Exact-Mean Decompositions for Cost-Preserving Single-Source Unsplittable Flows: Demand Scales and Local Interaction Structure

Sergey Nikolenko

August 13, 2026

Abstract

For a feasible fractional single-source flow $\mathbf{x}$ with demands $0 < d_i \leq D$ and arc costs $\mathbf{c} \geq \mathbf{0}$, we study unsplittable routings $\mathbf{y}$ that are simultaneously cost-preserving, $\mathbf{c}^\mathsf{T}\mathbf{y} \leq \mathbf{c}^\mathsf{T}\mathbf{x}$, and additively congestion-good, $\mathbf{y} \leq \mathbf{x} + CD\mathbf{1}$. We present an *exact-mean decomposition*, which is a probability law on unsplittable routings whose mean is $\mathbf{x}$ and whose atoms stay in a prescribed box. Its mean preserves every linear cost automatically. Conversely, on every graph class closed under deleting arcs, separation shows that the optimal cost-preserving constant is exactly the least radius of such a law. Cost therefore disappears from the structural problem.

We build these laws in two orthogonal ways. Demand arithmetic partitions the terminals into divisibility chains; independent products give the minimum chain-cover charge $\kappa_{\mathrm{div}}(\mathcal{V})$, computable by one maximum-weight matching. This yields, among other consequences, constants below 2 for two demand values and below $2m$ for at most m values per dyadic band. Interaction geometry instead factors exact-two-path systems through signed rows. Factor forests and arbitrary pair-row systems have sharp radius one. The first irreducible local cores have exact envelopes $9/8$ and $(299 - 41\sqrt{41})/32$ at arities three and four; a common bad-state/row-cover duality proves both. Globally, path colouring gives radius $2D$ for row arity at most three and an explicit bound at every fixed arity.

Finally, dyadic composition reduces the unrestricted universal-constant question to demands lying in one factor-two window, with only a factor-two loss in either direction. Thus a finite one-band constant is equivalent to a finite universal constant; the standing record implies $K^* \geq 2.5652$. We also prove a sharp constant one for two-exit outerplanar interval spines and raise the planar lower bound to 1.1735. All finite witnesses are checked in exact arithmetic, while the upper bound theorems do not rely on enumeration or solvers.

1 Introduction

Problem setting and current results. An instance of the *single-source unsplittable flow* problem consists of a digraph $G = (V, A)$, a source s, terminals $t_1, \ldots, t_k$ with demands $d_i > 0$, and nonnegative arc costs $\mathbf{c}$. Given a feasible fractional flow $\mathbf{x}$, we look for an *unsplittable* flow $\mathbf{y}$, i.e., a flow with one path per terminal carrying its whole demand, that does not exceed the fractional loads by too much. Dinitz, Garg, and Goemans [3] proved that fractional flows without costs can be routed in this way while incurring an additional maximal demand value: with no condition on the costs, there always exists $\mathbf{y}$ such that $\mathbf{y} \leq \mathbf{x} + D\mathbf{1}$, where $D = \max_i d_i$.

Goemans conjectured that cost preservation, $\mathbf{c}^\mathsf{T}\mathbf{y} \leq \mathbf{c}^\mathsf{T}\mathbf{x}$, comes for free with the same constant, and this conjecture has now been disproved: first Rybin [14] provided a seven-vertex instance with critical constant $16/15$, then it was improved to $\frac{299-41\sqrt{41}}{32} = 1.139747\ldots$ by Protti [13] and our previous work [11], and then our follow-up work [10] provided the current record, a seventeen-terminal common-point instance with an exact rational certificate that forces $C \geq 1.28249479798\ldots$; this example was later refined to $1.2826\ldots$ by de la Fuente [2], and this is the current standing record lower bound; its seventeen demands lie inside a single dyadic band.

Morell and Skutella's two-sided conjecture [9], together with the cost transformation of [17], would imply that the cost-preserving constant 2 suffices. Strong structured cases were already known: the cost-preserving constant 2 for planar digraphs [18], and an exact-mean two-sided radius-1 decomposition (hence cost-preserving constant 1) for series-parallel digraphs [1], in addition to the exact decomposition theorem for divisibility-chain demands of [9]. Before the present work, no finite constant was known for unrestricted digraphs and unrestricted demand-value structure.

The unifying idea of this work. In this work, we use the following principle throughout all upper bounds: *preserve the mean, then factor the error.* A fractional flow $\mathbf{x}$ is cost-preservingly routable once it is the mean of unsplittable routings in a small box around $\mathbf{x}$, because $\mathbb{E}[\mathbf{c}^\mathsf{T}\mathbf{Y}] = \mathbf{c}^\mathsf{T}\mathbf{x}$. The converse separation argument also holds: for restriction-closed classes this sufficient condition is an exact characterization. The cost vector is a witness that the good-routing hull misses $\mathbf{x}$, and nothing more.

With this basic idea, all our proofs use four techniques.

1. *Build a local law.* Integrality (see below) does this for equal demands, divisibility chains, and pair rows; probabilistic coupling does it for one row and for factor forests.
2. *Split until a local law applies.* Demand values split into divisibility chains, while path colouring splits bounded-arity interaction systems into pair-row blocks.
3. *Compose.* Obtain the block laws independently, then use the fact that exact means and the radii are additive, thus avoiding the costs and gluing of terminal paths.
4. *Expose an irreducible core.* When no known split applies, a bad Boolean state provides a row witnessing its excess. A matching, fractional cover, or Helly kernel then asks whether all the witness rows can be bad simultaneously. This *bad-state/cover duality* computes the local ternary and quaternary envelopes.

These techniques explain why two theorem families that look unrelated at first glance actually have the same form: arithmetic factorization charges a divisibility chain once, at its largest demand, and geometric factorization also charges a compatible interaction block once, at its largest demand. The main arithmetic obstruction is therefore many incomparable values in a single $[n, 2n]$ window; the main geometric obstructions are cyclic high-arity cores. The triangle and four-interval constants from our previous work [11, 10] reappear here because they are the first two local cores that survive all radius-one factorizations.

Figure 1 illustrates this proof structure. Note that the numerical certificates serve a much narrower role: they discover or verify finite boundary configurations. Claims about specific DAGs are rebuilt from arc lists with all simple paths rediscovered, and each such computation has a separately written implementation. The class and local-envelope theorems themselves do not rely on exhaustive programmatic enumeration or solvers.

Our contributions. In this work, we develop the exact-mean factorization principle along its two natural axes: arithmetic of demand values and geometry of support interactions. Our main results are as follows.

1. **Arithmetic factorization.** The block-composition and chain-partition theorems (Theorems 4.1–4.2) turn any cover by divisibility chains into an exact-mean two-sided decomposition. The best cover value $\kappa_{\text{div}}(\mathcal{V})$ is computed by one maximum-weight matching

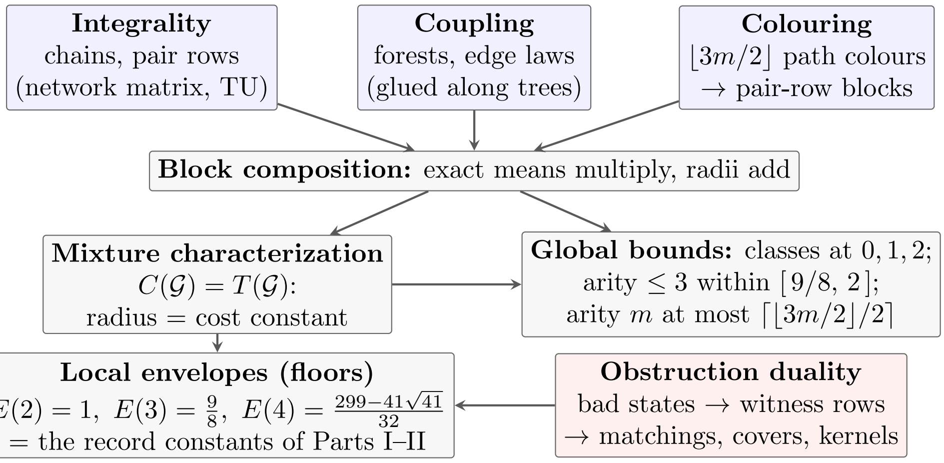

Figure 1: The proof grammar. Three mechanisms produce exact-mean block laws; composition adds their radii; the mixture characterization converts radii into cost-preserving constants in both directions. When these factorizations stop, bad-state/cover duality computes the surviving local core.

(Proposition 4.5), and we derive from this result the scale-sum, two-value, and bounded-band-density bounds.

