A near-linear preserver for bounded symmetric connectivity

Self-contained proof writeup

September 10, 2026

Abstract

For every finite directed unit-capacity multigraph G on n vertices and every positive integer k, there is a spanning deletion subgraph H preserving $\min\{\mathrm{flow}(s,t), \mathrm{flow}(t,s), k\}$ for every vertex pair, with

$$|E(H)| \le 2^{40} k^4 n \exp\bigl(16\sqrt{\log(n+1)}\bigr).$$

Thus the answer to the stated question is affirmative, with a polynomial in k of fixed degree and a subpolynomial factor depending only on n. The proof is an extremal count in an inclusion-minimal preserver. Its local ingredient uses minimum large cut shores to organize overlapping outgoing and incoming cut regions. A shallow hierarchy then permits summation of these local counts. All supporting lemmas are proved below; no conjectural or computational assumption is used.

Contents

1 Statement and proof structure

All logarithms without a subscript are natural. All graphs are finite. Each arc, including each parallel copy, has capacity one. For distinct vertices define

$$\lambda_G(s,t) = \min\{\mathrm{flow}_G(s,t), \mathrm{flow}_G(t,s)\}, \qquad \lambda_{G,k}(s,t) = \min\{\lambda_G(s,t), k\}.$$

Diagonal capped values may be taken to be k and impose no additional condition. A spanning deletion subgraph has the same vertex set and retains a subset of the original arc occurrences, with their original unit capacities.

Theorem 1.1. *For every directed unit-capacity multigraph G with n vertices and every integer $k \geq 1$, there exists a spanning deletion subgraph H such that, for all distinct s, t,*

$$\lambda_{H,k}(s,t) = \lambda_{G,k}(s,t),$$

and

$$|E(H)| \leq 2^{40} k^4 n \exp\!\left(16\sqrt{\log(n+1)}\right). \tag{1.1}$$

In particular the result holds for simple directed graphs.

For $n \geq 2$, the bound is of the required uniform form:

$$|E(H)| \leq 2^{40} k^4 n^{1+\varepsilon(n)}, \qquad \varepsilon(n) = \frac{16\sqrt{\log(n+1)}}{\log n} \longrightarrow 0.$$

Both the exponent 4 of k and the function ε are independent of the input graph, and ε is independent of k.

Take an inclusion-minimal preserver and fix, for each retained arc, an actual pair whose capped symmetric connectivity it is necessary to preserve. The following proof bounds every arc of that unchanged graph. An auxiliary pruning argument is used only to count its original critical arcs; local auxiliary graphs do not have to be combined.

The proof proceeds through four ingredients.

1. Section 2 proves sparse local transfer certificates, including certificates whose cost is charged only to a specified set of vertices, even when paths use exterior vertices.
2. Section 3 turns a small arc-and-vertex separator into a count of critical arcs whose endpoints and actual witness class do not fit in one remaining strongly connected component.
3. Section 4 proves the local theorem for several unbreakable terminal sets. Its central-shore lemma is what makes mixed incoming and outgoing localizations summable.
4. Section 5 constructs a shallow hierarchy, assigns whole actual witness triples to regions, and sums the local theorem. Its final calculation proves Theorem 1.1.

The notion of a connectivity preserver here is the one considered by Hoppenworth, Saranurak, and Wang [1], who give the bounds $O(k4^k n \log n)$ and $O(n\sqrt{kn})$. Rooted bounded-flow preservation is also studied by Bansal, Choudhary, Dhanoa, and Wardhan [2]. The proof below includes the static rooted statements it needs and supplies its own composition argument.

2 Transfer matrices and the local pruning lemmas

All graphs below are finite directed multigraphs. Arcs have capacity one, and parallel arcs remain distinct. A linkage has an unspecified matching between its prescribed source and sink multisets. Its paths are arc-disjoint, but may share original vertices. We use fields of rational functions over $\mathbb{Q}$; in particular all fields used for choosing new coefficients are infinite.

2.1 A generic matrix recording linkages

For a directed graph D whose individual arcs have independent indeterminates w_a, let A be its weighted adjacency matrix, with rows indexed by tails and columns by heads. Thus $(I - A)^{-1}_{xy}$ has the formal walk expansion from x to y. The determinant of $I - A$ has constant term one and is nonzero.

Lemma 2.1 (Generic inverse minors)**.** *For equal-size vertex sets S, J of D, the minor $\det(I - A)^{-1}[S, J]$ is nonzero if and only if D has a vertex-disjoint linkage from S to J. A nonzero specialized minor, at any specialization where $I - A$ is invertible, implies such a linkage in the specialized support.*

Proof. Jacobi's complementary-minor identity gives, up to sign,

$$\det(I - A)^{-1}[S, J] = \frac{\det(I - A)[V(D) \setminus J, V(D) \setminus S]}{\det(I - A)}.$$

In a nonzero monomial of the numerator, the off-diagonal choices form vertex-disjoint paths from $S \setminus J$ to $J \setminus S$, together with disjoint directed cycles; unused vertices contribute diagonal ones. Vertices of $S \cap J$ are deleted on both sides and supply zero-length paths. Hence absence of a linkage makes the numerator identically zero, including after specialization.

Conversely, in a vertex-disjoint linkage every vertex of $S \cap J$ must belong to the one path both starting and ending there; after discarding its closed walk this is a zero-length path. Specialize to a vertex-disjoint collection of the resulting nontrivial directed paths, giving their arcs weight one and every other arc weight zero. The matrix is then invertible and the indicated inverse minor is a signed permutation determinant. It is nonzero. This also proves the claimed equivalence without relying on absence of cancellation between different generic determinant terms. □

Lemma 2.2 (Rank-one original-arc representation)**.** *For a graph on n vertices and $K \geq 1$, there is a generic Kn by Kn matrix pencil with one rank-one scalar coordinate per original arc. Its inverse minors record all original endpoint-multiset linkages of at most K paths via suitable ordinary-copy source and sink sets. Keeping such an original nonzero minor at an invertible specialization preserves its linkage in the retained unit-capacity arc support.*

Proof. Replace every original vertex v by K ordinary copies. For every original arc $e = u \to v$, introduce one private bottleneck vertex b_e, arcs from every copy of u to b_e, and arcs from b_e to every copy of v. Give these incidences independent indeterminates. A vertex-disjoint linkage of at most K paths in this expanded graph projects to an arc-disjoint original linkage. Conversely any original linkage of at most K paths lifts: assign distinct copies to its at most K visits at each original vertex. Paths can first be made simple. The bottleneck vertex is the unit-capacity resource for the original arc.

The bottleneck-to-bottleneck block of $I - A$ is the identity. Eliminating it leaves an ordinary Kn by Kn matrix

$$M(z) = I - \sum_e z_e \alpha_e \beta_e^{\mathsf{T}}, \tag{2.1}$$

where α_e is supported on the tail copies and β_e on the head copies. The original graph corresponds to $z_e = 1$ for all e. The ordinary inverse block of the expanded matrix equals $M(z)^{-1}$. Crucially there is *one* variable z_e for an original arc. Setting it to zero removes that arc's bottleneck gadget; a nonzero value never increases its unit combinatorial capacity. By Lemma 2.1, original nonzero minors retained at an invertible specialization certify linkages in the retained unit-capacity original support. Signed coefficients and coefficients depending rationally on the original indeterminates are allowed. This conclusion uses only the implication from a nonzero minor to a linkage, not genericity of the new coefficients. □

2.2 Sparse solutions of affine equations

Lemma 2.3 (Affine support with a nonvanishing polynomial)**.** *Over an infinite field, suppose an affine system in m variables has rank at most a, and a polynomial p of degree at most b is nonzero at some solution. There is such a solution with at most $a + b$ nonzero coordinates.*

Proof. Choose a feasible solution with $p \neq 0$ and minimum support S, and write $t = |S|$. The affine space L of solutions supported on S has dimension $d \geq t - a$. Its restricted polynomial is not identically zero. Every nonconstant coordinate function on L defines a nonempty affine zero hyperplane. Minimality of S implies that $p|_L$ vanishes on this hyperplane, and hence its defining affine linear polynomial divides $p|_L$. This divisibility follows by choosing that linear polynomial as one coordinate and evaluating its constant coefficient over the infinite field.

