

THE EVEN CASES OF THE FOUR-UNIFORM BROWN–ERDŐS–SÓS ($2k + 2, k$)-PROBLEM

ABSTRACT. For integers r, n, s, k , let $f^{(r)}(n; s, k)$ denote the largest number of edges in an n -vertex r -uniform hypergraph in which every k distinct edges span more than s vertices. We prove that, for every fixed even integer $k \geq 2$,

$$\lim_{n \rightarrow \infty} \frac{f^{(4)}(n; 2k + 2, k)}{n^2} = \frac{1}{12}.$$

For even $k \geq 10$, the upper bound follows from an exact finite inequality for a strengthened forbidden family. Its proof decomposes the hypergraph according to pair-sharing connectivity, assigns two candidate vertex pairs to each link of suitable spanning trees, and verifies Hall's condition by an overlap-sensitive k -edge contradiction. Published packing results give the lower bound and the cases $k = 2, 4, 6, 8$.

1. INTRODUCTION AND REDUCTION

Throughout the finite argument, G denotes a fixed finite simple four-uniform hypergraph. We use the following notation throughout the paper.

Definition 1.1 (Basic hypergraph notation). An r -graph is a finite simple r -uniform hypergraph: every edge is an r -element vertex set and no edge is repeated. When the ambient vertex set is clear, we identify a hypergraph with its edge set. For an edge family A , its *rank* and *support* are

$$\text{rk}(A) := |A|, \quad V(A) := \bigcup_{a \in A} a.$$

Thus A spans $|V(A)|$ vertices. For edge families A, B , write $A + B := A \cup B$; for a single edge e , write $A + e := A \cup \{e\}$. All ranks of unions count distinct hyperedges.

For disjoint vertex sets X, Y , the notation

$$X \times Y := \{\{x, y\} : x \in X, y \in Y\}$$

denotes a set of unordered vertex pairs. A vertex pair p is *supported* by A when $p \in \binom{V(A)}{2}$.

Definition 1.2 (Extremal functions and forbidden families). For integers r, n, s, k , $f^{(r)}(n; s, k)$ is the maximum number of edges in an n -vertex r -graph containing no k distinct edges whose union has at most s vertices. This extremal question is called the r -uniform (s, k) -problem.

If $\mathcal{A}$ is a family of r -graphs, an r -graph is $\mathcal{A}$ -free if it contains no member of $\mathcal{A}$ as a subgraph, and $\text{ex}_r(n, \mathcal{A})$ is the largest number of edges in an n -vertex $\mathcal{A}$ -free r -graph. Finally, $\mathcal{H}_r(v, h)$ denotes the family of all h -edge r -graphs supported on at most v vertices.

Theorem 1.3 (Main theorem). *For every fixed even integer $k \geq 2$,*

$$\lim_{n \rightarrow \infty} \frac{f^{(4)}(n; 2k + 2, k)}{n^2} = \frac{1}{12}.$$

The finite argument concerns even $k \geq 10$. For such k , let

$$\mathcal{G}_k := \mathcal{H}_4(2k + 2, k) \cup \bigcup_{2 \leq h < k} \mathcal{H}_4(2h + 1, h).$$

The union here is a union of forbidden isomorphism types; its two displayed parts have different edge counts and hence are disjoint.

2020 *Mathematics Subject Classification.* 05C65, 05D05.

Key words and phrases. Brown–Erdős–Sós problem, extremal hypergraphs, pair packing, Hall's theorem.

Lemma 1.4 (The two support walls). *If G is $\mathcal{G}_k$ -free, then every nonempty $F \subseteq G$ with $|F| < k$ satisfies*

$$|V(F)| \geq 2|F| + 2, \quad 0 < |F| < k, \quad (1.1)$$

and every k -edge subfamily of G spans at least $2k + 3$ vertices. We call (1.1) the rank-below- k support wall and the second assertion the k -edge support wall.

Proof. The first assertion excludes every member of $\mathcal{H}_4(2h+1, h)$ for $2 \leq h < k$; for $h = 1$, a single edge has support $4 = 2h + 2$. The second assertion is precisely the exclusion of $\mathcal{H}_4(2k+2, k)$. $\square$

The finite statement from which the upper bound will follow is the next theorem.

Theorem 1.5 (Finite pair-capacity theorem). *Let $k \geq 10$ be even. Every n -vertex $\mathcal{G}_k$ -free 4-graph G satisfies*

$$6|G| \leq \binom{n}{2}.$$

For arbitrary integers $r \geq 3$ and $k \geq 2$, define

$$\mathcal{F}_k^{(r)} := \mathcal{H}_r((r-2)k+2, k) \cup \bigcup_{2 \leq h < k} \mathcal{H}_r((r-2)h+1, h). \quad (1.2)$$

In particular, $\mathcal{F}_k^{(4)} = \mathcal{G}_k$.

Proposition 1.6 (Published reduction). *For fixed $r \geq 3$ and $k \geq 3$,*

$$\limsup_{n \rightarrow \infty} n^{-2} f^{(r)}(n; (r-2)k+2, k) \leq \limsup_{n \rightarrow \infty} n^{-2} \text{ex}_r(n, \mathcal{F}_k^{(r)}). \quad (1.3)$$

Proof. This is Lemma 3.2 of Wang and Zeng [7]. The underlying reduction was proved by Delcourt–Postle for $r = 3$ [1] and by Shangquan for $r \geq 4$ [2]. $\square$

Definition 1.7 (Tight edge family). An edge family $A \subseteq G$ is *tight* if

$$|V(A)| = 2|A| + 2. \quad (1.4)$$

A *state* is a tight edge family singled out during one of the constructions below. The word “state” carries no additional hypothesis.

The proof of Theorem 1.5 occupies Sections 2–8. Section 9 then combines it with Proposition 1.6, a published lower bound, and the known small even cases.