2. **The one-band reduction.** Dyadic blocks reduce the unrestricted problem to the factor-two statement $\mathrm{FT}(K)$; separation gives the converse within a factor of two (Theorems 7.2–7.4). Therefore, a finite one-band constant exists if and only if a finite universal constant does. The seventeen-terminal record gives $K^* \geq 2.5652$. Rounding up fails because it loses terminal labels (Proposition 5.1); rounding down remains valid and yields the all-m shifted-grid family of Theorem 6.3.

3. **Interaction factorization.** The mixture characterization (Theorem 8.1) identifies class constants with exact-mean hull radii. Forest couplings round compatible local row laws at radius one, while total unimodularity rounds every exact-two-path pair-row system at radius one without a forest hypothesis (Theorems 8.7 and 8.11). This also settles the case of outerplanar two-exit interval spines with constant one.

4. **Irreducible local cores and global assembly.** The unique maximal ternary core has exact envelope $9/8$ (Theorem 8.14). The five realizable four-column forms have envelopes $9/8$ except for the interval form, whose value is exactly $(299 - 41\sqrt{41})/32$ (Theorem 9.1). Proper path colouring then partitions every global arity-three two-path system into two pair-row blocks, giving radius $2D$ (Theorem 8.18). Both local envelope proofs follow the same primal–dual obstruction scheme: good states build a marginal mixture, while bad states select rows whose cheap cover contradicts the active facet equation.

5. **Exact witnesses and audits.** We present a six-terminal witness that raises the *planar* lower bound to 1.1735 (Proposition 8.3), checking the result in exact arithmetic in its native representation.

Note that every counterexample instance of our previous work [11, 10] used at most nine distinct demand values, so the entire experimental record of the lower bound program is now positioned under a proven finite instance-dependent constant. The core of these examples, as it now turns out, lies in the non-divisibility of demands.

In particular, we show how divisibility chains lead to upper bounds additively. Values that divide one another can share one decomposition and are charged only once; e.g., demands

Demand values	Bound on C	Where
one value, or one divisibility chain	1	Cor. 4.4(1)
two values $D > d$	$1 + d/D < 2$	Cor. 4.4(1)
successive gaps of factor ≥ 2	< 2	Cor. 4.4(2)
$\leq m$ values per dyadic band	$< 2m$	Cor. 4.4(3)
arbitrary value set $\mathcal{V}$	$\kappa_{\mathrm{div}}(\mathcal{V}) \leq R/D$	Prop. 4.5
bounded loads $\mathbf{x} \leq LD\mathbf{1}$	$\sqrt{2L} + 2$	Cor. 6.4
unrestricted	open, equiv. to $K^* < \infty$	Thms. 7.2, 7.3

Table 1: Our results in terms of the structure of demand values. All bounds are unconditional in the graph, the costs, and the loads (except the bounded-load row, which is unconditional in the demand values instead).

Block type	Why one exact-mean law exists	Radius paid
divisibility chain	Morell–Skutella chain decomposition	$\max d_i$
factor-incidence forest	local row laws glued along singleton marginals	$\max d_i$
pair-row two-path block	one forbidden corner per row plus total unimodularity	$\max d_i$
ternary two-path block	four path colours merged into two pair-row blocks	$2 \max d_i$

Table 2: Factorization results: different mechanisms lead to the same result interface: an exact-mean block law and a radius. Theorem 4.1 only adds the radii.

$\{6, 3, 2\}$ have $D = 6$, and if the scale-sum bound charged every value once, we would get $C \leq (6 + 3 + 2)/6 = 11/6$, but since $\{6, 3\}$ forms a divisibility chain charged only at its maximum, the cover $\{6, 3\}, \{2\}$ costs $(6 + 2)/6$, and $C \leq 4/3$. Note that adding the value 1 does not change the upper bound (Proposition 4.5 finds the optimal cover by a maximum-weight matching). By contrast, for pairwise incomparable values such as $\{6, 5, 4\}$ no cover beats the plain sum, $C \leq 15/6$. This distinction always appears in the lower bounds: every record of [11, 10] uses pairwise non-divisible demands packed into a narrow band; we now back up this structure with theory.

The geometric counterpart of "charge a chain once" is "round a compatible interaction block once." A whole forest of signed row constraints costs one maximal demand, as does an arbitrarily cyclic system where every row meets at most two terminal choices. The first extra charge appears only at arity three, so the triangle, rather than a pairwise odd cycle, is the first core structure for a lower bound.

For unrestricted graphs and demands, the arithmetic reduction leaves one frontier: many incomparable values inside a single factor-two window. Table 1 summarizes that landscape. The *band constant* K^*, which we define as the best possible additive merging error for one window, in units of its lower endpoint, satisfies $C_{\mathrm{opt}} \leq K^* \leq 2C_{\mathrm{opt}}$ (Theorems 7.2 and 7.3), so a finite band constant is equivalent to a universal constant. The $k = 17$ record certifies $K^* \geq 2C_{17} \approx 2.565$ (Proposition 7.5), and its complete relaxed frontier is plotted in Figure 8.

AI use. Like the previous parts [11, 10], this work was written in close collaboration with GPT 5.6 Sol, Claude Fable 5, and Claude Opus 5, which contributed proofs, failed routes, adversarial reviews, code for the verification campaign, and more. This time, I cannot claim that it was a single tour de force by the models like Rybin's original counterexample [14]; it was a long journey with a lot of human involvement, but the key ideas were mostly provided by the LLMs. I have personally verified and edited this work in its entirety, and all errors are mine.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 fixes conventions and states the two decomposition primitives with full proofs. Section 4 proves the product theorem and its corollaries. Section 5 documents the round-up wall, Section 6 proves the round-down family, and Section 7 states the merger reductions, their factor-two converse, and the certified one-band lower bounds. Section 8 characterizes restriction-closed graph classes by their exact-mean mixture radius, develops forest and low-arity factorizations, and derives bounds for various classes. Section 9 closes the local arity-four envelope. Section 10 describes our exact verification pipeline, and Section 11 collects open problems and concludes.

2 Related work

Congestion without costs. Dinitz, Garg, and Goemans [3] proved the additive-D theorem for congestion; their augmenting-path proof is the template for most later positive results. Morell and Skutella [9] study arc-wise two-sided bounds: they prove the simultaneous multiplicative bounds $\mathbf{x}/2 - D\mathbf{1} \leq \mathbf{y} \leq 2\mathbf{x} + D\mathbf{1}$ for arbitrary demands, the exact convex decomposition theorem for demands forming a divisibility chain (the primitive quoted as Theorem 3.3 below), and they conjecture the additive two-sided analogue with error below D. Liu and Reis [6] recently confirmed a common-order case of the Morell–Skutella conjecture with the constant $(1 - \frac{1}{2m-2})D$ via weighted chairman assignment.

Costs. Kolliopoulos and Stein [5] introduced grouping demands into geometric classes; their round-up variant loses a constant factor in cost. Skutella [16, Theorems 3 and 4] obtained the cost-*preserving* bounds $\mathbf{y} \leq 2\mathbf{x} + D\mathbf{1}$ and $\mathbf{y} \leq \sqrt{2}\,\mathbf{x} + (1 + 1/\sqrt{2})D\mathbf{1}$ by rounding demands down; Theorem 6.3 is the all-m continuation of exactly these two constructions, and our stripping lemma is his cost argument stated in the form we need. Martens, Salazar, and Skutella [7] study convex combinations of unsplittable flows with *rounded* demands whose average recovers the original flow; the product theorem below differs in that no demand is ever modified and the mean is exact at the original demands. Swamy, Traub, Vargas Koch, and Zenklusen [17] give a general transformation producing cost guarantees from unweighted error-bounded variants; none of our theorems requires it, because exact product means and the stripping inequality supply the cost side directly.

Graph-structured positive results. Traub, Vargas Koch, and Zenklusen [18] prove the cost-preserving additive constant 2 for planar digraphs, along with simultaneous upper and lower bounds. Almoghrabi, Skutella, and Warode [1] prove the exact Morell–Skutella convex-decomposition conjecture for series-parallel digraphs, even for multiflows. Our bounds are orthogonal: the graph is unrestricted, while the parameter is the arithmetic structure of the demand values.