The linear parts of the coordinate functions span the dual of the translation space of L, because the coordinate embedding is injective. Choose d with independent linear parts. Their zero hyperplanes are distinct and give nonassociate linear factors of $p|_L$. Thus $b \geq d \geq t - a$, giving $t \leq a + b$. The case $d = 0$ is immediate. □

Lemma 2.4 (A locally charged two-sided terminal scaffold)**.** *Let W be a vertex set, let R be any set of terminals, and let E_* be selected original arcs with both endpoints in W. There is a subgraph retaining every arc outside E_* and at most*

$$K|W|(2K|R| + 1)$$

arcs of E_*, preserving every original endpoint-multiset linkage of at most K paths from R to arbitrary vertices and conversely. All linkages are in the ambient graph; R need not be contained in W.

Proof. Use (2.1), varying only the selected coefficients. Let E_R select all $K|R|$ terminal copies and fix the original matrices

$$Q = M(1)^{-1}E_R, \qquad P = E_R^{\mathsf{T}}M(1)^{-1}.$$

Impose $M(z)Q = E_R$ and $PM(z) = E_R^{\mathsf{T}}$. Only the $K|W|$ rows in the first system and $K|W|$ columns in the second change. Thus there are at most $2K^2|R||W|$ nonconstant scalar equations. Only those same $K|W|$ rows of $M(z)$ vary, so its determinant has degree at most $K|W|$ in the selected coordinates. Apply Lemma 2.3. At the resulting invertible matrix, both terminal inverse actions equal their original values. For any original linkage of at most K paths, choose one lift to ordinary copies; its source or sink copies in R are among the fixed coordinates. The corresponding nonzero minor is unchanged, and Lemma 2.1 supplies a linkage in the retained original support. Deletion gives the reverse inequality. □

2.3 Entering a region through a small return boundary

Lemma 2.5 (Selected entry pruning)**.** *Partition the vertices as $A\dot\cup B$, with b arcs from B to A. Let $R \subseteq A$ consist of p roots. Any selected set of arcs from A to B can be reduced to at most $b(Kp+1)$ arcs, freezing every other arc, while preserving every capped ordered root-to-A flow at cap K.*

Proof. In the ordinary matrix, factor the b return arcs and write

$$M = \begin{pmatrix} D & -X \\ -VW & E \end{pmatrix}, \qquad T = E^{-1}V.$$

The fixed generic principal blocks D, E are invertible. Only selected original scalar coordinates in X vary. The A Schur complement is $Q(X) = D - XTW$. If S selects the Kp root copies, fix the original row response $L = SQ(X)^{-1}$, and impose

$$LX'T = LXT.$$

There are at most Kpb linear equations. They give $LQ(X') = S$. Moreover

$$\det M' = \det E \det D \det(I_b - WD^{-1}X'T),$$

whose last factor is a polynomial of degree at most b in the selected coordinates, nonzero at the original point. The sparse-affine lemma retains at most $Kpb + b$ selected coordinates with this determinant nonzero. Then $SQ(X')^{-1} = L$, so all relevant original root-to-A minors stay nonzero. Decode their unit-capacity supports. If $b = 0$, the same argument retains no selected entry. □

Corollary 2.6 (External entry into a flat)**.** *Let F have b outgoing arcs and let $W \subseteq F$. Among selected arcs entering F which are individually critical for actual capped symmetric pairs in W, there are at most $b(K|W| + 1)$ arcs.*

Proof. Reverse the graph and apply Lemma 2.5 with outside shore $A = F$, roots W, and selected arcs from F to its complement. In the original orientation this preserves every capped flow from F to each member of W, hence every ordered W-to-W flow. Every selected critical arc must therefore be retained. The bound is independent of the number of exterior vertices. □

2.4 Rooted certificates charged only at designated vertices

Lemma 2.7 (Rooted critical localization)**.** *For a root r, at most K arcs entering any vertex $v \neq r$ are individually critical for any of the capped ordered flows from r. An arc $u \to v$ critical for some such flow is critical for the capped flow from r to v itself. No arc entering r is critical.*

Proof. Suppose $e = u \to v$ is critical for the target t. There is an incoming minimum-cut shore Y containing t, excluding r, and containing e in its boundary, of capacity $j \leq K$. Thus $v \in Y$, $u \notin Y$. Put $h = \text{flow}(r, v) \leq j$. Let T_v be the smallest incoming r-to-v minimum shore. It exists because the intersection of two such minimum shores is again minimum by cut submodularity. If $u \in T_v$, then

$$h + j \geq c^-(T_v \cap Y) + c^-(T_v \cup Y) \geq h + j.$$

The intersection contains v, the union contains t, and both exclude r, proving the second inequality. Equality makes $T_v \cap Y$ a smaller minimum v-shore excluding u, a contradiction. Hence e enters T_v, proving its criticality for v. All such arcs at v lie in this single boundary of size $h \leq K$. A simple path beginning at r never uses an arc entering r. □

An inclusion-minimal subgraph preserving all capped root-to-vertex flows therefore has indegree at most K at every nonroot vertex and no incoming root arcs: each retained arc is critical for some rooted requirement, so Lemma 2.7 applies in that subgraph. This supplies the capped rooted certificate used next, without invoking a separate branching theorem.

Lemma 2.8 (Charged rooted incidence)**.** *For any root r and any charged vertex set W, there is a subgraph preserving all capped flows $r \to w$, $w \in W$, with at most $K(K+2)|W|$ arcs incident to W. Every original arc with both endpoints outside W can be retained. The two-direction version costs at most $2K(K+2)|W|$.*

Proof. Take the capped rooted certificate just constructed. Put $A = W \cup \{r\}$ and $B = V \setminus A$. Keep its arcs whose heads lie in A, of total number at most $K|W \setminus \{r\}| \le K|W|$. Restore every original B-internal arc and every original A-to-B arc. The resulting graph contains the certificate. Its B-to-A boundary has $b \le K|W|$ arcs. Apply Lemma 2.5 with the single root r, retaining at most $b(K+1)$ A-to-B arcs. This preserves the original capped flows to A, and all arcs incident to W number at most $K|W| + b(K+1) \le K(K+2)|W|$. Finally restore any missing arc disjoint from W. This changes no charged count and cannot hurt preservation. Apply the same construction in the reversed graph and take a union for the last assertion. □

2.5 A one-shore exchange with no cap factor

Lemma 2.9 (One-shore reverse-arc pruning)**.** *Let $A \dot\cup B$ be a partition with b arcs from A to B. Any selected set of reverse arcs from B to A can be reduced to at most b entering each individual vertex of A, freezing all other arcs, while preserving every ordered flow between vertices of A, without truncation. Dually, reverse arcs critical for demands entirely in B have at most b selected arcs per tail in B.*

Proof. Index the b entry arcs into B by i. In an auxiliary graph on the resources of $H[B]$, use a private source port p_i entering the corresponding entry head. For each exit arc e from B to A, use a private absorbing terminal q_e attached to its tail. Add a private dummy d_i and the arc $p_i \to d_i$. Ports have no incoming arcs. Use a line-graph vertex for every internal B arc and every exit arc, with transitions for consecutive arcs. This converts unit arc disjointness, with unlimited original-vertex sharing, into ordinary vertex disjointness. Sources and dummies retain their private vertices.