2. ORDINARY COMPONENTS AND MAXIMAL TIGHT BLOCKS

Definition 2.1 (Ordinary adjacency and shadow). Two edges of G are *ordinarily adjacent* if they contain a common vertex pair. The *ordinary edge-intersection graph* has the edges of G as its vertices and joins ordinarily adjacent hyperedges. Its connected components, viewed again as edge families, are the *ordinary components* of G . An ordinary component is *positive* if it has at least two edges; a one-edge component is *singleton*.

An *ordinary path* is a sequence of hyperedges in which consecutive edges are ordinarily adjacent. Ordinary-path distance is graph distance in the ordinary edge-intersection graph. A *connected ordering* of a connected edge family A is an ordering $a_1, \dots, a_t$ such that, for every $i > 1$, the edge a_i is ordinarily adjacent to an earlier edge. When

$$a_i \cap V(\{a_1, \dots, a_{i-1}\})$$

has exactly two vertices, that pair is the *attachment pair* of a_i in the ordering. More generally, if an edge $e \notin A$ meets $V(A)$ exactly in a pair x , then x is the attachment pair of e to A .

The *ordinary shadow*, or *two-shadow*, of an edge family A is

$$\partial_2 A := \bigcup_{a \in A} \binom{a}{2}. \quad (2.1)$$

A pair is *shadowed in A* if it belongs to $\partial_2 A$, and is *unshadowed in A* otherwise. Relative to the fixed ambient graph G , it is *globally shadowed* if it belongs to $\partial_2 G$, and *globally nonshadow* otherwise. An edge containing a pair p is said to *shadow p*.

Lemma 2.2 (Structure of ordinary components). *Let O be an ordinary component and put $e := |O|$. Then:*

- (i) $e < k$;
- (ii) every nonempty connected subfamily of O is tight;
- (iii) two ordinarily adjacent edges intersect in exactly two vertices;
- (iv) every connected ordering has a well-defined attachment pair at each step after the first; and
- (v)

$$|\partial_2 O| = 5e + 1. \quad (2.2)$$

Moreover, the ordinary shadows of distinct ordinary components are disjoint.

Proof. If O had at least k edges, a tree traversal in its ordinary edge-intersection graph would give a connected ordering of k edges. The first edge has four vertices and every later edge adds at most two, so these k edges would span at most $2k + 2$ vertices, contrary to the k -edge support wall.

Now let $A \subseteq O$ be nonempty and connected, with $|A| = t < k$. A connected ordering gives $|V(A)| \leq 4 + 2(t - 1) = 2t + 2$, while (1.1) gives the reverse inequality. Thus A is tight. Applied to two adjacent edges, this shows that their intersection has exactly two vertices. At every subsequent step of a connected ordering, the same equality forces the new edge to meet the preceding support in exactly its attachment pair. Of the six pairs in that edge, precisely the attachment pair has appeared earlier; hence the edge contributes exactly five new pairs to the ordinary shadow. Starting with six pairs proves (2.2). Finally, a pair in the shadows of two components would make an edge of one ordinarily adjacent to an edge of the other, which is impossible. $\square$

Definition 2.3 (Proper tight blocks and closure). A *proper tight block* is a tight edge family of rank below k . It is *maximal* if it is inclusion-maximal among all proper tight blocks of G . An edge family $K \subseteq G$ is *pair-closed* if every ambient edge $a \in G$ that contains a pair supported by K already belongs to K . Two edge families are *incomparable* if neither contains the other. An edge family M *meets* an ordinary component O if $M \cap O \neq \emptyset$.

Lemma 2.4 (Closure of maximal blocks). *Every maximal proper tight block M is pair-closed. More generally, every pair-closed edge family contains each whole ordinary component that it meets. Distinct maximal proper tight blocks are incomparable.*

For a pair-closed K , define its nonshadow reservoir

$$Q(K) := \binom{V(K)}{2} \setminus \partial_2 K. \quad (2.3)$$

Every pair in $Q(K)$ is globally nonshadow.

Proof. Suppose $a \in G \setminus M$ contains a pair supported by M . Then

$$|V(M + a)| \leq 2(|M| + 1) + 2.$$

If $|M| + 1 < k$, the support wall forces equality, so $M + a$ is a larger proper tight block, contradicting maximality. If $|M| = k - 1$, then $M + a$ consists of k edges on at most $2k + 2$ vertices, contrary to the k -edge support wall. Thus M is pair-closed.

If a pair-closed family K contains one edge of an ordinary component, follow an ordinary path from that edge. Every first edge of the path outside K contains a pair supported by K , contradicting pair-closedness. Hence K contains the whole component. Maximality immediately makes two distinct maximal blocks incomparable.

Finally, if $p \in Q(K)$ were shadowed by an ambient edge, pair-closedness would put that edge in K , giving $p \in \partial_2 K$, a contradiction. $\square$

3. FUNDAMENTAL TOKENS AND PRIMARY PORTS

Definition 3.1 (Fundamental token). For each positive ordinary component O , fix a spanning tree of its ordinary edge-intersection graph. A tree edge, or *tree link*, joins two edges that, by Lemma 2.2, have the form

$$D = \{R \cup X, R \cup Y\}, \quad (3.1)$$

where R, X, Y are pairwise disjoint two-element vertex sets. The two members of D are its *source edges*, the two-edge family D is the *source diamond*, R is its *root pair*, and the component O is its *source component*. The tree link together with this data is a *fundamental token*. Its four *primary ports* are the four cross-pairs in $X \times Y$. Thus a component with e edges has exactly $e - 1$ fundamental tokens.

A token is *assigned* to an edge family K if its entire source component is contained in K . A primary port *survives* if it is globally nonshadow and is *consumed* if it is globally shadowed. An ambient ordinary component contained in K , other than the source component, is called an *outside component* for the token.

