be one colour class. Then

$$S_v \cap A = \varnothing \qquad (v \in A).$$

Consequently $H[A]$ contains no copy of F. Indeed, if φ were such a copy and $v = \varphi(r)$, its off-root part would have to meet S_v. All those off-root vertices lie in A, contradicting $S_v \cap A = \varnothing$. Since F is obligatory, every F-free triple system is countably chromatic. Thus each of the finitely many colour classes of D_B induces a countably chromatic subsystem. Tagging a vertex by its finite D_B-colour and its countable colour inside that class gives a countable proper colouring of $H[B]$.

If $H[R_m(F, r; H)]$ were countably chromatic as well, tagging the two sides would give a countable colouring of H. This contradicts $\chi(H) > \aleph_0$. □

Proposition A.12 (One-point amalgamation closure)**.** *The class of finite obligatory triple systems is closed under one-point amalgamations.*

Proof. Let F_1, F_2 be obligatory, with selected vertices r_1, r_2, and let F be their one-point amalgamation. If one factor consists only of its selected vertex, the assertion is immediate. Otherwise set

$$m = |V(F_2)|.$$

By Lemma A.11, the set

$$R = R_m(F_1, r_1; H)$$

induces an uncountably chromatic subsystem of every uncountably chromatic H. Hence $H[R]$ contains a copy of F_2; write φ for its embedding and let $v = \varphi(r_2)$.

At v there are m rooted copies of F_1 with pairwise disjoint off-root sets. The set $\varphi(V(F_2) \setminus \{r_2\})$ has $m-1$ vertices, so it meets at most $m-1$ of those off-root sets. One rooted F_1-copy therefore avoids the F_2-copy outside v. Their union is the required one-point amalgamation. □

Corollary A.8, Lemma A.9, Proposition A.12, and the trivial edgeless case prove

$$(2) \Longrightarrow (1). \tag{A.2}$$

A.5. **The sequence lift.** We next construct an uncountably chromatic host from a graph. Each new triple has two vertices at one sequence node and a third vertex, its apex, at a later node. The construction uses sequences of length below ω_1. We will group its hyperedges by their first node and record how those groups overlap; this description matches the formalisation. Li [3, Sections 3–4 and Theorem 4.6] gives an alternative one-apex and bridge-trace formulation.
The one-apex construction. Let $G = (X, A)$ be a graph. Regard its edge set $A = E(G)$ as an alphabet, so that each word records a sequence of graph edges. Let

$$T = A^{<\omega_1}$$

be the tree of all sequences of graph edges of countable ordinal length. For $s \in T$ and $x \in X$, write $x_s = (s, x)$.

Definition A.13 (The one-apex lift)**.** Define the triple system $\mathcal{L}(G)$ by

$$V(\mathcal{L}(G)) = T \times X.$$

For $s \in T$, an edge $a = \{x, y\} \in A$, a sequence $t \in T$ extending $s^\frown a$, and $z \in X$, put

$$\{x_s, y_s, z_t\} \in E(\mathcal{L}(G)). \tag{A.3}$$

The two vertices x_s, y_s form the base pair and z_t is the apex. Since t properly extends s, these are three distinct vertices. In the resulting triple, the sequence node occurring twice uniquely identifies s; its two graph labels identify the edge $a = \{x, y\}$, and the remaining vertex identifies t, z. Thus each hyperedge arises from unique data, up to interchanging x and y, and the system is simple and 3-uniform.

Chromatic and finite-trace properties.

Lemma A.14 (Chromatic lower bound)**.** *If $\chi(G) > \aleph_0$, then*

$$\chi(\mathcal{L}(G)) > \aleph_0.$$

Proof. Suppose, for a contradiction, that $c : T \times X \to \omega$ is a proper colouring of the lift. At each $s \in T$, the map $x \mapsto c(x_s)$ is a countable colouring of G. Hence there is an edge

$$a_s = \{x_s^0, x_s^1\} \in E(G)$$

whose two vertices at node s have a common colour k_s.
Construct $s_\alpha \in A^\alpha$ for $\alpha < \omega_1$ by

$$s_0 = \varnothing, \qquad s_{\alpha+1} = s_\alpha^\frown a_{s_\alpha},$$

and at limit ordinals take unions.