Take the generic inverse matrix of this finite auxiliary graph and its submatrix from the b source ports to all exit and dummy terminals. By Lemma 2.1, a set of its terminal columns is independent exactly when those terminals can be linked from distinct ports. These columns define a linear matroid of rank b, since the b dummies are independent.

For each head $v \in A$, let P_v be the columns of selected exit arcs entering v. Choose a column basis K_v for their span and retain these at most b exits, as well as every unselected exit. The following simultaneous exchange is elementary. If an independent terminal set contains $p \in P_v \setminus K_v$, not every member of K_v can lie in the span of the remaining columns: otherwise p would also lie there. Replace p by a member of K_v outside that span. Independence, the number of terminals in each head group, and all dummy terminals are preserved. Iterating removes every deleted exit from the terminal set.

Now take any integral s-to-t flow with $s, t \in A$, represented by arc-disjoint paths. Its maximal excursions through B use some set J of entry ports, at most b in total, and the same number of exit arcs. Their auxiliary linkage, together with the dummies d_i for $i \notin J$, is an independent set of size b. Apply the exchange. A new linkage to the resulting b terminals uses every port. Each retained dummy d_i forces its own port p_i, so the replacement excursions use exactly the original entry set J. They preserve the number of exits into each individual A vertex. Keep all original path portions in A. Together with the new excursions these form an integral flow of the same value, because divergences and all capacities are preserved. Entry-to-exit pairings need not be preserved. Discard any directed cycles. This proves the claim; reversing the graph proves the dual. ☐

2.6 Charged arcs in the opposite shore

We record the remaining response estimates used in charged cut localization. They are counts of individually critical arcs; separate temporary pruning constructions are never assumed to compose.

Lemma 2.10 (Two charged far-shore estimates)**.** *Let $A \dot\cup B$ have b arcs from A to B. Let $W_A \subseteq A$ and $W_B \subseteq B$. Suppose selected arcs are individually critical for actual capped symmetric pairs in W_A. Then the following bounds hold:*

1. *selected arcs with tail in B and head in W_B number at most $K(b+1)|W_B|$;*
2. *selected arcs with tail in W_B and arbitrary head number at most $Kb|W_A| + K(b+1)|W_B|$.*

Proof. Let S select all copies of W_A and let $L = SM^{-1}$ be the original response. The block L_B has rank at most b. Indeed the inverse cross-block from A to B factors through the original rank-at-most-b block M_{AB}, by the Schur inverse formula.

For the first assertion, factor $L_B = PQ$ with Q having at most b rows. The selected change Δ has rows in B and columns in W_B. Impose $Q\Delta = 0$, requiring at most $bK|W_B|$ equations. This implies $L\Delta = 0$. Only $K|W_B|$ columns of M vary, so its determinant has degree at most $K|W_B|$. The sparse-affine lemma retains at most $K(b+1)|W_B|$ selected arcs with M invertible and the full W_A source response fixed. All selected critical arcs must be retained.

For the second assertion put $T = W_A \cup W_B$ and eliminate $O = V \setminus T$. The block M_{OO} and the block M_{OT} are fixed, since only rows in W_B vary. Hence the Schur matrix

$$D = M_{TT} - M_{TO} M_{OO}^{-1} M_{OT}$$

is affine in the selected original scalar coordinates, with changing rows only in W_B. Its inverse is $M^{-1}[T,T]$. The original W_A-source response L_T has restriction to W_B of rank at most b, as above. Factor that restriction as $P'Q'$ and impose

$$Q'(D' - D)[W_B, T] = 0.$$

There are at most $bK(|W_A| + |W_B|)$ equations, and $\det D'$ has degree at most $K|W_B|$. At a resulting nonsingular solution the whole original W_A-to-W_A inverse block is fixed. Decoding and criticality give the second bound. The fixed generic principal matrix M_{OO} is invertible, so Schur nonsingularity also ensures nonsingularity of the full matrix. ☐

The coefficient fields in all these applications contain the original generic matrix entries and its inverse. The new scalar coefficients may depend on that entire field. All imposed conditions are genuine affine linear equations over that fixed field; no coefficientwise independence of the new scalars, singular specialization, or positivity of the new weights is assumed.

3 Counting critical arcs across guarded regions

All cuts and connectivity values in this section refer to one unchanged digraph D. Write $c(X) = |\delta^+(X)|$. Its cut function is submodular:

$$c(X) + c(Y) \geq c(X \cap Y) + c(X \cup Y).$$

This follows by checking the contribution of each arc. Also, directed connectivity satisfies

$$\text{flow}(x, z) \geq \min\{\text{flow}(x, y), \text{flow}(y, z)\}.$$

Indeed, every x–z source shore separates x from y or y from z. Consequently, symmetric connectivity satisfies the same inequality, and the relations $\lambda(x, y) \geq j$ are equivalence relations, with the diagonal pairs included by convention. Their classes form a laminar hierarchy as j increases.

Lemma 3.1 (Split classes consume cut capacity)**.** *For each $j \in \{1, \ldots, k\}$, label every symmetric j-class with its level, counting repeated vertex sets separately at different levels. The number $g(X)$ of these labelled classes split by X satisfies*

$$g(X) \leq c(X).$$

Proof. First take disjoint, incomparable classes $C_1, \ldots, C_p$ at levels $r_1, \ldots, r_p$, all split by X. For representatives $v_i \in C_i$, put $w_{ij} = \lambda(v_i, v_j) < \min(r_i, r_j)$, and let w_* be the weight of a maximum spanning tree on these representatives ($w_* = 0$ if $p = 1$). We claim

$$c(X) \geq \sum_i r_i - w_*. \tag{3.1}$$

For $p = 1$ this is the definition of a connectivity class. Otherwise let $w = \min_{i \neq j} w_{ij}$ and take a directed cut Y of capacity w separating a pair attaining it, in its weaker direction. Every C_i is whole in Y, and Y partitions the chosen classes into nonempty collections I, J. Every cross-collection value equals w: it is at least the chosen minimum and at most the capacity of Y. The maximum spanning-tree weight is therefore $w_I + w_J + w$, where w_I, w_J are the two internal maximum spanning-tree weights. This identity follows also by the usual tree edge-exchange argument, since all cross weights equal w and every internal weight is at least w. The cut $X \cap Y$ splits the classes in I, and $X \cup Y$ splits those in J. Induction and cut submodularity give

$$c(X) + w \geq c(X \cap Y) + c(X \cup Y) \geq \sum_i r_i - w_I - w_J,$$

which is (3.1).

Now choose the inclusion-minimal vertex sets among all split classes, using their highest split level when a vertex set repeats. These classes are an antichain. Every split labelled class is an ancestor of one of them, and every ancestor of one of them is split. The ancestor path of C_i has r_i labels, and two such paths share exactly w_{ij} labels. Their union thus has $\sum_i r_i - w_*$ labels. For a direct verification, partition at the minimum common level w as above: the two collections have disjoint labels above w and share one label at each of the first w levels; induction gives the formula. A virtual level-zero root handles different original strongly connected components. Applying (3.1) proves the lemma. □

A selected arc is called *critical for an actual j-pair* if its deletion lowers the symmetric connectivity of that pair below j, and the pair has symmetric connectivity at least j in D. For each selected arc fix such a pair, its level $j \leq k$, and its whole actual j-class. All these choices remain fixed in the arguments below.