Lower bounds. Goemans' conjecture—cost preservation at the DGG constant 1—was disproved by the seven-vertex 16/15 instance [14]; then it was improved to $\frac{299-41\sqrt{41}}{32} = 1.139747\ldots$ [13, 11], and then our follow-up work [10] provided the current record, a seventeen-terminal common-point instance with an exact rational certificate that forces $C \geq 1.28249479798\ldots$, later refined to $1.2826\ldots$ [2]. All these instances have pairwise non-divisible demand values in a narrow band. The results of this work explain one part of that pattern: divisibility-chain structure gives exact-mean mixtures at radius one, while a factor-two band contains no nontrivial divisibility relations. They do not claim that narrow bands or pairwise non-divisibility are necessary for lower bounds.

3 Preliminaries

3.1 Setting

$G = (V, A)$ is a directed acyclic graph with source s and terminals $t_1, \ldots, t_k$, regarded as private sinks (adding a private zero-cost sink arc per terminal changes nothing). Terminal i has demand $d_i > 0$; $D := \max_i d_i$. A *feasible flow* $\mathbf{p} \in \mathbb{R}^A_{\geq 0}$ has divergence d_i at t_i, $-\sum_i d_i$ at s, and 0 elsewhere. Every feasible single-source flow decomposes into per-terminal path flows $\mathbf{p} = \sum_i \mathbf{p}_i$, with $\mathbf{p}_i$ a d_i-flow from s to t_i; we fix such a decomposition when needed. An *unsplittable routing* $\mathbf{y}$ chooses one $s \to t_i$ path P_i per terminal and has load $y_a = \sum_{i:a \in P_i} d_i$; it is a *support routing of* $\mathbf{p}$ if every P_i lies in $\operatorname{supp}(\mathbf{p})$. Costs $\mathbf{c} \in \mathbb{R}^A_{\geq 0}$ act linearly: $\mathbf{c}^\mathsf{T}\mathbf{y} = \sum_i d_i\, c(P_i)$.

Write C_{opt} for the infimum of the constants C that work simultaneously for every instance and every nonnegative cost vector. We use C without a subscript only for a displayed valid bound or for the critical constant of a fixed instance; this distinction matters in Section 7.

Acyclicity comes without loss of generality for our purposes: cancelling directed cycles in a fractional flow only lowers loads and (nonnegative) cost, and support routings of the cancelled flow are routings of the original instance.

Definition 3.1 (exact-mean decomposition)**.** *An* exact-mean decomposition of radius ρ *for a feasible flow* $\mathbf{p}$ *is a finite probability law on support routings* $\mathbf{Y}$ *such that*

$$\mathbb{E}\mathbf{Y} = \mathbf{p}, \qquad \mathbf{Y} \leq \mathbf{p} + \rho\mathbf{1} \quad \textit{almost surely.}$$

It is two-sided *if also* $\mathbf{Y} \geq \mathbf{p} - \rho\mathbf{1}$ *almost surely. All distributions below are finitely supported, so "almost surely" is just an abbreviation for "on every atom of positive weight".*

Exact-mean decomposition leads to a conclusion about the costs for unsplittable flows: since for every arc cost vector $\mathbf{c}$, $\mathbb{E}[\mathbf{c}^\mathsf{T}\mathbf{Y}] = \mathbf{c}^\mathsf{T}\mathbf{p}$, there exists an atom (support routing) that costs no more than $\mathbf{p}$.

The rest of the paper is mostly devoted to two tasks made distinct by this definition: construct small-radius decompositions for elementary blocks, and then compose those blocks without changing their mean.

3.2 Two decomposition primitives

The first primitive is elementary flow integrality.

Lemma 3.2 (equal-demand convex decomposition)**.** *Suppose that every demand of a subinstance equals d, and let* $\mathbf{p}$ *be a feasible flow for that subinstance. Then* $\mathbf{p}$ *has a two-sided exact-mean decomposition of radius d, and, moreover, $|Y_a - p_a| < d$ (strictly) whenever $p_a/d \notin \mathbb{Z}$.*

Proof. Delete the arcs with $p_a = 0$; this can only shrink the set of available paths, so it suffices to prove the claim on $\operatorname{supp}(\mathbf{p})$. After scaling by d, the flow $\mathbf{g} = \mathbf{p}/d$ has integral divergences ($-n$ at s for n terminals, $+1$ at each terminal, 0 elsewhere). The polytope

$$\mathcal{Q} = \{\mathbf{h} :\ B\mathbf{h} = B\mathbf{g},\ \lfloor \mathbf{g} \rfloor \leq \mathbf{h} \leq \lceil \mathbf{g} \rceil\},$$

where B is the node-arc incidence matrix, has a totally unimodular constraint matrix and an integral right-hand side, hence is an integral polytope; it contains $\mathbf{g}$, so $\mathbf{g}$ is a convex combination of its integral points. An integral point is an acyclic integral flow with unit divergence at every terminal and therefore decomposes into exactly one unit $s \to t_i$ path per terminal, i.e., a routing. The difference between the floor and the ceiling gives $|d\mathbf{h} - \mathbf{p}| \leq d\mathbf{1}$, with strict inequality at non-integral coordinates, and deleted arcs carry zero in both $\mathbf{p}$ and $d\mathbf{h}$. $\square$

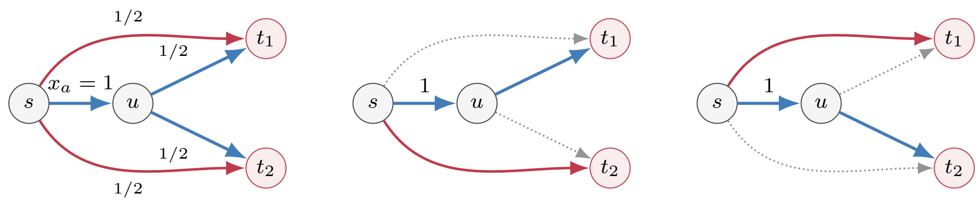

(a) fractional flow
each terminal is split 1/2–1/2

(b) atom 10
t_1 common, t_2 bypass

(c) atom 01
t_1 bypass, t_2 common

Figure 2: Lemma 3.2 on two unit-demand terminals: two atoms mixed with probability 1/2 each, common arc has max load one exactly, every terminal keeps its prescribed marginal, and each private choice arc changes by only $1/2 < D$. Pale dotted arcs belong to the support but are not selected in that atom.

For an illustration, let two unit-demand terminals each have one route through a common arc a and one private bypass, and let each use its common-arc route with share 1/2, as shown in Fig. 2. The fractional load of a is one. Mixing the two routings $(z_1, z_2) = (1, 0)$ and $(0, 1)$ with probability 1/2 preserves that load exactly and gives each terminal its prescribed marginal. Every private choice arc moves from its fractional load 1/2 to either zero or one, an error of $1/2 < D$. This is the smallest visible instance of the floor–ceiling box in Lemma 3.2: integrality couples the two terminal choices instead of rounding them independently.

The second primitive is the divisibility-chain decomposition of Morell and Skutella. A finite demand multiset forms a *divisibility chain* if its distinct values can be ordered $e_1 > e_2 > \cdots > e_h$ with $e_{j+1} \mid e_j$ (each distinct value divides the previous one; repeated values in the multiset itself are allowed).

Theorem 3.3 (Morell–Skutella chain decomposition [9]). *Let $\mathbf{p}$ be a feasible flow for a subinstance whose demand multiset forms a divisibility chain with largest value e_1. Then $\mathbf{p}$ admits a two-sided exact-mean decomposition of radius e_1 into unsplittable routings of that subinstance.*

Lemma 3.4 (support-preserving form). *In Theorem 3.3 the routings can be taken to be support routings of $\mathbf{p}$.*

Proof. Apply Theorem 3.3 to the instance restricted to the subgraph $\mathrm{supp}(\mathbf{p})$, on which $\mathbf{p}$ is still a feasible flow. Every routing produced lives in that subgraph, hence is a support routing; the load bounds are unchanged on $\mathrm{supp}(\mathbf{p})$ and vanish on deleted arcs. □

Lemma 3.2 is the one-value case of Theorem 3.3; we proved it separately to keep the equal-value applications of this work self-contained.

3.3 The box transfer

The following standard lemma converts exact-mean mixtures into cost-preserving routings; we state it with an explicit box width ρ because the relaxed merger reduction of Section 7 consumes the slack.