For a consumed port P , a *blocker* is an edge $c \in K$ that shadows P and belongs to an outside component. Its two-vertex remainder $c \setminus P$ will be denoted by B . A vertex pair B is *fresh relative to an edge family A* if $B \cap V(A) = \emptyset$. In a chronological construction, a *private fresh pair* for a newly added edge is a designated two-vertex remainder that is disjoint from the support of every earlier edge; private fresh pairs of edges added in the same round are additionally required to be pairwise disjoint.

Lemma 3.2 (Primary-port lemma). *Let a fundamental token with source component O be assigned to a pair-closed edge family K of rank below k . Then:*

- (i) *no edge of O shadows a primary port;*
- (ii) *each outside component shadows at most one primary port;*
- (iii) *each consumed port P_i has a blocker $c_i \in K$, and blockers for distinct consumed ports may be chosen in distinct outside components;*
- (iv) *$c_i \cap V(D) = P_i$, so $B_i := c_i \setminus P_i$ is fresh relative to D ; and*
- (v) *the pairs B_i are mutually disjoint, and $D + c_1 + \dots + c_t$ is tight for every chosen collection of $t \leq 4$ blockers.*

The family consisting of D and one blocker for every consumed port is called the token's full primary core. The chosen blockers c_i are its primary blockers, and each B_i is the corresponding blocker's private fresh pair in that core.

Proof. Suppose first that an edge of O shadows a primary port P . Choose such an edge at minimum ordinary-path distance from the two source edges of D , expose a shortest ordinary path to it, and include both source edges of D . Let A^- be the resulting edge family before the last edge a is exposed, and put $A := A^- + a$. The family A^- is connected and tight. The edge a contains both its predecessor attachment pair and P . These pairs are distinct, since otherwise the predecessor would shadow P at smaller distance. Their union has at least three vertices, so a adds at most one vertex to $V(A^-)$. Consequently $|V(A)| \leq 2|A| + 1$, contrary to (1.1); here $A \subseteq O$, so $|A| < k$ by Lemma 2.2. This proves (i).

For (ii), suppose an outside component shadows two distinct primary ports. The union of those ports has at least three vertices. The source component and that outside component are edge-disjoint tight families whose total rank is below k , so an intersection of at least three support vertices would again make their union violate (1.1). Thus an outside component shadows at most one port.

If P_i is consumed, an ambient edge shadows it. Since P_i is supported by K , pair-closedness puts this edge in K ; by (i) it lies in an outside component. Part (ii) makes the components for distinct ports distinct, proving (iii). If c_i met $V(D)$ outside P_i , then the three-edge family $D + c_i$ would span at most seven vertices, contrary to (1.1). Hence (iv) holds.

With t selected blockers, the family $D + c_1 + \dots + c_t$ has $2 + t < k$ edges and support at most $6 + 2t$. The support wall forces equality. Consequently no two blocker remainders intersect and every one is fresh relative to the preceding core. This proves (v). $\square$

Definition 3.3 (Token types). A token is *robust*, *one-port*, or *zero-port* according as at least two, exactly one, or no primary ports survive. A token is *hard* if it is one-port or zero-port. The unique surviving port of a one-port token is its *persistent port*. Every edge introduced in the token constructions below other than the two source edges is called *non-source*.

4. THE SIMULTANEOUS TERMINATING CORONA

Definition 4.1 (Corona terminology). A *pair group* is a finite set of vertex pairs. It is *unshadowed in A* if none of its pairs is shadowed in the edge family A.

Fix a hard token assigned to a pair-closed K of rank below k , and index its consumed-port blockers as c_0, c_1, c_2 in the one-port case and as c_0, c_1, c_2, c_3 in the zero-port case. The full primary core is therefore

$$H_5 := D + c_0 + c_1 + c_2 \quad \text{or} \quad H_6 := D + c_0 + c_1 + c_2 + c_3,$$

respectively. The subscript records the rank. Let P_0 be the port blocked by c_0 , and let $B_0 = c_0 \setminus P_0$. The *initial carrier* and its *initial frontier group* are

$$C_0 := D + c_0, \quad \mathcal{R}_0 := B_0 \times (V(D) \setminus P_0). \quad (4.1)$$

A *corona branch* is a sequence

$$(C_0, \mathcal{R}_0), (C_1, \mathcal{R}_1), \dots$$

formed by the following recursion. If $x \in \mathcal{R}_j$ is selected for expansion, choose an edge $e_x \in K$ that shadows x , put $B_x := e_x \setminus x$, and set

$$C_{j+1} := C_j + e_x, \quad \mathcal{R}_{j+1} := B_x \times (V(C_j) \setminus x).$$

Every selected edge e_x is a *corona edge*. The edge family C_j is the branch's *carrier at depth j*, $\mathcal{R}_j$ is its *frontier group*, and $C_0, \dots, C_j$ are its *branch prefixes*. If the branch selects $x_i \in \mathcal{R}_i$ at step i , its *branch edges* are $c_0, e_{x_0}, e_{x_1}, \dots$; the parent of e_{x_0} is c_0 , and the parent of $e_{x_{i+1}}$ is e_{x_i} .

The construction is run simultaneously on all current frontier groups. A one-port group is *terminal* when it contains a globally nonshadow pair; a zero-port group is terminal when it contains two. In a nonterminal one-port group every pair is *prescribed* for expansion; in a nonterminal zero-port group all of its globally shadowed pairs are prescribed, so at most one pair is not prescribed. A *frontier round* expands every prescribed pair in every current group, and the child groups form the next frontier. The *full cumulative state* T_j is the full primary core together with all corona edges selected in the first j rounds; thus $T_0 := H_5$ or H_6 . A globally nonshadow pair selected from a terminal group is a *corona resource*. A branch that ends at a terminal group is a *terminal branch*; its last carrier is the branch's *terminal carrier*.