Lemma 3.2 (Localizing charged arcs at one cut)**.** *Let W be a set of p charged vertices. Suppose every selected arc is incident to W and its fixed actual witness pair lies in W. For a cut $A \to B$ of capacity b, at most*

$$A_k(b)p + 2b, \qquad A_k(b) = 2k(k+2)b + 4k(b+1), \tag{3.2}$$

selected arcs have a split actual class, or have an endpoint on the shore opposite their whole actual class. Every other selected arc has its endpoints and its whole class together on one shore.

Proof. There are at most b split labelled classes by Lemma 3.1. For each one choose a root in its intersection with W. The two-sided charged rooted certificate of Lemma 2.8 uses at most $2k(k+2)p$ charged arcs and preserves every required pair in that class, by transitivity through the root. Every selected critical arc assigned to that class must belong to the certificate. Thus the split classes account for at most $2k(k+2)bp$ arcs.

Consider next classes wholly in A, and write $W_A = W \cap A$, $W_B = W \cap B$. We bound their selected arcs that touch B. Use the k-clone rank-one matrix M from Lemma 2.2, its generic inverse, and the rows

$$L = S_{W_A} M^{-1}.$$

The block L_B has rank at most b: the A–B inverse block factors through the cut block M_{AB}, whose rank is at most the number b of unit arcs in the cut. Factor $L_B = PQ$ with Q having at most b rows. All matrices here are over the original rational function field.

First vary only selected arcs with tails in B and heads in W_B. The equations $Q\Delta M_{B,W_B} = 0$, at most $bk|W_B|$ scalar equations, preserve L whenever the new matrix is nonsingular. Its determinant has degree at most $k|W_B|$ in the varying coordinates, since only those columns vary. Lemma 2.3 and criticality therefore bound this group by

$$k(b+1)|W_B|. \tag{3.3}$$

Second, vary selected arcs with tails in W_B, allowing arbitrary heads. Set $T = W_A \cup W_B$ and $O = V(D) \setminus T$, and eliminate the fixed block M_{OO}. The Schur pencil

$$D_T(z) = M_{TT}(z) - M_{TO}(z) M_{OO}^{-1} M_{OT}$$

is affine, with changes only in the W_B rows. Its inverse is $M(z)^{-1}[T,T]$. The original W_A source response restricted to W_B has rank at most b. Factor it as $P'Q'$ and impose

$$Q'\Delta D_T[W_B, T] = 0.$$

There are at most $bk(|W_A| + |W_B|)$ equations; the determinant degree is at most $k|W_B|$. The complete W_A–W_A response is preserved, so this group has at most

$$kb|W_A| + k(b+1)|W_B| \tag{3.4}$$

critical arcs. The fixed principal block M_{OO} is generically invertible, as required for the Schur elimination.

Third, vary selected arcs from B to W_A. The equations $Q\Delta M_{B,W_A} = 0$ and determinant degree at most $k|W_A|$ give at most

$$k(b+1)|W_A| \tag{3.5}$$

critical arcs. These arguments preserve the minors encoding each actual ordered witness flow, so individual criticality forces every selected arc into the respective retained support. Rational or signed support coefficients are harmless by Lemma 2.2.

Every incident-W arc touching B lies in one of these three groups, unless it goes forward from A to B; there are at most b of those. Overlapping groups may simply be overcounted. Summing (3.3)–(3.5) gives at most $2k(b+1)p+b$. Reverse the entire graph to obtain the same bound for classes whole in B. Adding the split-class charge proves (3.2). □

Lemma 3.3 (Guarded critical triples)**.** *Let Z be g arcs and T be t vertices. Let W have $p \geq 2$ vertices, and suppose every selected arc is incident to W and has its fixed actual witness pair in W. If no selected arc, together with its whole actual class, is contained in one SCC of $D - Z - T$, put $b = g + kt$. Then the number of selected arcs is at most*

$$4pA_k(b)\lceil \log_2(2p) \rceil + 8bp. \tag{3.6}$$

For $p \leq 1$ there are no selected distinct-pair critical arcs.

The same bound holds in a region F with a g-arc guard Z such that F is a union of SCCs of $D - Z$: use only $T \subseteq F$, selected arcs internal to F, and $W \subseteq F$, and require triple separation in $(D-Z)[F] - T$.

Proof. Split each vertex of T into an in-copy and an out-copy, joined by k parallel arcs. Incoming original arcs enter the in-copy, outgoing ones leave the out-copy; represent each original query vertex by its out-copy. This preserves every capped ordered flow, also after deleting any selected original arc: lift at most k simple paths using distinct bridge arcs, and project a packing in the other direction. Charge both copies only for vertices of W, obtaining W' of size $m \leq 2p$.

Delete the transformed copies of Z and all kt bridge arcs. Each split in-copy is a sink and each split out-copy a source, so all its SCCs are singletons; the remaining SCCs are those of $D - Z - T$. The actual split-graph class of a selected pair includes all representatives of its original actual class. Thus no selected endpoint-plus-class triple fits inside one of these SCCs. A split singleton cannot contain the two distinct witness representatives.

Fix a topological ordering of these SCCs. Every suffix union has outgoing boundary at most b in the full split graph: only the deleted arcs can leave a suffix. Recursively maintain an active charged subset $W_0 \subseteq W'$ and selected arcs whose witnesses are in W_0 and which are incident to W_0. Choose a median SCC by its W_0 weight, so the charged mass strictly before and strictly after it is at most $|W_0|/2$. Apply Lemma 3.2 at the two suffix boundaries adjacent to this SCC. This costs at most

$$2A_k(b)|W_0| + 4b.$$

Every survivor has its entire triple either strictly before the median, inside it, or strictly after it. The middle possibility is excluded by the hypothesis. Recurse on the two side sets of charged vertices. All classes, cut capacities, and criticality statements remain those of the unchanged full split graph; auxiliary counting deletions at different recursion nodes are never combined.

The total charged mass at each depth is at most m, the number of nontrivial depths is at most $\lceil \log_2 m \rceil$, and there are at most m internal recursion nodes: each median discards at least one charged vertex. Hence the total is at most

$$2A_k(b)m\lceil \log_2 m \rceil + 4bm,$$

which implies (3.6).

For the regional assertion, no SCC of $D - Z$ crosses F. Deleting the additional vertices cannot create a crossing SCC, so the global SCCs meeting F are precisely the local ones relevant to the hypothesis. The first assertion applies. In particular a region with outgoing boundary h is covered by taking $Z = \delta^+(F)$ and $g = h$. □

4 Simultaneous unbreakability and central large shores

All capacities in this section are capacities in the same unchanged unit-capacity directed multigraph H. Parallel arcs are counted separately. Write $c^+(X) = |\delta_H^+(X)|$ and $c^-(X) = |\delta_H^-(X)|$. Write $\varphi_H(x, y)$ for the directed maximum flow, with the reflexive convention $\varphi_H(x, x) = +\infty$, and put $\lambda_H(x, y) = \min\{\varphi_H(x, y), \varphi_H(y, x)\}$ for distinct vertices. A set U is (Q, k)-*unbreakable* if every outgoing cut of capacity at most k has at most Q members of U on at least one shore. The same assertion holds for incoming shores by complementation.

4.1 One rank for several small-shore constraints

Lemma 4.1 (Simultaneous capped cut rank)**.** *Let $U_1, \ldots, U_L$ be (Q, k)-unbreakable sets, where $Q, k \geq 1$ are integers and $|U_i| > 3Q$. Fix a nonempty index set $I \subseteq \{1, \ldots, L\}$. Call Y feasible if $|Y \cap U_i| \leq Q$ for every $i \in I$, and set*

$$\rho(S) = \min\{k,\ \min\{c^+(Y) : S \subseteq Y,\ Y \text{ feasible}\}\}, \tag{4.1}$$

where an empty inner minimum is $+\infty$. Then ρ is a normalized, monotone, integer-valued submodular rank, with $\rho(V(H)) = k$.