Lemma 3.5 (box transfer). *Let $\mathbf{x}$ be feasible, $\mathbf{c} \geq \mathbf{0}$, $\rho \geq 0$, and let $\mathbf{q}$ minimize $\mathbf{c}^\mathsf{T}\mathbf{p}$ over feasible flows with $\mathbf{p} \leq \mathbf{x} + \rho D \mathbf{1}$. Then there exists a vector $\boldsymbol{\lambda} \geq \mathbf{0}$ such that:*

(1) $\mathbf{q}$ *minimizes $(\mathbf{c}+\boldsymbol{\lambda})^\mathsf{T}\mathbf{p}$ over all feasible flows, and every support path of $\mathbf{q}$ is $(\mathbf{c}+\boldsymbol{\lambda})$-shortest for its terminal;*

(2) $\mathbf{c}^\mathsf{T}\mathbf{q} \leq \mathbf{c}^\mathsf{T}\mathbf{x} - \rho D \|\boldsymbol{\lambda}\|_1$*;*

(3) for every random support routing $\mathbf{Y}$ *of* $\mathbf{q}$*,*

$$\mathbb{E}\big[\mathbf{c}^\mathsf{T}\mathbf{Y}\big] = \mathbf{c}^\mathsf{T}\mathbf{q} + \boldsymbol{\lambda}^\mathsf{T}\big(\mathbf{q} - \mathbb{E}[\mathbf{Y}]\big).$$

Proof. Take $\boldsymbol{\lambda}$ to be optimal dual multipliers of the box constraints $\mathbf{p} \le \mathbf{x} + \rho D\mathbf{1}$ in the linear program defining $\mathbf{q}$; they exist by LP duality. Complementary slackness makes $\mathbf{q}$ optimal for the Lagrangian objective $(\mathbf{c} + \boldsymbol{\lambda})^\mathsf{T}\mathbf{p}$ over the unconstrained flow polyhedron, which is statement (1). In the dual of this min-cost flow, every positive-flow arc is tight. Therefore every $s \to t_i$ path contained in $\mathrm{supp}(\mathbf{q})$, including a path obtained by splicing support arcs from different members of a decomposition, has the same potential-telescoped length and is shortest for terminal i.

For (2), strong duality gives

$$\mathbf{c}^\mathsf{T}\mathbf{q} = \min_{\mathbf{p}\text{ feasible}} (\mathbf{c} + \boldsymbol{\lambda})^\mathsf{T}\mathbf{p} - \boldsymbol{\lambda}^\mathsf{T}(\mathbf{x} + \rho D\mathbf{1}) \le (\mathbf{c} + \boldsymbol{\lambda})^\mathsf{T}\mathbf{x} - \boldsymbol{\lambda}^\mathsf{T}\mathbf{x} - \rho D\|\boldsymbol{\lambda}\|_1 = \mathbf{c}^\mathsf{T}\mathbf{x} - \rho D\|\boldsymbol{\lambda}\|_1.$$

For (3), every support routing has

$$(\mathbf{c} + \boldsymbol{\lambda})^\mathsf{T}\mathbf{Y} = \sum_i d_i \,\mathrm{dist}_{\mathbf{c}+\boldsymbol{\lambda}}(t_i) = (\mathbf{c} + \boldsymbol{\lambda})^\mathsf{T}\mathbf{q}$$

by (1); then we subtract $\boldsymbol{\lambda}^\mathsf{T}\mathbf{Y}$ and take expectations. □

At $\rho = 0$ the lemma says that if a mixture of support routings of $\mathbf{q}$ has mean exactly $\mathbf{q}$, its mean cost is exactly $\mathbf{c}^\mathsf{T}\mathbf{q} \le \mathbf{c}^\mathsf{T}\mathbf{x}$, so some atom is cost-preserving. For mixtures built at $\mathbf{x}$ itself even this is not needed: the mean cost is $\mathbf{c}^\mathsf{T}\mathbf{x}$ directly.

4 Exact-mean product decompositions

Let us begin by stating the common mechanism that will be applied throughout this section.

Theorem 4.1 (block composition)**.** *Let $\mathcal{P}$ be any partition of the terminals, and suppose every block flow $\mathbf{x}^B$ has an exact-mean decomposition of radius ρ_B. Then $\mathbf{x}$ has an exact-mean decomposition of radius $\sum_B \rho_B$. If every block decomposition is two-sided, so is the conclusion.*

Proof. Draw the blocks independently and add. Atom bounds and means add across blocks, and every terminal lies in exactly one block. The cost assertion is part of Definition 3.1. □

Theorem 4.2 (chain-partition product)**.** *Fix a partition $\mathcal{P}$ of the terminals into blocks such that the demand multiset of every block B forms a divisibility chain, and let*

$$S(\mathcal{P}) := \sum_{B\in\mathcal{P}} \max_{i\in B} d_i.$$

Then every feasible flow $\mathbf{x}$ is a convex combination of support routings with exact mean and

$$|\mathbf{y} - \mathbf{x}| \le S(\mathcal{P})\,\mathbf{1} \qquad \textit{componentwise.}$$

Consequently, for every cost vector $\mathbf{c}$ (even with signed entries), some support routing satisfies

$$\mathbf{c}^\mathsf{T}\mathbf{y} \le \mathbf{c}^\mathsf{T}\mathbf{x}, \qquad \mathbf{x} - S(\mathcal{P})\mathbf{1} \le \mathbf{y} \le \mathbf{x} + S(\mathcal{P})\mathbf{1}.$$

Proof. Fix a path decomposition $\mathbf{x} = \sum_i \mathbf{x}_i$ and let $\mathbf{x}^B = \sum_{i\in B}\mathbf{x}_i$. Theorem 3.3 and Lemma 3.4 give block B an exact-mean two-sided decomposition of radius $\max_{i\in B} d_i$. Apply Theorem 4.1. The radii add to $S(\mathcal{P})$, and exact mean gives the cost assertion even for a signed cost vector. □

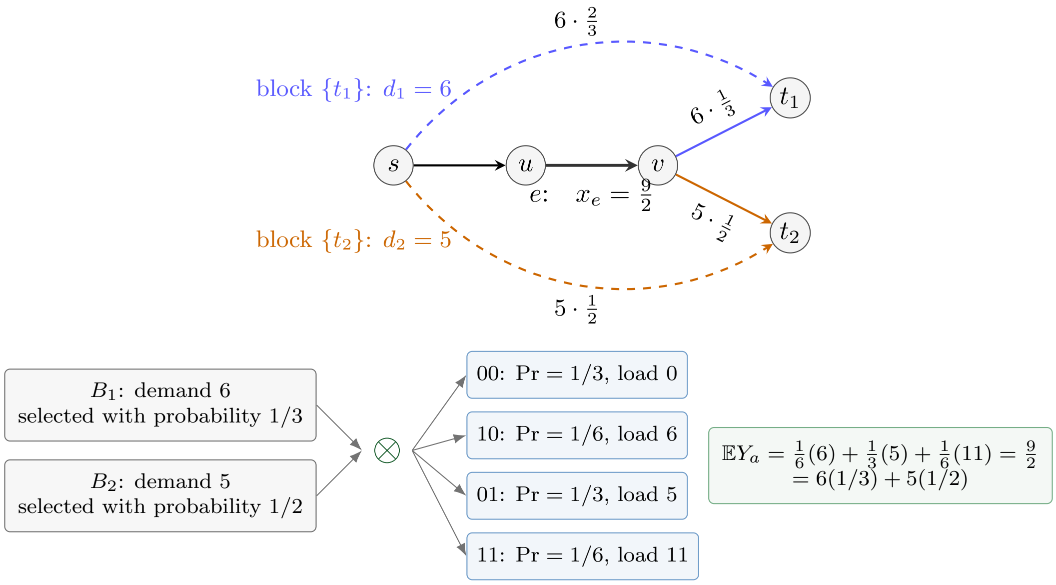

Figure 3: The independent product in Theorem 4.1 for the numerical example in text. Top: an instance realizing the example. Terminal t_1 (demand 6, one block) selects the route through the shared arc e with share $\frac{1}{3}$, terminal t_2 (demand 5, the other block) with share $\frac{1}{2}$; dashed arcs are the alternative routes, and arc labels are fractional flows, so $x_e = 6 \cdot \frac{1}{3} + 5 \cdot \frac{1}{2} = \frac{9}{2}$. Bottom: the product law. The four atom probabilities are products of the two block probabilities; their weighted arc load is exactly the sum of the two fractional block loads.

Remark 1. No terminal is split between blocks and no path is reassembled. The only probabilistic operation is an independent product of finitely many finite laws; the two-sided error is simply the sum of their radii.