Lemma 4.2 (Simultaneous corona lemma). *Let K be a pair-closed edge family of rank below k , and let a hard token be assigned to K . Then the construction in Definition 4.1 is well-defined and terminates. More precisely, at depth j ,*

$$\text{rk}(C_j) = 3 + j, \quad |\mathcal{R}_j| = 8 + 4j, \quad (4.2)$$

the carrier C_j is tight, and $\mathcal{R}_j$ is unshadowed in C_j . Every full cumulative state T_j is tight; its frontier groups are mutually disjoint and unshadowed in T_j . Every new corona edge has a private fresh pair disjoint from the preceding full cumulative state and from the other fresh pairs chosen in the same round.

At termination, a one-port token has a corona resource z , while a zero-port token has two distinct corona resources z_1, z_2 . All these pairs are globally nonshadow and supported by K .

Proof. Lemma 3.2 shows that C_0 is tight of rank three. The set $V(D) \setminus P_0$ has four vertices, so $\mathcal{R}_0$ has eight pairs. It is unshadowed not only in C_0 but in the full primary core: the other blockers have fresh pairs disjoint from B_0 and contain no pair of $B_0 \times (V(D) \setminus P_0)$.

Suppose inductively that C_j is a tight rank- $(3+j)$ carrier and $\mathcal{R}_j$ is unshadowed in C_j . If $x \in \mathcal{R}_j$ is globally shadowed, then x is supported by K ; pair-closedness supplies an edge

$$e_x = x \cup B_x \in K, \quad (4.3)$$

where $B_x := e_x \setminus x$. Because x is unshadowed in C_j , the edge e_x is new. The family $C_j + e_x \subseteq K$ has rank below k , and it spans at most two more vertices than C_j . Equality in (1.1) is therefore forced: e_x meets $V(C_j)$ exactly in x , the pair B_x is fresh relative to C_j , and

$$C_{j+1} := C_j + e_x, \quad \mathcal{R}_{j+1} := B_x \times (V(C_j) \setminus x). \quad (4.4)$$

The new group is unshadowed in C_{j+1} , and since $|V(C_j)| = 2(3+j) + 2 = 8 + 2j$, it has $2(6 + 2j) = 12 + 4j$ pairs. This proves the branch assertions (4.2) by induction.

We next verify the simultaneous invariant. In a nonterminal round, every pair of a one-port frontier group is globally shadowed, while all but at most one pair of a zero-port group are globally shadowed. Choose one shadowing edge for every prescribed pair. Each chosen edge lies in K by pair-closedness and outside the current full cumulative state because the current frontier groups are unshadowed there. Such an edge shadows at most one prescribed frontier pair: two distinct prescribed pairs have at least three endpoints, so an edge shadowing both would meet the current tight support in at least three vertices and violate (1.1). Thus all chosen edges are distinct.

Adding them simultaneously introduces at most two vertices per edge. The resulting family remains a subfamily of K of rank below k , so the support wall forces equality. The new edge remainders are therefore pairwise disjoint private fresh pairs, and the new full cumulative state is tight. The next frontier groups are mutually disjoint and unshadowed in that entire state. Indeed, an old edge contains no vertex of a new fresh pair; the edge that created B_x contains B_x only together with x , whereas its child group uses

$$B_x \times (V(C_j) \setminus x);$$

and the other same-round edges have disjoint fresh pairs. Hence every edge chosen in the following round is genuinely new.

Every nonterminal round adds at least one previously unused edge of the finite family K , so there are at most $|K| - |T_0|$ such rounds. If a current full cumulative state exhausted K , then every pair unshadowed in that state would be globally nonshadow by pair-closedness. Each frontier group has at least eight pairs, contradicting the relevant nonterminal condition. Thus a terminal group appears earlier. Selecting one or two of its globally nonshadow pairs gives z or z_1, z_2 , respectively. Their endpoints lie in the terminal carrier and hence in $V(K)$. $\square$

5. EVEN-RANK STATES

Definition 5.1 (Base and ancestry). The *stopping depth* j of a hard token is the depth of a chosen terminal corona branch. Its terminal carrier has rank $3 + j$. If this rank is even, the terminal carrier is the token's *selected base*. If it is odd, add an unused primary blocker; the resulting tight family is the *promoted base* and is taken as the selected base. In both cases the selected base has even rank

$$\beta := 2 \left\lceil \frac{3+j}{2} \right\rceil. \quad (5.1)$$

The transitive closure of the parent relation on branch edges from Definition 4.1 is the *ancestry relation*. An edge set is *ancestrally closed* if it contains every ancestor of each corona edge it contains. An *ancestry-first order* places every ancestor before its descendants. Relative to a chosen terminal branch, an *off-branch edge* is a corona edge not on that branch. Two branches *diverge* at the first depth at which they use different child edges of the same carrier.

The promoted base in Definition 5.1 is well-defined. Every hard token has at least three primary blockers, and only c_0 lies on the chosen branch. Lemma 3.2 makes an unused blocker's private pair

disjoint from $D + c_0$. Moreover, the simultaneous construction starts from the full primary core, so Lemma 4.2 makes every later corona edge's private fresh pair disjoint from every unused blocker. Adding such a blocker to an odd-rank terminal carrier therefore preserves tightness.

Definition 5.2 (Resource lists and rooted states). For each token τ assigned to a pair-closed edge family K of rank below k , define its *two-resource list* $L(\tau)$, a two-element set of globally nonshadow pairs supported by K , as follows:

- for a robust token, choose any two surviving primary ports;
- for a one-port token, use $\{p, z\}$, where p is its persistent port and z is a corona resource; and
- for a zero-port token, use its corona resources $\{z_1, z_2\}$.