For this rank write

$$\mathrm{cl}_\rho(S) = \{x : \rho(S \cup \{x\}) = \rho(S)\}.$$

Every closure F of rank less than k is itself feasible and satisfies $c^+(F) = \rho(F)$. If $F_x = \mathrm{cl}_\rho(\{x\})$ and $W_F = \{x : F_x = F\}$, the nonempty low singleton cells W_F are disjoint and satisfy $W_F \subseteq F$. If $B \subsetneq F$ are flats, then $\rho(B) < \rho(F)$ and $W_F \cap B = \varnothing$. All statements have incoming-cut analogues.

Proof. Normalization, integrality, monotonicity, and the total rank are immediate. To prove submodularity, if either $\rho(A)$ or $\rho(B)$ equals k, use the cap for their union and monotonicity for their intersection. Otherwise choose feasible minimizing shores X, Y for A, B. Their intersection is feasible. If $c^+(X \cup Y) \geq k$, then $\rho(A \cup B) \leq k \leq c^+(X \cup Y)$. If $c^+(X \cup Y) < k$, then for each $i \in I$ the union has at most $2Q$ terminals of U_i and its complement has more than Q. Unbreakability forces the union to have at most Q terminals of each such bag. Thus it is feasible and again $\rho(A \cup B) \leq c^+(X \cup Y)$. Together with $\rho(A \cap B) \leq c^+(X \cap Y)$ and directed cut submodularity, this proves the claim.

For completeness, submodularity implies the standard closure facts used here. If adjoining each of x, y separately does not increase the rank of S, then submodularity applied to $S + x, S + y$ shows that adjoining both does not increase it. Repeating gives $\rho(\mathrm{cl}_\rho(S)) = \rho(S)$. The same diminishing returns inequality proves monotonicity of closure. Its closure is closed, and any flat properly containing another has strictly larger rank.

Let a closed F have rank $r < k$. A feasible minimizing shore $Y \supseteq F$ has $c^+(Y) = r$. Monotonicity and the definition of ρ give $\rho(Y) = r$, so closedness of F implies $Y \subseteq F$. Hence $Y = F$, proving the actual-capacity assertion. The statements about singleton cells follow from closure monotonicity: if $x \in B$ and B is closed, then $F_x \subseteq B$. □

Lemma 4.2 (Actual witness interval)**.** *Suppose $e = u \to v$ is individually critical for $\lambda_H(s, t) \geq j$, where $1 \leq j \leq k$. Let C_e be the entire symmetric j-class containing s, t, and let $d = \varphi_H(u, v)$. Then $1 \leq d \leq j$ and, for every $c \in C_e$,*

$$\varphi_H(u, c) \geq d, \qquad \varphi_H(c, v) \geq d. \tag{4.2}$$

In particular e and C_e lie in one original SCC.

Proof. Deleting one unit arc can decrease an ordered maximum flow by at most one. The lost direction of the witness pair therefore has value exactly j in H. It has a tight outgoing shore A of capacity j crossed by e, and A splits C_e. It follows that $d \le j$. The arc itself gives $d \ge 1$.

If $\varphi_H(u,c) = a < d$, a minimum u–c shore X contains u, excludes all of C_e, and contains v: a cut below j cannot split a symmetric j-class, and a cut below d cannot separate u from v. Then $A \cap X$ separates u, v, and $A \cup X$ splits C_e. Submodularity would give

$$d + j \le c^+(A \cap X) + c^+(A \cup X) \le j + a < j + d,$$

a contradiction. Reversing every arc and complementing A proves the other inequality. The critical arc belongs to a witness flow path: otherwise that path packing survives its deletion. A reverse witness path supplies a return, so its endpoints and the whole positive-level class belong to one SCC. □

4.2 A shared root set gives two simultaneous low shores

Lemma 4.3 (Joint root localization)**.** *Choose $R_i \subseteq U_i$ with $|R_i| = 2Q + 1$ and put $R = \bigcup_i R_i$. Suppose $J \subseteq H$ preserves every unpaired endpoint-multiset linkage of at most $k + 1$ paths from R to arbitrary vertices and from arbitrary vertices to R. For an arc $e = u \to v$ outside J satisfying Lemma 4.2, there is a partition O, I of the bag indices and outgoing and incoming shores Y_+, Y_- containing $\{u, v\} \cup C_e$ such that*

$$c^+(Y_+) < d, \quad c^-(Y_-) < d, \qquad |Y_+ \cap U_i| \le Q \ (i \in O), \quad |Y_- \cap U_i| \le Q \ (i \in I). \tag{4.3}$$

There are at most 2^L such partitions. If one part is empty, its unconstrained shore can be taken to be $V(H)$.

Proof. Take a minimum u–v outgoing cut X of capacity d. It loses e in J, so its capacity there is at most $d - 1$. Split the entire root set into $R \cap X$ and $R \setminus X$. The linkage guarantee implies

$$\varphi_H(u, R \setminus X) \le d - 1, \qquad \varphi_H(R \cap X, v) \le d - 1$$

whenever the respective root set is nonempty. Indeed, d edge-disjoint paths would specify an endpoint multiset of size $d \le k + 1$ which must also be linked in J, contradicting its cut capacity.

A minimum cut for the first aggregate flow gives Y_+, containing u and avoiding $R \setminus X$. The interval inequalities force it also to contain v and C_e. Complementing a minimum cut for the second aggregate flow gives Y_- containing the same triple and avoiding $R \cap X$. An empty aggregate root set uses $V(H)$ in its place.

Assign i to O if $|R_i \setminus X| \ge Q + 1$, and to I otherwise; the latter case has $|R_i \cap X| \ge Q + 1$. Unbreakability applied to Y_+ or Y_- proves the asserted smallness. The aggregation over the full union R, rather than separate linkage guarantees for each R_i, is essential. □

Fix one partition with nonempty O, I, and form the outgoing rank of Lemma 4.1 using the constraints in O, and the incoming rank using the constraints in I. Denote their singleton flats and cells by F_x, W_F and G_x, Z_G, respectively. For every arc in this partition, (4.3) gives low ranks for all its endpoints and witness vertices. Every minimizing outgoing shore containing u has capacity below d and therefore contains C_e and v. Likewise, every minimizing shore containing a member of C_e contains v. The incoming argument is dual. Thus the common witness flats satisfy

$$F_v \subseteq F_{C_e} \subseteq F_u, \qquad G_u \subseteq G_{C_e} \subseteq G_v. \tag{4.4}$$

All vertices of C_e have the same flats: their symmetric connectivity is at least $j \ge d$, while the minimizing cuts in question have capacity below d. In particular, $C_e \subseteq W_{F_{C_e}} \cap Z_{G_{C_e}}$.

4.3 A principal shore contained in every large low flat

Lemma 4.4 (Central large shore)**.** *Fix the outgoing rank for $O \neq \varnothing$ and an index $i \notin O$. If a feasible shore large on U_i, meaning $|Y \cap U_i| > Q$, has capacity less than k, let b_i be the minimum capacity among* all *such large feasible shores. Close one minimizer in the outgoing rank and call the result B_i. Then B_i is a uniquely determined feasible flat with*

$$c^+(B_i) = \rho_+(B_i) = b_i < k.$$

Every low outgoing flat F large on U_i contains B_i. The central shore B_i need not be a singleton flat.

Proof. Every feasible superset of a large minimizer is still large on U_i, so its rank cannot drop below b_i. Closing it preserves rank b_i and, by Lemma 4.1, gives an actual feasible shore of capacity b_i.