For an illustration, take two one-terminal blocks with demands 6 and 5, and let their selected-route shares be $1/3$ and $1/2$ (Fig. 3). Independent rounding gives states $00, 10, 01, 11$ with probabilities $1/3, 1/6, 1/3, 1/6$. On an arc used by both selected routes the four integral loads are $0, 6, 5, 11$, while their weighted mean is

$$\tfrac{1}{6}\,6 + \tfrac{1}{3}\,5 + \tfrac{1}{6}\,11 = \tfrac{9}{2} = 6 \cdot \tfrac{1}{3} + 5 \cdot \tfrac{1}{2}.$$

The theorem's radius $6 + 5 = 11$ is just a worst-case triangle bound; the important invariant is the exact mean, which is additive even without matching labels or paths across the two blocks.

Remark 2. Theorem 4.2 is the case in which every block carries the chain primitive ($\rho_B = \max_{i \in B} d_i$); the merger reduction of Section 7 is the case in which every dyadic block carries a hypothetical $\mathrm{FT}(K)$ primitive ($\rho_{B_j} = Kh_j$). Any future per-block mean-exact decomposition, which we will consider for structured graph classes, special value patterns, and a new rounding primitive, will plug into the same composition with its ρ_B added to the budget.

Corollary 4.3 (scale sum)**.** *Let $\mathcal{V} = \{d_i\}$ be the set (not multiset) of distinct demand values and let $R = \sum_{v \in \mathcal{V}} v$. Then*

$$\mathbf{x} \in \operatorname{conv}\big\{\mathbf{y} : \ \mathbf{y} \text{ is a support routing}, \ |\mathbf{y} - \mathbf{x}| \le R\mathbf{1}\big\},$$

and in particular the cost-preserving additive constant of the instance is at most R/D.

Proof. Partition terminals by equal demand value; every block is a one-value chain, and $S = \sum_{v \in \mathcal{V}} v = R$. Apply Theorem 4.2; the equal-value case needs only Lemma 3.2. $\square$

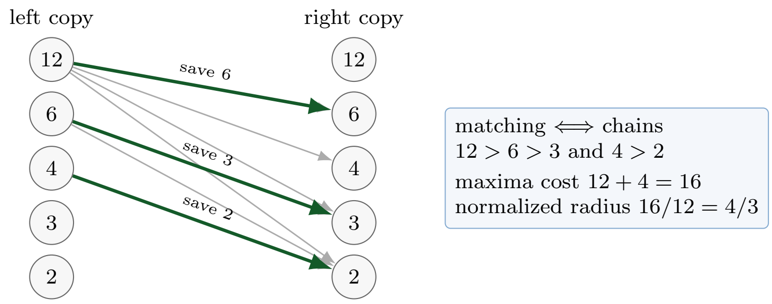

Figure 4: The matching reduction in Proposition 4.5 on $\{12, 6, 4, 3, 2\}$. Gray edges are divisibility links, green edges form the maximum-weight matching. A value used once as a right endpoint and once as a left endpoint becomes an internal vertex of a chain.

Corollary 4.4 (immediate special cases)**.**

1. *If the demands take two values $D > d$, then $C \le 1 + d/D < 2$; if additionally $d \mid D$, then $C \le 1$.*
2. *If successive distinct values decrease by factors at least two, then $R < 2D$ and $C < 2$.*
3. *If every dyadic band $(2^{-(j+1)}D,\ 2^{-j}D]$ contains at most m distinct demand values, then $R < 2mD$ and $C < 2m$; the total number of values may be unbounded.*

Proof. (1) The set $\mathcal{V} = \{D, d\}$ gives $R = D + d$; under $d \mid D$ the whole demand multiset is one chain with maximum D, so Theorem 4.2 applies with $S = D$. (2) $R \le D \sum_{j \ge 0} 2^{-j} = 2D$, and the inequality is strict unless the values are exactly geometric and infinite. (3) Sum the band bounds: $R < \sum_{j \ge 0} m\, D 2^{-j} = 2mD$. □

The instance-optimal form of the theorem is an efficiently computable parameter. On the distinct values $\mathcal{V}$ define

$$\kappa_{\mathrm{div}}(\mathcal{V}) := \frac{1}{D} \min_{\mathcal{P}} \sum_{B \in \mathcal{P}} \max B,$$

the minimum over partitions of $\mathcal{V}$ into divisibility chains.

Proposition 4.5 (optimal chain cover)**.** *$C \le \kappa_{\mathrm{div}}(\mathcal{V})$, and $\kappa_{\mathrm{div}}(\mathcal{V})$ is computable by one maximum-weight bipartite matching on $O(|\mathcal{V}|^2)$ edges.*

Proof. The bound is Theorem 4.2 applied to the partition of the terminals induced by an optimal partition of the values (terminals sharing a value travel together; repeats inside a chain are allowed). For the computation, make left and right copies of $\mathcal{V}$ and put an edge $u_L v_R$ of weight v whenever $u > v$ and $v \mid u$. A matching is exactly a set of consecutive links of vertex-disjoint decreasing divisibility chains, and conversely. With no edges every value is a chain maximum, at total weight $\sum_v v$; matching $u_L v_R$ makes v a non-maximum, saving exactly v. Hence $\min_{\mathcal{P}} \sum_B \max B = \sum_v v -$ (max-weight matching). □

For a specific example of the chain cover, consider $\mathcal{V} = \{12, 6, 4, 3, 2\}$ with matching links $12 \to 6 \to 3$ and $4 \to 2$ (Fig. 4). Its saved weight is $6 + 3 + 2 = 11$, so the cost of unmatched maxima is $(12+6+4+3+2) - 11 = 16$, corresponding to the two chains $12 > 6 > 3$ and $4 > 2$. Thus this demand set has the immediate normalized bound $C \le 16/12 = 4/3$; a tempting cover $12 > 4 > 2$ and $6 > 3$ would cost 18 instead.

Remark 3 (two-sidedness)**.** All bounds in this section are two-sided: the atoms satisfy $\mathbf{y} \ge \mathbf{x} - S\mathbf{1}$ as well. The results therefore also contribute to the two-sided program of [9], whose conjectured universal two-sided error is below D: e.g., two-valued instances admit routings with $|\mathbf{y} - \mathbf{x}| \le (1 + d_{(2)}/D)D\mathbf{1}$.

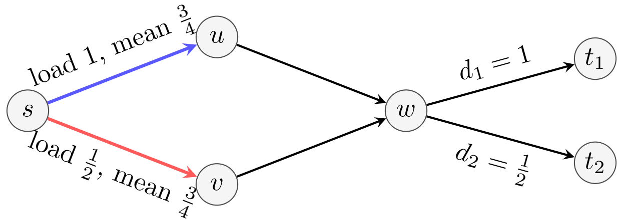

Figure 5: The labelled-marginal wall (Proposition 5.1). With both demands rounded up to 1, the routing "t_1 via u, t_2 via v" matches the padded aggregate flow exactly on every arc, so it is a valid aggregate decomposition by itself. Restoring the true demands gives loads $(1, \frac{1}{2})$ on the two source arcs against the true fractional means $(\frac{3}{4}, \frac{3}{4})$: the merge at w has erased which terminal used which prefix.

5 Why rounding up fails: the labelled-marginal wall

It is tempting to reduce unequal demands to Lemma 3.2 by rounding demands *up* to a common value d, decomposing the padded aggregate flow, and restoring the true demands on the selected paths. The next proposition shows that this idea has a fundamental flaw: an aggregate decomposition does not know which terminal rides which path, and after unequal demands are restored that information becomes necessary. This is a brief section to close this specific (quite tempting) path.

Proposition 5.1 (two-terminal labelled-marginal wall)**.** *Aggregate mean preservation after round-up does not imply preservation of the terminal-weighted mean. For a specific example, consider a DAG with arcs $s \to u$, $s \to v$, $u \to w$, $v \to w$, $w \to t_1$, $w \to t_2$, let each terminal fractionally route half of its demand via u and half via v, with true demands $d_1 = 1$, $d_2 = 1/2$, both rounded up to 1 (Fig. 5). The single routing (terminal 1 via u, terminal 2 via v) is by itself a valid convex decomposition of the padded aggregate flow, yet its true loads on $(s\,u, s\,v)$ are $(1, \frac{1}{2})$ while the true fractional loads are $(\frac{3}{4}, \frac{3}{4})$.*

Proof. The padded aggregate flow places load 1 on each of $s \to u$, $s \to v$, load 2 on the merged middle, and load 1 on each sink arc. The stated routing has exactly these padded loads, so as an aggregate integral point it decomposes the padded flow with coefficient 1. Restoring $(d_1, d_2) = (1, \frac{1}{2})$ gives loads $(1, \frac{1}{2})$ on the two source arcs. The true fractional loads are $(\frac{3}{4}, \frac{3}{4})$. Thus, the merge at w has erased the labels: aggregate correctness cannot see which terminal used which prefix. □

Proposition 5.1 does not refute any inequality claimed in this note, it refutes a specific proof technique. But it does mean that every valid route below either keeps demands unmodified inside blocks (Section 4) or rounds *down* and pays a multiplicative restoration cost (Section 6).