Each member of $L(\tau)$ is called a *resource* of τ . A tight state is *p-rooted* if p is supported by that state.

Lemma 5.3 (Even-rank state lemma). *Let K be a pair-closed edge family of rank below k , and let a hard token assigned to K have stopping depth j and selected base rank β . In the one-port case, write p for its persistent port; in the zero-port case, write z_1, z_2 for its two corona resources.*

- (i) *In the one-port case, there is a tight p -rooted state at every even rank*

$$2, 4, \dots, \beta. \quad (5.2)$$

- (ii) *In the zero-port case, the full cumulative state T_j satisfies*

$$\begin{aligned} |T_j| &\geq 6 + \sum_{\ell=0}^{j-1} (7 + 4\ell) \\ &= 2j^2 + 5j + 6 \geq 2\beta - 2, \end{aligned} \quad (5.3)$$

and there is a tight state supporting both z_1, z_2 at every even rank

$$\beta, \beta + 2, \dots, 2\beta - 2. \quad (5.4)$$

In the zero-port case, in particular,

$$2\beta - 2 < k. \quad (5.5)$$

Proof. For a one-port token, the source diamond has rank two and supports p . The even-rank branch prefixes, together with the possible promoted base, give every rank in (5.2); the persistent port stays supported throughout.

For a zero-port token, a nonterminal frontier group at depth ℓ has $8 + 4\ell$ pairs and at most one globally nonshadow pair. It therefore contributes at least $7 + 4\ell$ distinct shadowing edges to the cumulative construction. Adding these contributions to the full primary core H_6 gives the first two quantities in (5.3); the final inequality follows directly from $\beta = 2\lceil(3 + j)/2\rceil$.

The selected base is an ancestrally closed subfamily of T_j . Starting from it, add every missing primary blocker and then all off-branch corona edges in ancestry-first order. By (5.3), this order has at least $|T_j| - \beta \geq \beta - 2$ edges. At the first divergence of two branches, their fresh pairs are disjoint, and descendants of either branch use only its ancestors and later fresh pairs. Hence the private pair of an off-branch edge is absent from the selected base and from all earlier off-branch edges. Once its ancestors are present, its attachment pair is supported, so the edge adds exactly its private fresh pair. The first $\beta - 2$ additions therefore produce tight states supporting both corona resources at the ranks in (5.4). The top state is a subfamily of K , whose rank is below k , and (5.5) follows. $\square$

Definition 5.4 (Boundary terminology). Fix a hard token with selected base rank β . The states supplied by Lemma 5.3 are the token's *boundary states*; their attainable ranks form its *boundary interval*. A *displayed base or boundary* is a particular such edge family together with the source diamond and the chronological order used in its construction. In the zero-port case, a *complementary state* for an integer S is the boundary state of rank $k - S$, when $\beta \leq k - S \leq 2\beta - 2$ and $k - S \equiv \beta \pmod{2}$.

6. BOUNDARY TRANSVERSALITY AND SOURCE RECURRENCE

Definition 6.1 (Resource multigraph and carriers). Given a finite token family with two-resource lists, its *resource multigraph* has the globally nonshadow resource pairs as vertices and one edge for each token, joining the two resources on that token's list. Parallel token edges are allowed. The *token line graph* has the token edges as its vertices and joins two distinct token edges when they share a resource vertex. A *resource path* is a path in this line graph, equivalently a sequence of distinct token edges in which consecutive edges share a resource. The *resource-graph distance* between two token edges is their distance in the token line graph. A *resource cycle* is a cycle in the resource multigraph; two parallel token edges form a cycle of length two. A path *omits* a token if it does not use that token's edge. If a token edge is deleted from a resource cycle, the remaining resource path is the token's *alternate path*; it omits that token.

The *base carrier* of a robust token is its source diamond, of rank two. The base carrier of a hard token is its selected base, of rank β . A *realized carrier family* is a token family together with the base carrier chosen for each token. It is a *realized carrier path* when its token edges form a resource path. Once the path is oriented, for each pair of consecutive carriers choose one resource that both support; this pair is the later carrier's *incoming resource*. A realized carrier family is *connected* if its token edges induce a connected subgraph of the resource multigraph.

For a realized carrier family $\mathcal{U}$, let $E(\mathcal{U})$ be the union of the hyperedge sets of its carriers. Its *nominal rank*, also called its *weight*, is the sum of its carrier ranks with multiplicity; its *actual rank* is $|E(\mathcal{U})|$. Equivalently, each token edge has weight equal to the rank of its base carrier, and the weight of a resource path, resource cycle, or token subfamily is the sum of its token-edge weights. The family is *edge-disjoint* if its nominal and actual ranks agree. Given a displayed boundary A and another realized carrier family $\mathcal{U}$, their *repeated edge set* is $J := A \cap E(\mathcal{U})$, and $d := |J|$ is the number of repeated edges.

Lemma 6.2 (Private-pair transversality). *Fix a displayed base or boundary A of a token assigned to a pair-closed edge family K of rank below k . Order the non-source edges of A by placing the primary blockers first and the corona edges afterward in construction order. Every non-source edge has a private fresh pair disjoint from the source diamond and from all earlier displayed edges. Moreover, distinct non-source edges lie in distinct ambient ordinary components, and none lies in the source component.*

Consequently, let $J \subseteq A$ and put $d := |J|$. If $d \geq 2$ and J contains at most one source edge, then

$$|V(J)| \geq 7 + 2(d - 2) = 2d + 3. \quad (6.1)$$

For $d = 1$, if q is any globally nonshadow pair, then the single edge in J together with q spans at least five vertices.

Proof. The private-pair assertion follows from Lemma 3.2 for primary blockers and from the simultaneous invariant in Lemma 4.2 for corona edges.