If $\alpha < \beta$, then s_β extends $s_\alpha^\frown a_{s_\alpha}$. If $k_{s_\alpha} = k_{s_\beta}$, take either endpoint z of the monochromatic edge at s_β. By (A.3),

$$\{(s_\alpha, x^0_{s_\alpha}), (s_\alpha, x^1_{s_\alpha}), (s_\beta, z)\}$$

is a monochromatic hyperedge. Therefore the colours k_{s_α}, $\alpha < \omega_1$, are pairwise distinct, impossible in ω. □

We now describe the finite linear subhypergraphs of the lift. The proof groups hyperedges with the same base node, identifies each group as a graph expansion, and then controls the overlaps between groups. Li gives a parallel bridge-trace formulation [3, Theorem 4.6].

Theorem A.15 (Finite linear trace decomposition)**.** *Every finite, not necessarily induced, linear subhypergraph K of $\mathcal{L}(G)$ is obtainable, by finite disjoint unions and one-point amalgamations, from systems J^+ where J is a finite subgraph of G.*

Proof. We first name the pieces used in the proof. A *base node* is a sequence s at which a lift hyperedge begins. The two vertices with coordinate s are its *core vertices*; the third vertex, with coordinate properly extending s, is its *apex*. The *base fibre* K_s consists of the hyperedges of K based at s, together with their incident vertices. A node is *active* if its fibre is nonempty. We will show that every fibre is a graph expansion and that distinct fibres can be joined along single points without forming an overlap cycle.

First set aside every isolated vertex of K. Each is a one-vertex edgeless factor, hence isomorphic to J^+ for a one-vertex edgeless subgraph J of G, and may be restored by disjoint union at the end. We may therefore assume that every vertex of K lies in a hyperedge.

Every hyperedge of $\mathcal{L}(G)$ has a unique base node: exactly two vertices have sequence coordinate s, while the apex has a strictly longer coordinate. Grouping the hyperedges of K by their base node gives the fibres K_s just defined.

Claim A.15.1 (each base fibre is one expansion piece). For every active base node s, there is a finite subgraph J_s of G such that $K_s \cong J_s^+$.

Proof of Claim A.15.1. For a fixed graph edge $a = \{x, y\} \in E(G)$, at most one hyperedge of K_s can have base pair $\{x_s, y_s\}$, since two such hyperedges would share two vertices and violate linearity. Let J_s be the finite subgraph of G whose edges are precisely the graph edges used by hyperedges of K_s, with vertex set consisting of their endpoints.

Define a map $J_s^+ \to K_s$ by sending each graph vertex $x \in V(J_s)$ to the core vertex x_s, and by sending the private vertex of $a = \{x, y\}$ to the apex of the unique hyperedge of K_s with base pair $\{x_s, y_s\}$. Distinct graph edges have distinct apices: if z_t were the apex for a and a', the coordinate of t immediately after s would have to equal both a and a'. No apex is a core vertex, because an apex has sequence coordinate strictly extending s, whereas every core vertex has coordinate s. The map is therefore injective, and by construction it is bijective on hyperedges and preserves each incidence. Hence $K_s \cong J_s^+$. □

Claim A.15.2 (two fibres have at most one common point). For distinct active base nodes s and u,

$$|V(K_s) \cap V(K_u)| \leq 1.$$

Proof of Claim A.15.2. If a point z_t lies in both K_s and K_u, then both s and u are prefixes of t, so they are comparable. Suppose $s \subsetneq u$. A common point cannot be a core point of K_s; it is an apex of K_s. The graph edge used by its hyperedge in K_s is the coordinate of u immediately following s. Hence all points in $K_s \cap K_u$ would be apices of the same base edge of K_s. There is at most one such apex. Therefore

$$|V(K_s) \cap V(K_u)| \leq 1. \tag{A.4}$$

□

Let Q be the bipartite incidence graph of the base fibres. Its left class is the set of active base nodes. Its right class consists of points that lie in at least two fibres. Join a base node s to a point p exactly when $p \in V(K_s)$.

Claim A.15.3 (the overlap graph is a forest). The graph Q is acyclic.

Proof of Claim A.15.3. Suppose that Q has a cycle, and choose on it a base node s of minimum sequence length. Its two neighbouring base nodes on the cycle are proper descendants of s. The two intervening shared points are distinct, so they are distinct apices in K_s, arising from two distinct graph edges $a \neq a'$. Thus the two neighbouring base nodes lie in different immediate branches below s, one beginning with a and the other with a'.