6 Rounding down: stripping, grids, and shifted chains

Rounding demands down preserves cost if the removed fraction is stripped from the most expensive paths first. This is Skutella's method [16]; we state the exact form we need. Unlike the product theorems, this route deliberately changes the demands and therefore gives up exact mean. The stripping inequality pays for that change on the cost side, while the ratio between true and rounded demands becomes the multiplicative load term.

Lemma 6.1 (stripping expensive paths)**.** *Let $\mathbf{x}$ satisfy demands d_i and choose targets $0 < \bar{d}_i \le d_i$. There is a subflow $\bar{\mathbf{x}} \le \mathbf{x}$ satisfying the demands $\bar{d}_i$ such that for every choice of paths*

$P_i \subseteq \operatorname{supp}(\bar{\mathbf{x}})$,

$$\sum_i (d_i - \bar{d}_i)\, c(P_i) \;\le\; \mathbf{c}^\mathsf{T}(\mathbf{x} - \bar{\mathbf{x}}).$$

Proof. Process the terminals in any order. For terminal i, repeatedly select a maximum cost path $s \to t_i$ inside the positive support of the current aggregate flow and subtract as much flow as possible along it, until the divergence at t_i has decreased by exactly $d_i - \bar{d}_i$ (the last subtraction may be partial). This operation changes only the divergences at s and t_i. A nonnegative acyclic flow with positive residual divergence at t_i contains a path $s \to t_i$, so the process cannot get stuck; a maximum-cost such path exists because the graph is finite and acyclic.

Let $\bar{\mathbf{x}}$ be the final aggregate flow and fix any $P_i \subseteq \operatorname{supp}(\bar{\mathbf{x}})$. At every moment of terminal i's processing all arcs of P_i carried positive aggregate flow since loads only decrease, and P_i survives to the end. Thus, P_i was one of the available paths $s \to t_i$. Therefore, every path stripped while processing terminal i had cost at least $c(P_i)$. The amount stripped for terminal i is in total $d_i - \bar{d}_i$, hence its cost is at least $(d_i - \bar{d}_i)c(P_i)$. Summing over terminals and noting that the stripped pieces are disjoint portions of $\mathbf{x} - \bar{\mathbf{x}}$, we get the lemma. □

Theorem 6.2 (geometric grid)**.** *For every $r > 1$, every feasible $\mathbf{x}$, and every $\mathbf{c} \ge \mathbf{0}$ there is an unsplittable routing $\mathbf{y}$ with*

$$\mathbf{c}^\mathsf{T}\mathbf{y} \le \mathbf{c}^\mathsf{T}\mathbf{x}, \qquad \mathbf{y} \le r\,\mathbf{x} + \frac{r}{r-1}\,D\mathbf{1}.$$

Proof. Round each d_i down to $\bar{d}_i$, the largest value of the grid $\{Dr^{-j} : j \ge 1\}$ not exceeding it; then $\bar{d}_i \le d_i \le r\bar{d}_i$. Apply Lemma 6.1 to get $\bar{\mathbf{x}}$, a feasible flow for the rounded demands. The distinct rounded values sum to at most $\bar{R} \le D\sum_{j\ge1} r^{-j} = D/(r-1)$. By Corollary 4.3 applied to $(\bar{\mathbf{x}}, \bar{d})$, there is a mixture of support routings of $\bar{\mathbf{x}}$ with mean $\bar{\mathbf{x}}$; its atoms $\bar{\mathbf{y}}$ obey $\bar{\mathbf{y}} \le \bar{\mathbf{x}} + \bar{R}\mathbf{1}$, and since the mean rounded cost is $\mathbf{c}^\mathsf{T}\bar{\mathbf{x}}$, some atom has rounded cost $\sum_i \bar{d}_i\, c(P_i) \le \mathbf{c}^\mathsf{T}\bar{\mathbf{x}}$.

Let us restore the true demands on that atom's paths. The cost is then

$$\mathbf{c}^\mathsf{T}\mathbf{y} = \sum_i \bar{d}_i\, c(P_i) + \sum_i (d_i - \bar{d}_i)c(P_i) \le \mathbf{c}^\mathsf{T}\bar{\mathbf{x}} + \mathbf{c}^\mathsf{T}(\mathbf{x} - \bar{\mathbf{x}}) = \mathbf{c}^\mathsf{T}\mathbf{x},$$

using Lemma 6.1 for the second sum. As for the load, on every arc we have

$$y_a \le r\bar{y}_a \le r\bar{x}_a + r\bar{R} \le rx_a + \frac{r}{r-1}D.$$

□

Coupling the grid levels into shifted divisibility chains halves the additive term.

Theorem 6.3 (m shifted chains)**.** *Let $b \ge 2$ and $m \ge 1$ be integers, and let $r = b^{1/m}$. For every feasible $\mathbf{x}$ and $\mathbf{c} \ge \mathbf{0}$ there exists a routing $\mathbf{y}$ with*

$$\mathbf{c}^\mathsf{T}\mathbf{y} \le \mathbf{c}^\mathsf{T}\mathbf{x}, \qquad \mathbf{y} \le r\,\mathbf{x} + A_{b,m}D\mathbf{1}, \qquad A_{b,m} = \frac{r\,(1 - b^{-1})}{r-1}.$$

In particular, for $b = 2$

$$\mathbf{y} \le 2^{1/m}\mathbf{x} + \frac{2^{1/m}}{2\,(2^{1/m} - 1)}\,D\mathbf{1}. \tag{1}$$

Proof. Round down to the grid $\{Dr^{-j} : j \ge 1\}$ and strip as in Theorem 6.2. Partition the grid levels by the residue of j modulo m: within one residue class consecutive levels differ by the integer factor $r^m = b$, so the rounded demands of that class form a divisibility chain, and

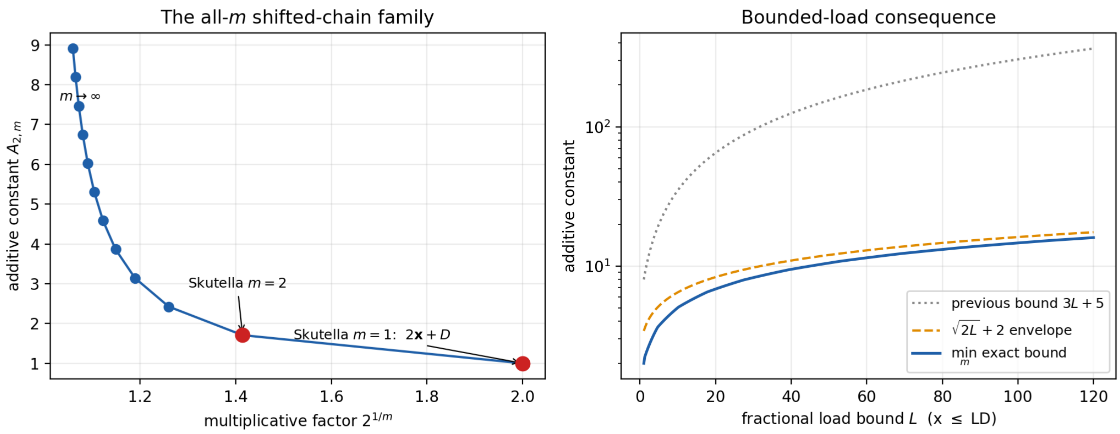

Figure 6: Left: the all-m shifted-chain family of Theorem 6.3 in the (multiplicative, additive) plane; Skutella's two documented bounds are the first two points of the curve. Right: the bounded-load consequence (Corollary 6.4): the exact $\min_m$ bound, its $\sqrt{2L}+2$ envelope, and the earlier $3L+5$ bound, on a logarithmic scale.

the largest level in residue class $h \in \{1, \ldots, m\}$ is at most Dr^{-h}. Theorem 4.2 applied to this m-block partition of the rounded instance gives a mean-exact mixture with atomwise error at most

$$S \le D \sum_{h=1}^{m} r^{-h} = D\,\frac{1-b^{-1}}{r-1}.$$

Choose an atom whose rounded cost is at most $\mathbf{c}^{\mathsf{T}}\bar{\mathbf{x}}$ and restore true demands: the cost bound follows from Lemma 6.1 exactly as before, and $\mathbf{y} \le r\bar{\mathbf{y}} \le r\mathbf{x} + rS\mathbf{1} = r\mathbf{x} + A_{b,m}D\mathbf{1}$. $\square$

For $(b, m) = (2, 1)$, inequality (1) becomes $\mathbf{y} \le 2\mathbf{x} + D\mathbf{1}$, which is Skutella's first cost-preserving bound. For $(b, m) = (2, 2)$ it reads $\mathbf{y} \le \sqrt{2}\mathbf{x} + (1 + 1/\sqrt{2})D\mathbf{1}$, which is his second bound. Theorem 6.3 is a generalization of these two bounds; Figure 6 plots the family and its bounded load corollary.