Let F be any low flat large on U_i. Since both F and B_i are low large shores, unbreakability gives $|U_i \setminus F|, |U_i \setminus B_i| \leq Q$. Hence

$$|F \cap B_i \cap U_i| \geq |U_i| - 2Q > Q.$$

The intersection is feasible. If its capacity is below k, the defining minimum b_i applies; otherwise its capacity is also at least b_i. Consequently submodularity gives

$$c^+(F \cup B_i) \leq c^+(F) + b_i - c^+(F \cap B_i) \leq c^+(F) < k.$$

For each OUT bag this union has at most $2Q$ terminals and more than Q outside. Unbreakability makes it feasible on every OUT bag. Thus $\rho_+(F \cup B_i) \leq \rho_+(F)$, and closedness forces $B_i \subseteq F$. Applying this conclusion with another closed minimizer in place of F proves uniqueness.

It is important that b_i was minimized over all large feasible shores, not merely the singleton flats: their intersection need not have a nonempty singleton cell. ☐

4.4 Counting in a fixed accounting region

We use three previously proved estimates, stated here with the precise orientations needed. Lemma 2.6 bounds arcs entering an outgoing b-shore, each individually critical for an ordered capped-k demand between vertices of a set W inside the shore, by $b(k|W| + 1)$. The one-shore estimate of Lemma 2.9 bounds arcs entering an outgoing b-shore, critical for demands wholly *outside* that shore, by b per outside *tail*. Both statements allow all other arcs and vertices of the ambient graph to remain present.

For the third estimate define

$$\mathcal{A}(k, b) = 2k(k+2)b + 4k(b+1), \qquad \mathcal{D}(k, b, N) = 4\mathcal{A}(k, b)\lceil \log_2(2N) \rceil + 8b. \tag{4.5}$$

Lemma 3.3 has the following consequence. If M is a union of SCCs after a g-arc deletion, at most t vertices are then deleted inside M, and $p \leq N$ charged vertices contain a witness pair and at least one endpoint of each selected M-internal critical arc, the number of such arcs is at most $\mathcal{D}(k, g+kt, N)p$, provided no selected endpoint plus *whole actual witness class* fits in one remaining SCC. For $p \leq 1$ there are no distinct-pair demands.

Lemma 4.5 (Mixed-pattern accounting)**.** *Let $V_0 \subseteq V(H)$ have $N \geq 2$ vertices. Select individually actual capped-k symmetric critical arcs with endpoints and entire chosen actual classes C_e in V_0.*

Let $U_1, \dots, U_L$ be the large bags above, and let T_{small} be a union of $h - L$ further bags each of size at most $3Q$, where $h \geq L$. Put

$$T = T_{\text{small}} \cup \bigcup_{i=1}^{L} U_i.$$

Suppose no selected triple $\{u, v\} \cup C_e$ lies in one SCC of $H - T$. Suppose also that all selected arcs have the same mixed root partition and satisfy (4.4). With

$$b_* = (h+1)(k-1) + 3kQh,$$

their number is at most

$$\Big[2(k-1)(k+1) + (k-1)L + (L+1)\mathcal{D}(k, b_*, N)\Big]N. \tag{4.6}$$

In particular this is $O(k^3 Q(h+1)^2 N \log(N+1))$, with an absolute constant.

Proof. Write $F = F_{C_e}$ and $G = G_{C_e}$. First count arcs with tail outside F by the external-entry estimate, charging the set $W_F \cap V_0$. These sets are disjoint, and every active group has a nonempty charged cell. Their total is at most $(k-1)(k+1)N$. Apply the reversed estimate to arcs with head outside G for the same cost. A survivor satisfies

$$u \in W_F \cap G, \qquad v \in F \cap Z_G, \qquad C_e \subsetneq W_F \cap Z_G. \tag{4.7}$$

For example $u \in F$ implies $F_u \subsetneq F$, whereas (4.4) gives the reverse inclusion.

Construct the central shore B_i of Lemma 4.4 for each $i \in I$ whenever a low feasible large shore exists. There are at most L distinct exceptional singleton flats $F = B_i$. A central shore which is not a singleton flat creates no exception.

Fix one exceptional F and group (4.7) by G. Use the region and charged set

$$M_{F,G} = F \cap G, \qquad P_{F,G} = Z_G \cap F \cap V_0.$$

The selected head and whole witness class belong to $P_{F,G}$. For fixed F these charged sets are disjoint, of total size at most N. Delete $\delta^+(F) \cup \delta^-(G)$, at most $2(k-1)$ arcs. A cycle crossing F needs a deleted outgoing arc; a cycle crossing G needs a deleted incoming arc. Thus $M_{F,G}$ is a union of residual SCCs. Its intersection with every large bag has size at most Q, using F for OUT bags and G for IN bags. Its intersection with T therefore has size at most $3Qh$. After deleting these local vertices every SCC meeting the region stays in it, avoids all of T, and is contained in an SCC of $H - T$. The original triple-separation assumption applies. Since the pattern is mixed, $L \geq 2$ and $2(k-1) + 3kQh \leq b_*$. The charged estimate bounds this exceptional-F family by $\mathcal{D}(k, b_*, N)N$. All exceptions cost at most L times this quantity.

Consider a nonexceptional F large on U_i. The principal shore satisfies $B_i \subsetneq F$. Closure monotonicity gives $W_F \cap B_i = \varnothing$, so the selected tail and the entire witness class in (4.7) are outside B_i. If its head belongs to B_i, the arc enters that shore and is critical for an outside demand. The one-shore estimate permits at most b_i such arcs per outside tail. All selected tails belong to V_0, so this costs at most $b_i N$ per central shore, and at most $(k-1)LN$ in total. Repeated counting can only increase this upper bound.

For every remaining nonexceptional flat put

$$J(F) = \{i \in I : |F \cap U_i| > Q\}, \quad M_F = F \setminus \bigcup_{i \in J(F)} B_i,$$

and delete the guard

$$Z_F = \delta^+(F) \cup \bigcup_{i \in J(F)} \delta^+(B_i).$$

All the central shores in this union are proper subflats of F. Hence $W_F \subseteq M_F$. After the preceding entry count, all surviving selected arcs of this group are internal to M_F; their tails and actual classes are in $W_F \cap V_0$.

The guard has at most $(L+1)(k-1)$ arcs and makes M_F a union of SCCs. A cycle leaving F requires a deleted exit of F; a cycle entering one of the removed B_i must subsequently leave it to return to M_F, using a deleted exit of B_i. This argument permits arbitrary overlap among the central shores. Every large bag is small on M_F: an OUT bag is small on F; an IN bag outside $J(F)$ is small on F; and for an IN bag in $J(F)$, the low large shore B_i leaves at most Q terminals outside itself. Thus $|T \cap M_F| \leq 3Qh$.

Every SCC meeting M_F after the guard and this local vertex deletion lies in an SCC of $H - T$, exactly as above. Apply the charged estimate with $W_F \cap V_0$. These sets are disjoint over all nonexceptional F and have total size at most N. The entire surviving family therefore costs at most $\mathcal{D}(k, b_*, N)N$. Combining the four contributions proves (4.6).

Each F uses its own region and guard only for an independent count of original critical arcs. No modified graphs are combined, and there is no additional enumeration of $J(F)$. All charged vertices belong to V_0, even when the low flats extend far outside it. ☐

Theorem 4.6 (Local counting with several unbreakable bags)**.** *Let $h, Q, k \geq 1$ be integers. Let V_0 have N vertices in an arbitrary unit-capacity directed graph H. Suppose a set E_0 of individually actual capped-k symmetric critical arcs has its endpoints and chosen entire actual witness classes in V_0. Let T be the union of h bags, each either of size at most $3Q$ or globally (Q,k)-unbreakable. Assume no selected endpoint-plus-class triple is contained in one SCC of $H - T$. Then, with natural logarithms,*

$$|E_0| \leq 2^{12}\, 2^h k^3 Q(h+1)^2 N \log(N+1). \tag{4.8}$$

The estimate is independent of every vertex outside V_0.