We prove the ambient-component assertion explicitly. When a non-source edge is added, its attachment pair is unshadowed in the preceding displayed family and its private pair is new. Consequently it is not ordinarily adjacent to an earlier displayed edge; the same observation for every later edge shows that each non-source edge is a one-edge ordinary component of A , distinct from the component containing the source diamond.

Suppose that two of these components, or one of them and the component of A containing the source diamond, belong to the same ambient ordinary component of G . Let O_1, O_2 be two such components of A , and choose a shortest ambient ordinary path

$$g_0, g_1, \dots, g_m, \quad g_0 \in O_1, \quad g_m \in O_2,$$

whose internal edges do not lie in A . We have $m \geq 2$, since $O_1 \neq O_2$. Lemma 2.4 and pair-closedness imply that the whole path lies in K . Starting with the tight family A , adjoin $g_1, \dots, g_{m-2}$, if any, in order. Each edge contains its predecessor pair and hence adds at most two vertices. Since every resulting family is a rank-below- k subfamily of K , (1.1) forces equality at each step, so every

prefix remains tight. The last internal edge g_{m-1} contains both its predecessor pair and a pair of g_m . The two pairs are distinct by shortestness, and their union has at least three vertices. Thus g_{m-1} adds at most one vertex to a tight prefix, contrary to (1.1). Hence these components are distinct even in G .

Now delete the non-source edges of J in reverse construction order. At each deletion, the edge's private fresh pair disappears from the remaining support. The last two edges cannot be the two source edges and belong to different ordinary components, so they span at least seven vertices. Restoring the other $d - 2$ edges adds at least two vertices per edge and proves (6.1). If $d = 1$, a globally nonshadow pair q is not contained in the repeated ambient edge, so their endpoint union has at least five vertices. $\square$

Lemma 6.3 (Source-recurrence exclusion). *Let an edge-disjoint realized carrier path of nominal rank below k omit a token with source edges a, b and root pair $R = a \cap b$. The path repeats at most one of a, b . Hence its overlap with any displayed boundary of the omitted token satisfies the hypotheses of Lemma 6.2.*

Proof. Every prefix of the path is tight. Indeed, the first carrier is tight, and each later carrier is edge-disjoint from the preceding prefix and shares its incoming resource pair with that prefix. The resulting support has size at most $2\rho + 2$, where $\rho < k$ is the prefix's nominal rank, so (1.1) forces equality.

Each link of each fixed spanning tree defines only one token, so no other carrier has the same source diamond. A carrier cannot contain both a and b as non-source edges, because its distinct non-source edges lie in distinct ambient ordinary components, whereas a, b lie in the same source component. The mixed case is impossible as well: if a carrier uses one of a, b in its own source diamond, all its non-source edges lie outside that source ordinary component, so it cannot contain the other.

It remains to exclude two path carriers containing a and b separately. Expose the later carrier. Its support meets the previous path in the root pair R , because the earlier support contains the other source edge, and also in its incoming globally nonshadow resource. The incoming resource differs from R , since R is shadowed by both a and b . The two pairs have at least three endpoints. Thus the later tight carrier meets the preceding tight support in at least three vertices, which violates (1.1). Therefore at most one source edge recurs. $\square$

7. THE TOKEN SYSTEM OF DISTINCT REPRESENTATIVES

Definition 7.1 (Hall terminology). Given a finite token family $\mathcal{T}$ and a subfamily $\mathcal{S} \subseteq \mathcal{T}$, let

$$N(\mathcal{S}) := \bigcup_{\tau \in \mathcal{S}} L(\tau),$$

where $L(\tau)$ is the two-resource list of τ . The family $\mathcal{T}$ satisfies *Hall's condition* if $|N(\mathcal{S})| \geq |\mathcal{S}|$ for every $\mathcal{S} \subseteq \mathcal{T}$. A subfamily $\mathcal{S}$ is *Hall-deficient* if $|N(\mathcal{S})| < |\mathcal{S}|$, and is *inclusion-minimal deficient* if none of its proper subfamilies is deficient. A *system of distinct representatives*, abbreviated *SDR*, is an injective choice of one resource from each token's list.

The *cycle rank* of a finite multigraph with c connected components is $|E| - |V| + c$. A token edge is *cyclic* if it belongs to a resource cycle. Given an oriented weighted resource path and a positive threshold, its *first-crossing prefix* is the shortest initial subpath whose weight is at least that threshold.

Let $\mathcal{K}$ be a family of pairwise incomparable pair-closed edge families. An assignment of tokens to $\mathcal{K}$ is *whole-component* if all fundamental tokens of each positive ordinary component are assigned to one member of $\mathcal{K}$ containing that entire component.

Lemma 7.2 (Carrier rigidity). *Let $\mathcal{U}$ be a connected realized carrier family of nominal rank at most k . Then $\mathcal{U}$ is edge-disjoint. If its nominal rank is below k , then $E(\mathcal{U})$ is tight.*

Proof. First observe that an edge-disjoint realized carrier path of nominal rank below k is tight. Each individual carrier is tight, and two consecutive carriers support their common globally nonshadow

resource. The two endpoints of that resource lie in both supports, while no ambient edge contains the pair. Thus adjoining the next carrier creates support of size at most $2\rho + 2$, where ρ is the new total rank; the support wall forces equality. Induction proves the observation.