Starting from one neighbour of s, every base node on the remaining part of the cycle lies in the same immediate branch below s. Indeed, consecutive base nodes share a point and are therefore comparable. For two comparable proper descendants of s, the first sequence coordinate after s is the same, so a length-two step through a shared point cannot move from one immediate branch to another. Induction along the remaining path therefore keeps every base node in the first branch, contradicting that its final node lies in the other branch. Hence Q is acyclic. □

Root each component of Q at a base node, and order its base nodes by nondecreasing even distance from the root. For a nonroot base node s, let p_s be the point-node adjacent to s on the path to the root.

Claim A.15.4 (Base-fibre running intersection)**.** *Every earlier factor that meets K_s meets it at p_s.*

Proof of Claim A.15.4. Suppose K_u is earlier and contains a point $p \in K_s$. Then Q contains the length-two path $s - p - u$. Let the distance of s from the root be $2d$. The point-node p has distance either $2d - 1$ or $2d + 1$. In the second case u has distance $2d + 2$, so it cannot be earlier. Thus p has distance $2d - 1$, and the uniqueness of the path to the root in the tree Q gives $p = p_s$. This also covers an earlier sibling u, which shares the same parent point p_s. □

Claim A.15.5 (the ordered fibres assemble to K). Let U_s be the union of the factors preceding K_s. Equation (A.4) and Claim A.15.4 give

$$V(K_s) \cap U_s \subseteq \{p_s\}.$$

Starting with the root factor, add the remaining factors in this order by disjoint unions or one-point amalgamations as appropriate. At each step the natural maps into K identify exactly the displayed common point and nothing else, so they assemble to an isomorphism onto the union of the factors already added. Different components of Q, and the isolated factors removed at the start, are combined by disjoint union. This is the running-intersection step: Claim A.15.1 supplies the expansion pieces, while Claims A.15.2–A.15.4 certify that each new piece is attached along at most one previous point. □

Corollary A.16 (Restrictions on finite linear traces)**.** *Every finite linear $K \subseteq \mathcal{L}(G)$ satisfies:*

1. *every hyperedge-node of $I(K)$ is incident with a bridge;*
2. *every Berge cycle of K is contained in a single factor J_s^+;*
3. *a Berge cycle of length ℓ in K yields an ordinary cycle of length ℓ in G.*

Proof. In every J_s^+, each incidence with a private vertex is a bridge, and these incidences remain bridges after one-point amalgamation. A simple cycle in the Levi graph cannot pass through a one-point cut and return without repeating that point. Thus every Berge cycle stays in one factor. The Berge cycles in J_s^+ are exactly the ordinary graph cycles in J_s. □

A.6. **A classical high-odd-girth input.** The following classical graph theorem supplies the odd-cycle obstruction. We apply it directly, rather than replace it by a shift-graph construction.

Theorem A.17 (Erdős–Hajnal)**.** *For every positive integer m there is a graph G such that*

$$\chi(G) > \aleph_0$$

and G contains no odd cycle of length at most m.

Source. Theorem 7.4 of Erdős and Hajnal [7, p. 76], as quoted in Erdős et al. [8, Theorem C, p. 428], gives for every positive integer i an uncountably chromatic graph containing no cycle C_{2j+1} for $0 < j < i$. Choose i so that $2i - 1 \geq m$. □

A.7. **Avoidance hosts.** The intrinsic criterion can fail in three ways: nonlinearity, a missing bridge, or an odd Berge cycle. Li [3, Sections 2–4] organises the negative direction in a parallel way. We prove the first two avoidance constructions in our notation and apply the Erdős–Hajnal theorem from Appendix A.6 to the third.

We prove

$$\neg(3) \Longrightarrow \neg(1).$$

Nonlinearity obstruction.

Proposition A.18 (Avoidance of every nonlinear finite triple system)**.** *There is an uncountably chromatic linear triple system. Consequently, every finite nonlinear triple system is non-obligatory.*

Proof. We use the Erdős–Rado relation [12, Theorem 4(i), formula (95), p. 471]

$$(2^{\aleph_0})^+ \longrightarrow (\aleph_1)^2_{\aleph_0}. \tag{A.5}$$

Let $\kappa = (2^{\aleph_0})^+$. Define a triple system T_κ whose vertices are the pairs in $[\kappa]^2$, and whose hyperedges are

$$\{\{\alpha,\beta\},\{\alpha,\gamma\},\{\beta,\gamma\}\}, \qquad \alpha<\beta<\gamma<\kappa.$$

This triple system is linear. Two graph edges belonging to a triangle determine its three vertices and its third edge, so distinct triangles of a complete graph cannot share two edges.