Proof. For $N \leq 1$ no selected distinct-pair demand exists. Suppose $N \geq 2$. Treat every bag of size at most $3Q$ as small, and let $L \leq h$ be the number of the remaining large bags. If $L = 0$, apply the charged-triple estimate directly to $W = V_0$, with no arc guard and at most $3Qh$ deleted vertices. The bound below covers this case.

Otherwise take $2Q+1$ roots in each large bag and let R be their union. Apply Lemma 2.4 with $K = k+1$, allowing deletion only from E_0 and freezing every other original arc. It retains at most

$$KN(2K|R|+1) \leq 26k^2Q(h+1)N$$

selected arcs and preserves the full two-sided endpoint-multiset linkages of size at most K. Count these retained selected arcs as the scaffold. Every other selected arc remains critical *in H*; Lemmas 4.2 and 4.3 assign it one of at most 2^L partitions. The scaffold is used to establish a cut inequality in H, not to change the graph for the subsequent counting.

A mixed partition is bounded by Lemma 4.5. If all large bags are OUT, use only the outgoing rank. Count arcs entering their gate flat by the external-entry estimate. Every survivor is internal to that flat and has its tail and actual class in $W_F \cap V_0$. The outgoing boundary has size at most $k-1$, and the flat contains at most $3Qh$ vertices of T. Apply the charged-triple estimate in each flat and sum the disjoint charged cells. This costs no more than the mixed bound. The all-IN case is the reversed argument.

Here is a coarse absolute calculation. Set $b_* = (h+1)(k-1) + 3kQh$. Since $h, Q, k \geq 1$,

$$3 \leq b_* \leq 4kQ(h+1), \qquad \mathcal{A}(k, b_*) \leq 12k^2 b_* \leq 48k^3 Q(h+1).$$

For $N \geq 2$, $\lceil \log_2(2N) \rceil \leq 4 \log(N+1)$ and $\log(N+1) > 1$. Therefore

$$\mathcal{D}(k, b_*, N) \leq 800k^3 Q(h+1) \log(N+1).$$

The other terms in (4.6) contribute at most $3k^3 Q(h+1)^2 N \log(N+1)$. Thus each pattern costs at most $803k^3 Q(h+1)^2 N \log(N+1)$, also covering the homogeneous and all-small cases. Add the scaffold, sum at most 2^h patterns, and enlarge the numerical coefficient to 2^{12}. This proves (4.8). □

Scope of the local theorem. The hypotheses concern the entire actual witness class at its lost symmetric level j, not an arbitrarily enlarged directed own-strength class. They also concern SCCs of the full graph after all the specified bags are deleted. No assertion about an induced region replaces these hypotheses. Once a hierarchy assigns each actual triple to a region satisfying them, the right side of (4.8) charges only that region's vertices; neither exterior paths nor repeated low-flat overlaps create an additional ambient-size factor.

5 A shallow hierarchy and global composition

All graphs in this section are finite directed multigraphs with unit arc capacities. For an arc set D, write $\deg_D(v)$ for the number of its incidences at v, counting both ends. For a nonnegative vertex weighting w, write $w(X) = \sum_{v \in X} w(v)$. A weighting is ϕ-expanding in a digraph J if

$$|\delta_J^+(X)| \geq \phi \min\{w(X), w(V(J) \setminus X)\} \qquad (X \subseteq V(J)).$$

Lemma 5.1 (Monotone weighted decomposition)**.** *Fix a digraph H on at most $n \geq 2$ vertices and a number $\phi > 0$. Suppose the vertex weights w and a forced arc set D increase monotonically. One can maintain a monotone arc set $B \supseteq D$ so that, after processing all current updates, every SCC Q of $H - B$ has ϕ-expanding weights in $(H - B)[Q]$. At every finite stage,*

$$|B \setminus D| \leq \phi \lceil \log_2 n \rceil \, w(V(H)). \tag{H1}$$

Proof. Add newly forced arcs to B. Whenever an SCC Q of $H - B$ has a proper shore $X \subset Q$ satisfying

$$|\delta^+_{(H-B)[Q]}(X)| < \phi \min\{w(X), w(Q \setminus X)\},$$

add that outgoing boundary to B. This strictly refines the SCC partition: after the deletion no path from X to $Q \setminus X$ can stay in Q, and a path leaving and returning to Q would contradict that Q was an SCC. A proper cut of a nontrivial SCC has positive capacity, so every such paid update adds an arc.

Represent each paid refinement by the binary partition $(X, Q \setminus X)$; represent the additional SCC refinements and forced refinements by free binary partitions. This gives a binary partition tree on the original vertices. Fix any finite stage and use its final weighting to evaluate the entire history. A paid refinement costs at most ϕ times the smaller final child weight, because weights only increased.

For any nonnegative weighting, the sum of the smaller child weights over a binary partition tree is at most $w(V(H))\log_2 n$. Indeed, with $0\log_2 0 = 0$, define

$$\Psi(A) = w(A)\log_2 w(A) - \sum_{v\in A} w(v)\log_2 w(v).$$

For a split $A = A_1 \dot\cup A_2$, the decrease of this potential is at least $\min\{w(A_1), w(A_2)\}$, by the binary entropy inequality. The initial potential is at most $w(V(H))\log_2 n$, and the final potential is nonnegative. The resulting bound on the total paid arcs proves (H1); any paid arc later put into D only decreases its left side. At quiescence there is no violating SCC cut, proving the expansion assertion. $\square$

Lemma 5.2 (Coupled feedback hierarchy)**.** *Let H have $n \geq 2$ vertices and $m \geq 1$ arcs. Set*

$$\ell = \lceil \log_2 n \rceil, \qquad 0 < r \leq \tfrac{1}{4}, \qquad \phi = \frac{r}{4\ell}, \qquad a = 2\phi\ell = \frac{r}{2}.$$

There exist nested original arc sets

$$E(H) = C_0 \supseteq C_1 \supseteq \cdots \supseteq C_{L+1} = \varnothing$$

such that

$$|C_i| \leq mr^i \tag{H2}$$

and, for every $i \geq 0$ and every SCC Q of $H - C_{i+1}$, the weighting $\deg_{C_i}$ restricted to Q is ϕ-expanding in $H[Q]$. One may take

$$L + 1 \leq 2 + \frac{\log m}{\log(1/r)}. \tag{H3}$$

Proof. Use a copy of Lemma 5.1 for every integer $i \geq 0$. In copy i, the terminal weighting is $w_i = \deg_{C_i}$, the forced set is $D_i = C_{i+2}$, and its output is C_{i+1}. Start with $C_0 = E(H)$ and all other sets empty. Propagate weight changes and newly forced arcs to adjacent copies, and process violating cuts, one arc-adding update at a time. At any finite stage only finitely many indices are nonempty. Processing forced updates eventually ensures $C_{i+2} \subseteq C_{i+1}$; $C_1 \subseteq C_0$ holds automatically.

Put $x_i = |C_i|$. The historical cost bound gives, at every finite stage,

$$x_{i+1} \leq x_{i+2} + a x_i. \tag{H4}$$

Pending forced additions do not invalidate this inequality: an addition to C_{i+2} increases its right side, and when the forced arc is copied to C_{i+1} it is already covered by that term. Paid updates satisfy the historical bound of their own copy. Thus the inequalities hold without requiring all pending nesting updates to have been processed at the observation time.