Suppose now that $\mathcal{U}$ repeats a hyperedge. Among all pairs of overlapping carriers, choose one at minimum distance in the token-line subgraph induced by $\mathcal{U}$. Let C be the last carrier on a shortest path in that induced subgraph, and let $\mathcal{P}$ be the realized carrier path formed by its predecessors. The carriers in $\mathcal{P}$ are edge-disjoint, have nominal rank at most $k - 2$, and form a tight union by the preceding observation. Put $J := E(\mathcal{P}) \cap C$ and $d := |J| \geq 1$. Lemma 6.3 shows that J contains at most one source edge of the last carrier. If $d \geq 2$, then the last carrier cannot be robust, since a robust carrier consists only of its two source edges. It is therefore a displayed hard-token base, and Lemma 6.2 gives $|V(J)| \geq 2d + 3$. If $d = 1$, the repeated edge together with the incoming globally nonshadow resource gives at least five common support vertices, which is the same bound $2d + 3$.

Both the earlier path and the last carrier are tight. Their union therefore has support at most

$$2(\text{nominal rank}) + 4 - (2d + 3) = 2(\text{actual rank}) + 1.$$

Because the actual rank is smaller than k , this contradicts (1.1). Hence no edge repeats. If the nominal rank is below k , choose a spanning tree of their induced token line graph and root it. Order its vertices parent before child. Every carrier after the first then shares a resource with the preceding union. The same support-wall induction used for paths shows that the full union is tight. $\square$

Lemma 7.3 (Resource-cycle bounds). *Every resource cycle has even weight at least $k + 2$. In addition, no resource cycle consists entirely of robust tokens.*

Proof. Suppose a resource cycle had weight at most k . By Lemma 7.2, its carriers are edge-disjoint. Delete one carrier. The remaining carrier path is tight, and the deleted carrier shares with it the two distinct resources at its endpoints. Those two pairs have at least three endpoints. Closing the cycle therefore creates support of size at most $2w + 1$, where w is the cycle weight. If $w < k$, this violates (1.1); if $w = k$, it violates the k -edge support wall. Thus $w > k$, and since every carrier rank is even, $w \geq k + 2$.

If a robust-only cycle has length at most $k/2$, its weight is at most k , which was just excluded. If it is longer, take $k/2$ consecutive rank-two carriers. Lemma 7.2 makes these k hyperedges distinct. Starting with six vertices and adjoining consecutive carriers across their shared resource pairs shows that they span at most $6 + 4(k/2 - 1) = 2k + 2$ vertices, contrary to the k -edge support wall. $\square$

Theorem 7.4 (Token SDR theorem). *Let $\mathcal{K}$ be a finite family of pairwise incomparable pair-closed edge families, each of rank below k . Suppose a whole-component assignment of fundamental tokens to $\mathcal{K}$ is given. Then the token lists admit an SDR. Thus one can choose, from every token's list, a globally nonshadow pair supported by its assigned edge family so that all chosen pairs are distinct.*

Proof. For each robust token, use the two surviving primary ports chosen in Definition 5.2 and its rank-two source diamond. For each one-port token, use $\{p, z\}$; for each zero-port token, use $\{z_1, z_2\}$. Give every hard token the selected even base supplied by Lemmas 4.2 and 5.3. All listed resources are globally nonshadow and are supported by the assigned pair-closed family.

Suppose Hall's condition fails, and let $\mathcal{S}$ be an inclusion-minimal deficient token family. Represent each token by its edge in the resource multigraph. If $t := |\mathcal{S}|$, minimality gives exactly $t - 1$ resource vertices: deficiency gives at most $t - 1$, while deleting any token gives at least $t - 1$. The multigraph is connected, since otherwise one proper connected component would already be deficient. Its cycle rank is therefore $t - (t - 1) + 1 = 2$, so it contains a resource cycle.

By Lemma 7.3, every resource cycle contains a hard token. Choose a cyclic hard token λ whose selected base rank β is maximum among cyclic hard tokens, and choose a resource cycle through λ . Every other token on that cycle has base-carrier rank at most β : this follows from maximality for hard tokens and from the rank-two definition for robust tokens. The alternate path obtained by deleting λ has weight at least $k + 2 - \beta$, by Lemma 7.3.

Suppose first that λ is one-port, with persistent port p . Orient the alternate path from p , and let S be the weight of the first-crossing prefix for threshold $k - \beta$. Its edge weights are even and at most

β , so

$$S \in \{k - \beta, k - \beta + 2, \dots, k - 2\}. \quad (7.1)$$

Attach the p -rooted state of rank $k - S$ supplied by (5.2).

Suppose instead that λ is zero-port. Orient the alternate path from either resource of λ , and let S be the weight of the first-crossing prefix for threshold $k - (2\beta - 2)$. The threshold is positive by (5.5), and the alternate path reaches it because its weight is at least $k + 2 - \beta$. The first-crossing weight satisfies

$$S \in \{k - (2\beta - 2), k - (2\beta - 4), \dots, k - \beta\}. \quad (7.2)$$

Attach the complementary state of rank $k - S$ from (5.4).

In either case, the path prefix and the attached boundary have nominal rank exactly k and share the resource from which the path was oriented. The bounds (7.1) and (7.2) give $S < k$, so Lemma 7.2 makes the path prefix edge-disjoint and tight. The prefix omits λ . If their hyperedge sets are disjoint, their tight supports overlap in the two endpoints of this resource, so the k edges span at most $2k + 2$ vertices, contrary to the k -edge support wall.

Otherwise, let $d \geq 1$ be the number of repeated boundary edges. Lemma 6.3 permits the use of Lemma 6.2. For $d \geq 2$, the common support has at least $2d + 3$ vertices; for $d = 1$, the repeated edge and shared resource give five common vertices. In both cases the actual rank is $k - d$, and the support is at most

$$2k + 4 - (2d + 3) = 2(k - d) + 1, \quad (7.3)$$

contrary to (1.1). Thus no deficient family exists. Hall's marriage theorem now supplies the required SDR. $\square$

8. PROOF OF THE FINITE THEOREM

Proof of Theorem 1.5. Let $C(G)$ be the set of all ordinary components of G , including singleton components, and let m be its cardinality. Every positive component O is a proper tight block by Lemma 2.2. Because G is finite, extend it to a maximal proper tight block $M(O)$. If the same block is chosen more than once, retain only one copy in the family of assigned blocks. By Lemma 2.4, the resulting distinct blocks are pairwise incomparable and pair-closed.