It is uncountably chromatic. A countable vertex-colouring of T_κ is a countable edge-colouring of K_κ. By (A.5) there is a monochromatic subset of cardinality $\aleph_1$, and in particular a monochromatic graph triangle. This is a monochromatic hyperedge of T_κ.

If a finite triple system F is nonlinear, it has two distinct edges sharing at least two vertices. No injective embedding of those edges can exist in the linear host T_κ. Thus T_κ is F-free. □

Missing-bridge obstruction.

Proposition A.19 (Avoidance when the bridge condition fails)**.** *Let F be finite and linear, and suppose some hyperedge-node of $I(F^\circ)$ is incident with no bridge. Then F is non-obligatory.*

Proof. Use the sequence lift of the complete graph K_{ω_1}. Li gives an alternative missing-bridge obstruction [3, Section 4]. By Lemma A.1 it suffices to avoid F°, so assume that F has no isolated vertices. Take $G = K_{\omega_1}$. Then $\chi(G) > \aleph_0$, and Lemma A.14 gives

$$\chi(\mathcal{L}(G)) > \aleph_0.$$

If F appeared in $\mathcal{L}(G)$, retain exactly the host hyperedges corresponding to the edges of F. Because the embedding is injective and F is linear, these selected host hyperedges form a finite linear subhypergraph isomorphic to F. Corollary A.16 says that every hyperedge-node of such a trace is incident with a bridge, contradicting the chosen hyperedge of F. Therefore $\mathcal{L}(K_{\omega_1})$ is an uncountably chromatic F-free host. □

Odd-cycle obstruction.

Proposition A.20 (Avoidance of an odd Berge cycle)**.** *Let F be finite and linear, and suppose F° has an odd Berge cycle. Then F is non-obligatory.*

Proof. Suppress isolated vertices and put $m = |E(F)|$. By Theorem A.17 choose a graph G with $\chi(G) > \aleph_0$ and no odd cycle of length at most m. Lemma A.14 gives $\chi(\mathcal{L}(G)) > \aleph_0$.

Suppose that F embeds in $\mathcal{L}(G)$ and retain exactly the host hyperedges corresponding to the edges of F. The selected trace is finite and linear. Corollary A.16 localises every Berge cycle in one expansion fibre and sends a Berge cycle of length ℓ to an ordinary cycle of the same length in G. The given odd Berge cycle has $\ell \leq m$, contradicting the choice of G. Thus $\mathcal{L}(G)$ is an uncountably chromatic F-free host. □

The three propositions cover every failure of the intrinsic conditions and therefore prove

$$\neg(3) \Longrightarrow \neg(1). \tag{A.6}$$

A.8. **Proof of the classification.** The expansion and closure arguments establish (2) $\Longrightarrow$ (1), Section 2 establishes (2) $\Longleftrightarrow$ (3), and the avoidance hosts prove the contrapositive $\neg(3) \Longrightarrow \neg(1)$. Together these implications prove all equivalences in Theorem A. In particular,

$$F \text{ is obligatory} \iff F \in \mathcal{B},$$

$$F \in \mathcal{B} \iff \begin{array}{l} F^\circ \text{ is linear,} \\ \text{every hyperedge-node of } I(F^\circ) \text{ is incident with a bridge,} \\ \text{every Berge cycle of } F^\circ \text{ has even length.} \end{array}$$

All embeddings are non-induced, so additional host hyperedges on the image vertices are irrelevant. The proof uses only ZFC, graph-colouring compactness [11], the classical Erdős–Hajnal high-odd-girth theorem [7], and the standard Erdős–Rado relation (A.5) [12]. It assumes neither CH nor GCH and uses no forcing axiom or large-cardinal hypothesis.