We claim $x_i \leq mr^i$ at every finite stage. Otherwise the finite-support sequence x_i/r^i has a maximum $M > m$ at an index $j \geq 1$. Applying (H4) at $j-1$ gives

$$x_j \leq x_{j+1} + a x_{j-1} \leq Mr^j \left(r + \frac{a}{r}\right) \leq \frac{3}{4} M r^j,$$

contradicting $x_j = Mr^j$.

Since the x_i are integers, every index with $mr^i < 1$ is permanently empty. There are therefore only finitely many possible nonempty indices, and only finitely many possible arc additions. Process the finite pending updates and violations until quiescence. The forced rules give nesting, and every residual SCC has expansion in $(H - C_{i+1})[Q]$, hence also in $H[Q]$. This proves (H2); choosing the first guaranteed empty index gives the deliberately loose bound (H3). $\square$

Lemma 5.3 (Unbreakable bags with explicit depth)**.** *Let $k \geq 1$, $n \geq 2$, and $0 < \epsilon \leq 1$. Suppose H has at most $kn(n-1)$ arcs and $k < \sqrt{n}$. There is a hierarchy as in Lemma 5.2 with at most*

$$h_{\max} \leq \frac{6}{\epsilon} \tag{H5}$$

levels of regions. A level-i region Q is an SCC of $H - C_{i+1}$, its children are the SCCs of $H - C_i$ it contains, and its bag is

$$U_{i,Q} = \{v \in Q : v \text{ is incident to an arc of } C_i\}.$$

Every bag is globally (q,k)-unbreakable for the integer

$$q = \lceil 16k \lceil \log_2 n \rceil n^{\epsilon} \rceil \leq 49kn^{\epsilon} \log(n+1). \tag{H6}$$

At each level the regions partition the vertices, so the sum of all region cardinalities is at most $h_{\max} n$.

Proof. If H has no arcs, take one level of singleton regions with empty bags; all assertions are immediate. Otherwise use $r = \frac{1}{4}n^{-\epsilon}$ in Lemma 5.2; then $\phi = n^{-\epsilon}/(16\ell)$. Since $m < n^{5/2}$ and $\log(1/r) = \log 4 + \epsilon \log n$, the bound (H3) implies

$$L + 1 \leq 2 + \frac{5}{2\epsilon} \leq \frac{6}{\epsilon}.$$

Every member of $U_{i,Q}$ has terminal weight at least one. An outgoing cut of $H[Q]$ with at most k arcs has a side of terminal weight at most k/ϕ, and hence a side with at most q bag vertices. Thus the bag is unbreakable in $H[Q]$. The restriction of any global outgoing cut to Q has no greater capacity, so the bag is also globally unbreakable in H. Finally $\ell \leq 3\log(n+1)$, and $1 \leq kn^{\epsilon}\log(n+1)$, giving the coarse bound in (H6). The partition-size assertion follows from the SCC partitions at each level. □

Lemma 5.4 (Ancestor bags give full-graph triple separation)**.** *In the preceding hierarchy, let Z be a set of at least two vertices contained in one original SCC. Assign Z to a deepest region Q containing it. Let T_Q be the union of Q's bag and every bag on its ancestor path. No SCC of the full graph $H - T_Q$ contains Z.*

Proof. Before processing a path region A, every full-graph SCC after the preceding path-bag deletions that meets A is contained in A. At the root this follows from the original SCC partition. Deleting $U_{i,A}$ cannot merge SCCs. Every internal C_i arc of A has an endpoint in that bag, so any resulting SCC meeting A is a strongly connected subgraph of $H - C_i$. It therefore lies in a global child SCC. This proves the induction down the path. In particular it accounts for all outside excursions: a return excursion through an exterior vertex would put that vertex in the same full SCC, contradicting containment in A.

If Z meets a deleted bag, the conclusion is immediate. Otherwise a residual SCC containing Z would lie in one child of Q, contradicting the deepest-region assignment. Such an assignment always exists above the bottom: $H - C_0$ has singleton SCCs and $|Z| \geq 2$. Repeated regions and overlapping bags cause no problem, since deletions cannot merge SCCs. □

Theorem 5.5 (A uniform connectivity preserver)**.** *Theorem 4.6 gives the following local estimate: in a strongly connected digraph, selected individually critical arcs with both endpoints and their whole assigned actual symmetric witness classes in a set W of N vertices number at most*

$$2^{12} 2^h k^3 q (h+1)^2 N \log(N+1), \qquad q \geq 1, \tag{H7}$$

provided their endpoint-plus-class triples are separated by deleting the union of h globally (q, k)-unbreakable bags. Consequently every n-vertex digraph has a spanning subgraph preserving all capped symmetric k-connectivities with at most

$$2^{40} k^4 n \exp\bigl(16\sqrt{\log(n+1)}\bigr) \tag{H8}$$

arcs. In fact the calculation below gives 2^{26} and 8 in place of 2^{40} and 16.

Proof. Take an inclusion-minimal spanning subgraph H preserving the original capped symmetric values. Each remaining arc is critical for some actual pair at a level $j \le k$; fix the pair and its whole actual symmetric j-class throughout the counting. The arc and the class lie in one original SCC. Indeed, the arc occurs on a witness path, and a reverse witness path supplies return reachability. Arcs between SCCs and loops are therefore absent. Parallel multiplicity is at most k, since an arc from a larger bundle cannot be critical at any level at most k.

For $k \ge \sqrt{n}$, keeping these arcs already costs at most $kn(n-1) \le nk^3$, which satisfies (H8). For $k < \sqrt{n}$, use Lemma 5.3 on the unchanged graph H, componentwise with the same global parameters. Assign each critical arc to the deepest region containing its endpoints and its whole chosen class. Lemma 5.4 supplies exactly the full-graph SCC hypothesis of the local estimate (H7).

Apply (H7) to a region with the active set equal to that region. All witnesses, cuts, and criticalities remain those of the same ambient H; independently pruned local graphs are never combined. The total region cardinality and the bound $h \le h_{\max}$ give

$$|E(H)| \le 2^{12} 2^{h_{\max}} k^3 q (h_{\max}+1)^2 h_{\max} n \log(n+1).$$

Set $\epsilon = 1/\sqrt{\log(n+1)}$. By (H5) and (H6),

$$h_{\max} \le 6\sqrt{\log(n+1)}, \qquad q \le 49k\, e^{\sqrt{\log(n+1)}} \log(n+1).$$

Write $t = \log(n+1) > 1$. Then $h_{\max}(h_{\max}+1)^2 \le 294 t^{3/2}$, so the displayed count is at most

$$2^{12} \cdot 49 \cdot 294\, k^4 n\, t^{7/2} \exp\bigl((1 + 6\log 2)\sqrt{t}\bigr).$$

Here $2^{12} \cdot 49 \cdot 294 = 59{,}006{,}976 < 2^{26}$. The elementary inequality $\log t \le 2\sqrt{t}/e$, for $t > 0$, shows $t^{7/2} \le \exp((7/e)\sqrt{t})$, and

$$1 + 6\log 2 + 7/e < 8.$$

This proves the stronger bound with constants 2^{26} and 8, and in particular (H8). The cases $n \le 1$ require no arcs. The exponent of k is fixed and the factor involving n alone is $n^{o(1)}$. $\square$

References

[1] Gary Hoppenworth, Thatchaphol Saranurak, and Benyu Wang. *Near-Optimal Fault-Tolerant Strong Connectivity Preservers.* FOCS 2025; arXiv:2510.02562, 2025. `https://arxiv.org/abs/2510.02562`.

[2] Shivam Bansal, Keerti Choudhary, Harkirat Dhanoa, and Harsh Wardhan. *Fault-Tolerant Bounded Flow Preservers.* arXiv:2404.16217, 2024. `https://arxiv.org/abs/2404.16217`.