Corollary 6.4 (bounded fractional load)**.** *If* $\mathbf{x} \le LD\mathbf{1}$ *with* $L \ge 1$*, then*

$$C \;\le\; \min_{m \ge 1}\left\{(2^{1/m}-1)L + \frac{2^{1/m}}{2(2^{1/m}-1)}\right\} \;\le\; \sqrt{2L} + 2.$$

Proof. Subtracting $\mathbf{x}$ from (1) turns the multiplicative term into $(2^{1/m}-1)L$. For the explicit estimate set $a = \log 2$, $m = \lceil a(\sqrt{2L}+1)\rceil$, $t = a/m \le 1/(\sqrt{2L}+1)$. Then

$$e^t - 1 \le \frac{t}{1-t} \le \frac{1}{\sqrt{2L}} \quad\Longrightarrow\quad (2^{1/m}-1)L \le \sqrt{L/2},$$

and, using $1 - e^{-t} \ge t(1 - t/2)$,

$$\frac{2^{1/m}}{2(2^{1/m}-1)} = \frac{1}{2(1-e^{-t})} \le \frac{1}{2t} + \frac{1}{2} \le \frac{m}{2a} + \frac{1}{2} \le \sqrt{L/2} + 1 + \frac{1}{2\log 2}.$$

The sum is below $\sqrt{2L} + 2$. $\square$

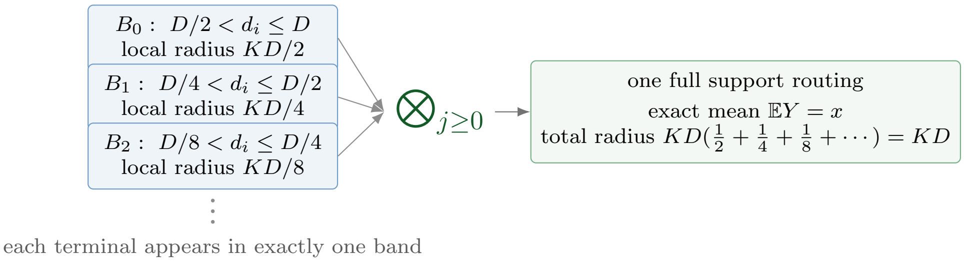

Figure 7: Illustration for Theorem 7.2: a factor-two primitive is applied independently inside each demand band, its local radii form a geometric series summing to KD, and exact means add to the original flow.

7 The factor-two merger reduction

In this section, we prove a product theorem that localizes the entire universal problem into one statement about a single factor-two window of demand values.

Definition 7.1. *For $K \ge 0$ and $\beta \ge 0$, let* $\mathrm{FT}(K,\beta)$ *denote the following statement:* for every $h > 0$ and every feasible flow $\mathbf{p}$ of an instance whose demands all lie in $(h, 2h]$, there is a distribution over support routings $\mathbf{Y}$ of $\mathbf{p}$ with

$$\mathbb{E}[\mathbf{Y}] \ \ge\ \mathbf{p} - \beta h \mathbf{1}, \qquad \mathbf{Y} \ \le\ \mathbf{p} + Kh\mathbf{1} \text{ almost surely.}$$

Write $\mathrm{FT}(K) := \mathrm{FT}(K,0)$ *with the mean required to be exact,* $\mathbb{E}[\mathbf{Y}] = \mathbf{p}$.

Theorem 7.2 (dyadic merger reduction)**.** *If* $\mathrm{FT}(K)$ *holds, the universal cost-preserving additive constant satisfies $C \le K$.*

Proof. Given $\mathbf{x}$ and $\mathbf{c} \ge \mathbf{0}$, partition the terminals into dyadic blocks $B_j = \{i : 2^{-(j+1)}D < d_i \le 2^{-j}D\}$ and decompose $\mathbf{x} = \sum_j \mathbf{x}^j$ along a path decomposition. Apply $\mathrm{FT}(K)$ to each nonempty block with $h_j = 2^{-(j+1)}D$ and draw independently. Every product atom is a full support routing with

$$\mathbf{Y} \le \mathbf{x} + K\sum_{j\ge 0} h_j \mathbf{1} \le \mathbf{x} + KD\mathbf{1},$$

because $\sum_{j\ge 0} 2^{-(j+1)} = 1$ (see Fig. 7 for an illustration); the product mean is exactly $\mathbf{x}$, so the mean cost is $\mathbf{c}^\mathsf{T}\mathbf{x}$ and some atom is cost-good. □

Define the *band constant*

$$K^* := \inf\{K : \mathrm{FT}(K) \text{ holds}\} \in [0, \infty].$$

Theorem 7.3 (converse up to band normalization)**.** *If C_0 is a valid universal cost-preserving additive constant, then* $\mathrm{FT}(2C_0)$ *holds. Consequently*

$$C_{\mathrm{opt}} \ \le\ K^* \ \le\ 2C_{\mathrm{opt}},$$

with the convention that the right side is infinite when no universal constant exists. In particular, $K^* < \infty$ is equivalent to existence of a finite universal constant.

Proof. Fix a flow $\mathbf{p}$ with demands in $(h, 2h]$, and restrict the instance to $\mathrm{supp}(\mathbf{p})$. Let

$$\mathcal{S} = \{\mathbf{y} :\ \mathbf{y} \text{ is a routing in } \mathrm{supp}(\mathbf{p}),\ \mathbf{y} \le \mathbf{p} + 2C_0 h\mathbf{1}\}.$$

The assumed theorem with the zero cost vector first shows that $\mathcal{S} \neq \varnothing$, since $D \leq 2h$. We prove $\mathbf{p} \in \operatorname{conv} \mathcal{S}$. Otherwise finite-dimensional strict separation supplies $\mathbf{w} \in \mathbb{R}^A$ such that

$$\mathbf{w}^\mathsf{T}\mathbf{y} > \mathbf{w}^\mathsf{T}\mathbf{p} \qquad \text{for every } \mathbf{y} \in \mathcal{S}. \tag{8}$$

Since the support graph is acyclic, we can choose vertex potentials π in topological order with

$$c_{uv} := w_{uv} + \pi_v - \pi_u \geq 0 \qquad \text{for every support arc } (u, v);$$

for example, after fixing potentials at vertices with no incoming support arc, take $\pi_v \geq \max_{(u,v)}(\pi_u - w_{uv})$. Every routing and $\mathbf{p}$ have the same divergence vector, hence the potential term has the same value on both and

$$\mathbf{c}^\mathsf{T}(\mathbf{y} - \mathbf{p}) = \mathbf{w}^\mathsf{T}(\mathbf{y} - \mathbf{p}). \tag{9}$$

Apply the assumed universal theorem with the nonnegative cost $\mathbf{c}$ to the support-restricted instance. It returns a routing $\mathbf{y}$ with

$$\mathbf{y} \leq \mathbf{p} + C_0 D \mathbf{1} \leq \mathbf{p} + 2C_0 h \mathbf{1}, \qquad \mathbf{c}^\mathsf{T}\mathbf{y} \leq \mathbf{c}^\mathsf{T}\mathbf{p}.$$

Thus $\mathbf{y} \in \mathcal{S}$, while (9) contradicts (8). Therefore $\mathbf{p} \in \operatorname{conv} \mathcal{S}$, which is exactly $\mathrm{FT}(2C_0)$. Taking the infimum and using Theorem 7.2 gives the desired inequalities. □

Theorem 7.4 (relaxed merger reduction)**.** *If* $\mathrm{FT}(K, \beta)$ *holds, then* $C \leq K + \beta$*.*