Appendix B. Lean interfaces and reproducibility

The public source is in the `Erdos593` repository, which includes the self-contained `Erdos593SelfContained.lean` file. The selected checkpoint uses Lean 4.32.0 and pinned Mathlib. The preserved validation record reports successful compilation of 250 internal modules and checks of 401 declarations in the axiom audit. Their reported dependencies are subsets of `propext`, `Classical.choice` and `Quot.sound`. The standalone file was regenerated and checked in a separate replay. Later source comparisons and the repository's read-only snapshot checker identify the precise files covered by those records. Checking their hashes does not compile the proofs again.

The list below omits the namespace prefix `Erdos593.TripleSystem`.

- Theorem A: `isObligatory_iff_isolatedReduction_intrinsic` and `isObligatory_iff_constructible_isolatedReduction`. The target system is finite. Host systems retain the documented ambient-universe quantification of the formalization.

- The canonical atom normal form: `CanonicalAtom.canonicalAtom_reconstruction`, `exists_canonicityTransport` and the associated statements on generation by atoms and on minimal generators. Reconstruction requires the intrinsic conditions and no isolated points. The canonicity result describes how the original edge sets correspond under an incidence isomorphism.
- The finite spectra: `obligatory_spectrum_necessity`, `exists_obligatory_spectrum`, `exists_obligatory_fixed_order_iff`, `exists_obligatory_cycleRank_iff` and `CanonicalAtom.exists_connected_atom_count_iff`, together with the accepted maximum-atom-count interfaces. The zero-rank and rank-one cases retain their separate conditions.
- The supported product: `CanonicalAtom.obligatory_supported_decomposition_product`, its standard-product and lattice-bound interfaces, and `supportedDecomposition_iff_manuscript`. Each product factor uses the actual set of canonical atoms incident with one shared point.
- Theorem 4.2: `supportedPieceDeficit_eq_sum` and `supportedPieces_add_deficit`, backed by the original-edge piece-accounting identities.
- Theorem 4.3: `supportedPartitions_covBy_iff_exists_unique`, `supportedPartitions_block_card_of_covBy`, `supportedPartitions_covBy_iff_block_card` and `supportedPartitions_covBy_iff_piece_deficit`. These results describe covers among supported decompositions. Their converses still require strict comparability; they do not assert a cover relation in the full ambient partition lattice.

The Lean proof follows a different route from Appendix A: its negative direction uses a shift-graph host rather than the displayed high-odd-girth input, and its closure proof uses a layering that also handles singular cardinals. Both routes prove the same classification statements. The original classification and the finite structural extensions are separate from Li's Lean development [3]; no shared code or common private development history is asserted.

The repository preserves the exact theorem types, source manifests, compiler logs and axiom reports. A table accompanying this author-review source matches the manuscript statements to their Lean declarations. Kernel acceptance alone does not establish novelty or show that a formal statement matches its prose formulation. No additional run of an external kernel checker is claimed [22].

Acknowledgments

The author thanks Tom de Groot for advice on improving the exposition. The author received no funding for this work and declares no competing interests.

AI assistance

OpenAI's GPT-5.6 Pro, GPT-6 Pro and Codex, and Aristotle [23], were used in proof development, checking, editorial revision and the Lean formalization. The author takes responsibility for the arguments, attribution, citations and final text.

References

[1] Paul Erdős. On some problems in combinatorial set theory. *Publikacije Instituta Matematičkog (Beograd) (N.S.)*, 57(71):61–65, 1995.

[2] Thomas F. Bloom. Erdős problem 593. `https://www.erdosproblems.com/593`, 2026. Accessed 11 July 2026.

[3] Eric Li. A resolution of Erdős Problems 593 and 1177: Obligatory triple systems and exact spectra. arXiv:2606.24882, submitted 23 June 2026. doi: 10.48550/arXiv.2606.24882. URL `https://arxiv.org/abs/2606.24882`.

[4] Hassler Whitney. Non-separable and planar graphs. *Trans. Amer. Math. Soc.*, 34(2):339–362, 1932. doi: 10.1090/S0002-9947-1932-1501641-7.

[5] Reinhard Diestel. *Graph Theory*, fifth edition. Graduate Texts in Mathematics, volume 173. Springer, Berlin, 2017. doi: 10.1007/978-3-662-53622-3.

[6] Richard P. Stanley. *Enumerative Combinatorics*, volume 1, second edition. Cambridge Studies in Advanced Mathematics, volume 49. Cambridge University Press, Cambridge, 2012. doi: 10.1017/CBO9781139058520.