Assign all $|O| - 1$ fundamental tokens of each positive component O to $M(O)$, and apply Theorem 7.4 simultaneously. Singleton components contribute zero tokens, so the theorem supplies

$$\sum_{O \in C(G)} (|O| - 1) = |G| - m \quad (8.1)$$

distinct globally nonshadow vertex pairs. These pairs are disjoint from every ordinary shadow. Lemma 2.2 and the disjointness of component shadows show that those shadows contain

$$\sum_{O \in C(G)} (5|O| + 1) = 5|G| + m \quad (8.2)$$

distinct pairs. The two classes in (8.1)–(8.2) therefore exhibit $6|G|$ distinct pairs of vertices of G . Consequently,

$$6|G| \leq \binom{n}{2}, \quad (8.3)$$

as required. $\square$

9. ASYMPTOTIC CONCLUSION

Let $k \geq 10$ be even. Theorem 1.5 gives $\text{ex}_4(n, \mathcal{F}_k^{(4)}) \leq \binom{n}{2}/6$. Proposition 1.6 therefore yields

$$\limsup_{n \rightarrow \infty} \frac{f^{(4)}(n; 2k + 2, k)}{n^2} \leq \frac{1}{12}. \quad (9.1)$$

We state the published lower input in the exact form needed here. Let $r \geq 3$, let F be a finite r -graph, and let q be a positive integer. Define $P_{\leq q}(F)$ to be the set of pairs $p \in \binom{V(F)}{2}$ for which there exist an integer i and a subfamily $A \subseteq F$ satisfying

$$1 \leq i \leq q, \quad |A| = i, \quad |V(A) \cup p| \leq (r-2)i + 2.$$

Proposition 9.1 (Published supported-packing bound). *Let $r \geq 3$ and $k \geq 2$ be fixed, put $q := \lfloor k/2 \rfloor$, and let F be a nonempty finite $\mathcal{F}_k^{(r)}$ -free r -graph. Then*

$$\liminf_{n \rightarrow \infty} \frac{f^{(r)}(n; (r-2)k + 2, k)}{n^2} \geq \frac{|F|}{2|P_{\leq q}(F)|}. \quad (9.2)$$

Proof. This is Lemma 3.1 of Wang and Zeng [7], specialized to the notation above. Their proof invokes the supported-packing theorem of Glock–Joos–Kim–Kühn–Lichev–Pikhurko [4]; the underlying conflict-free matching method is developed separately by Glock–Joos–Kim–Kühn–Lichev [8]. $\square$

Take $r = 4$, fix an even $k \geq 2$, and let F consist of one edge. It is $\mathcal{F}_k^{(4)}$ -free, and $P_{\leq k/2}(F)$ consists of precisely the six pairs in that edge. Proposition 9.1 gives

$$f^{(4)}(n; 2k + 2, k) \geq \left(\frac{1}{12} - o(1) \right) n^2, \quad (9.3)$$

where $o(1)$ denotes a quantity tending to zero as $n \rightarrow \infty$.

The exact cases $k = 2, 4, 6, 8$ are already known. Specifically, Rödl's approximate-Steiner packing theorem gives $k = 2$ [3]; Glock–Joos–Kim–Kühn–Lichev–Pikhurko give $k = 4$ for every $r \geq 4$ [4]; Glock–Kim–Lichev–Pikhurko–Sun give $k = 6$ for every $r \geq 4$ [5]; and Pikhurko–Sun give $k = 8$ for every $r \geq 4$ [6].

Proof of Theorem 1.3. For every even $k \geq 10$, the upper estimate (9.1) and lower estimate (9.3) agree. The four published cases just listed cover the remaining positive even integers. Therefore, for every fixed even $k \geq 2$,

$$\lim_{n \rightarrow \infty} \frac{f^{(4)}(n; 2k + 2, k)}{n^2} = \frac{1}{12}. \quad (9.4)$$

$\square$

REFERENCES

- [1] M. Delcourt and L. Postle, *The limit in the $(k+2, k)$ -problem of Brown, Erdős and Sós exists for all $k \geq 2$* , Proc. Amer. Math. Soc. **152** (2024), 1881–1891.
- [2] C. Shangquan, *Degenerate Turán densities of sparse hypergraphs II: a solution to the Brown–Erdős–Sós problem for every uniformity*, SIAM J. Discrete Math. **37** (2023), 1920–1929.
- [3] V. Rödl, *On a packing and covering problem*, European J. Combin. **6** (1985), 69–78.
- [4] S. Glock, F. Joos, J. Kim, M. Kühn, L. Lichev, and O. Pikhurko, *On the $(6, 4)$ -problem of Brown, Erdős and Sós*, Proc. Amer. Math. Soc. Ser. B **11** (2024), 173–186.
- [5] S. Glock, J. Kim, L. Lichev, O. Pikhurko, and S. Sun, *On the $(k+2, k)$ -problem of Brown, Erdős and Sós for $k = 5, 6, 7$* , Canadian J. Math. (2025), 1–43.
- [6] O. Pikhurko and S. Sun, *On the quadratic 8-edge case of the Brown–Erdős–Sós problem*, European J. Combin. **135** (2026), 104364.
- [7] Y. Wang and J. Zeng, *On the $(k+2, k)$ -problem of Brown, Erdős and Sós for even integers k* , arXiv:2603.19345 (2026).
- [8] S. Glock, F. Joos, J. Kim, M. Kühn, and L. Lichev, *Conflict-free hypergraph matchings*, J. London Math. Soc. **109** (2024), e12899.