[7] Paul Erdős and András Hajnal. On chromatic number of graphs and set-systems. *Acta Math. Acad. Sci. Hungar.*, 17(1–2):61–99, 1966. doi: 10.1007/BF02020444.

[8] Paul Erdős, Fred Galvin, and András Hajnal. On set-systems having large chromatic number and not containing prescribed subsystems. In *Infinite and Finite Sets (Keszthely, 1973)*, volume 10 of *Colloquia Mathematica Societatis János Bolyai*, pages 425–513. North-Holland, 1975.

[9] Péter Komjáth. Some remarks on obligatory subsystems of uncountably chromatic triple systems. *Combinatorica*, 21(2):233–238, 2001. doi: 10.1007/s004930100021.

[10] Christian Reiher. Obligatory hypergraphs. arXiv:2403.11223, 2024. doi: 10.1090/proc/17021. URL `https://arxiv.org/abs/2403.11223`.

[11] Nicolaas Govert de Bruijn and Paul Erdős. A colour problem for infinite graphs and a problem in the theory of relations. *Nederl. Akad. Wetensch. Proc. Ser. A*, 54(5):371–373, 1951. doi: 10.1016/S1385-7258(51)50053-7.

[12] Paul Erdős and Richard Rado. A partition calculus in set theory. *Bull. Amer. Math. Soc.*, 62(5):427–489, 1956. doi: 10.1090/S0002-9904-1956-10036-0.

[13] M. Amin Bahmanian and Mateja Šajna. Connection and separation in hypergraphs. *Theory Appl. Graphs*, 2(2):Article 5, 2015. doi: 10.20429/tag.2015.020205. URL `https://arxiv.org/abs/1504.04274`.

[14] R. Endre Tarjan. A note on finding the bridges of a graph. *Inform. Process. Lett.*, 2(6):160–161, 1974. doi: 10.1016/0020-0190(74)90003-9.

[15] Robert Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–160, 1972. doi: 10.1137/0201010.

[16] Frank Simon, Peter Tittmann, and Martin Trinks. Counting connected set partitions of graphs. *Electron. J. Combin.*, 18(1):P14, 2011. doi: 10.37236/501. URL `https://arxiv.org/abs/1005.1726`.

[17] Richard P. Stanley. An introduction to hyperplane arrangements. In Ezra Miller, Victor Reiner, and Bernd Sturmfels, editors, *Geometric Combinatorics*, volume 13 of *IAS/Park City Mathematics Series*, pages 389–496. American Mathematical Society, Providence, RI, 2007. URL `https://math.mit.edu/~rstan/arrangements/arr.html`.

[18] George Grätzer. *Lattice Theory: Foundation*. Birkhäuser, Basel, 2011. doi: 10.1007/978-3-0348-0018-1.

[19] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In André Platzer and Geoff Sutcliffe, editors, *Automated Deduction—CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, Cham, 2021. doi: 10.1007/978-3-030-79876-5_37.

[20] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 367–381. Association for Computing Machinery, New York, 2020. doi: 10.1145/3372885.3373824.

[21] Jeremy Avigad and John Harrison. Formally verified mathematics. *Commun. ACM*, 57(4):66–75, 2014. doi: 10.1145/2591012.

[22] Mario Carneiro. Lean4Lean: Verifying a typechecker for Lean, in Lean. arXiv:2403.14064v3, revised 14 September 2025. URL `https://arxiv.org/abs/2403.14064v3`.

[23] Tudor Achim, Alex Best, Kevin Der, Mathïs Fédérico, Sergei Gukov, Daniel Halpern-Leistner, Kirsten Henningsgard, Yury Kudryashov, Alexander Meiburg, Martin Michelsen, Riley Patterson, Eric Rodriguez, Laura Scharff, Vikram Shanker, Vladimir Sicca, Hari Sowrirajan, Aidan Swope, Matyas Tamas, Vlad Tenev, Jonathan Thomm, Harold Williams, and Lawrence Wu. Aristotle: IMO-level automated theorem proving, 2025. URL `https://arxiv.org/abs/2510.01346`. arXiv:2510.01346.

[24] G. Rossi. Weighted paths between partitions. *arXiv:1509.01852v1*, 2015. `https://arxiv.org/abs/1509.01852v1`.

École normale supérieure–PSL, 45 rue d'Ulm, 75005 Paris, France
Email address: `samuil.petkov@ens.psl.eu`