Proof. Given $\mathbf{x}, \mathbf{c}$, set $\rho = \beta$ and let $\mathbf{q}, \boldsymbol{\lambda}$ be as in Lemma 3.5 for the box $\mathbf{p} \leq \mathbf{x} + \beta D \mathbf{1}$. Apply $\mathrm{FT}(K, \beta)$ blockwise to $\mathbf{q}$ and take the independent product $\mathbf{Y}$: a support routing of $\mathbf{q}$ with

$$\mathbb{E}[\mathbf{Y}] \;\geq\; \mathbf{q} - \beta \sum_j h_j \mathbf{1} \;\geq\; \mathbf{q} - \beta D \mathbf{1}, \qquad \mathbf{Y} \;\leq\; \mathbf{q} + KD\mathbf{1} \;\leq\; \mathbf{x} + (K + \beta) D \mathbf{1}.$$

By Lemma 3.5(3) and $\boldsymbol{\lambda} \geq \mathbf{0}$,

$$\mathbb{E}[\mathbf{c}^\mathsf{T}\mathbf{Y}] = \mathbf{c}^\mathsf{T}\mathbf{q} + \boldsymbol{\lambda}^\mathsf{T}(\mathbf{q} - \mathbb{E}[\mathbf{Y}]) \leq \mathbf{c}^\mathsf{T}\mathbf{q} + \beta D \|\boldsymbol{\lambda}\|_1 \leq \mathbf{c}^\mathsf{T}\mathbf{x},$$

using Lemma 3.5(2) at $\rho = \beta$. As usual, some atom is cost-good. □

Remark 4 (normalization)**.** The window is normalized by its *lower* endpoint h. If a merger lemma is proved with error K' times the upper endpoint $2h$, then $\mathrm{FT}(2K')$ holds and the reduction yields $C \leq 2K'$. Corollary 4.4(3) shows $\mathrm{FT}(2m)$ holds whenever every window contains at most m distinct values; and inside a factor-two window no two distinct values divide one another, so the chain machinery of Section 4 does not yield any results there.

Theorem 7.2 gives $C_{\mathrm{opt}} \leq K^*$, while Theorem 7.3 gives the reverse comparison within a factor two. The record itself gives a substantially stronger lower bound than the fixed-radius KKT extraction.

Proposition 7.5 (exact one-band lower bound)**.** *Let*

$$C_{17} = \frac{16032508626563604534001872727 1}{125000000000000000000000000000} = 1.28260069012508836272014981816 8\ldots.$$

Then

$$K^* \geq 2C_{17} = 2.5652013802501767254402996363 36\ldots.$$

In particular $\mathrm{FT}(K)$ *is false for every* $K < 2C_{17}$*.*

Proof. All demands of the standing refined $k = 17$ record [2] lie in $[0.788277893510092, 1] \subset (\frac{1}{2}, 1]$, so take $h = \frac{1}{2}$ and $D = 1$. Its exact hull-entry certificate says that the least D-normalized one-sided radius whose routing sublevel has the fractional flow in its convex hull is exactly C_{17}. The solver-free lower certificate is quite short: all 8763 routings of overload strictly below C_{17} select at most six late paths, whereas the sum of the seventeen fractional late shares is $6 + 10^{-15}$. Hence their convex hull cannot contain the fractional share vector. Conversely, an 18-atom rational mixture has exactly that barycenter and every atom has overload at most C_{17}. An $\mathrm{FT}(K)$ mixture on this instance would have atom radius $Kh = K/2$. Thus no such mixture exists when $K/2 < C_{17}$, proving the claim. $\square$

The earlier KKT extraction remains useful because it constrains the *relaxed* merger frontier rather than merely exact mean.

Proposition 7.6 (relaxed bound at $K = 2$)**.** *There is an explicit one-band flow* $\mathbf{q}$ *for which every distribution over support routings with atoms* $\mathbf{Y} \le \mathbf{q} + D\mathbf{1}$ *satisfies*

$$\max_a \big(q_a - \mathbb{E}[Y_a]\big) \;\ge\; \delta = 0.297915757320558\ldots,$$

where δ *is an exactly certified rational attained by a* 15*-atom mixture and proved optimal by an* 11*-arc dual separator. Consequently*

$$\mathrm{FT}(2, \beta) \quad\Longrightarrow\quad \beta \ge 2\delta = 0.595831514641116\ldots.$$

For this fixed flow, write $\beta_q(K)$ *for the least lower-band-normalized coordinatewise mean allowance among mixtures with atom bound* Kh*, and put* $\beta_q(K) = \infty$ *if that upper box contains no routing. On this same flow, however, the exact-mean hull entry is*

$$\rho_q = \frac{1000000002257139951}{10^{18}} D = \big(1 + 2.257139951 \cdot 10^{-9}\big) D.$$

Equivalently, the local exact-mean band threshold is $K_q = 2.000000004514279902$*. This value is exact: a* 16*-atom mixture attains it, and a rational terminal-marginal separator excludes the immediately preceding routing level* $1.000000002129669614D$*. Moreover, the complete local relaxed frontier on this flow has exactly* 15 *constant-defect segments before exact mean. Over all* $K \ge 0$*, its minimum of* $K + \beta_q(K)$ *is attained at the exact-mean endpoint and equals* K_q*. The best segment with positive mean defect has*

$$\begin{aligned} K &= 1.473335255532544\ldots, \\ \beta_q(K) &= 0.614334394476161\ldots, \\ K + \beta_q(K) &= 2.087669650008705\ldots. \end{aligned}$$

Proof. Take the KKT flow of the $k = 17$ record at capacity width $\gamma = C_{17} - 1 - 10^{-9}$. Its demands are the same one-band demands and $D = 1$. The exact fixed-box certificate[1] enumerates all 32768 positive-support routings, identifies the 4420 routings with $\mathbf{Y} \le \mathbf{q} + D\mathbf{1}$, and supplies matched rational primal and dual certificates for the displayed δ. Since $h = \frac{1}{2}$, the coordinatewise mean allowance in $\mathrm{FT}(2, \beta)$ is $\beta h = \beta/2$; therefore it cannot be smaller than δ.

For the exact-mean entry, enumerate the same 32768 support routings and sort their exact upper-overload levels. At ρ_q the displayed 16-atom rational mixture has the prescribed terminal marginals and, independently, mean $\mathbf{q}$ on every raw arc. At the immediately preceding level, the exact rational separator[2] is strictly larger on the target marginal vector than on every

[1] `goalA_fixedbox_k17_deficit_exact_20260813`; see Appendix A.

[2] `goalA_band_threshold_q_exact_20260809`; see Appendix A.

eligible routing. This proves both sides of the claimed threshold without relying on floating LP feasibility.

For the full frontier, sort all exact routing levels up to ρ_q (there are 95), and resolve the mean-defect LP only when the preceding dual ceases to dominate a newly admitted route. This produces 15 matched rational primal–dual segments. On each segment the same primal is feasible from its left endpoint and the same dual is checked against every route through its right endpoint; hence the defect is constant throughout. Comparing the 15 exact values of $2(t + \delta(t))$ proves the last assertion. $\square$

The flow-level frontier, however, is not the binding constraint on the relaxed route: the instance-level bound below supersedes it.

Corollary 7.7 (instance-level relaxed wall)**.** *If* $\mathrm{FT}(K, \beta)$ *holds, then*

$$K + \beta \;\geq\; 2C_{17} \;=\; 2.5652013802501767254402996363336\ldots.$$

Proof. Let the instance have all demands in one dyadic band, so $h = D/2$ and the proof of Theorem 7.4 has a single block. With box width $\rho = \beta h/D = \beta/2$, the mean allowance is $\beta h = (\beta/2)D$ and the atom bound is $Kh = (K/2)D$, so the instance's cost-preserving constant is at most $(K + \beta)/2$. The $k = 17$ record is such an instance and forces the constant C_{17}; hence $(K + \beta)/2 \geq C_{17}$. $\square$

Remark 5 (the refined record)**.** The standing record [2] shares the one-band structure: all seventeen demands lie in $(\frac{1}{2}, 1]$ (ten distinct values). We rebuilt its 67-arc DAG, rediscovered both paths of every terminal, enumerated all 2^{17} routings, and certified the hull entry at its published constant. The entry sublevel has 8782 routings and an exact 18-atom mixture; the immediate predecessor is separated by the rank functional with exact gap 10^{-15}. A second implementation reconstructed the interval rows and private exits without using the graph generator.

Remark 6 (finer windows do not help)**.** Windows of ratio $r < 2$ do not escape the normalization loss. Defining $\mathrm{FT}_r(K)$ on windows $(h, rh]$, the reduction over the r-adic tower costs $\sum_{j\geq 0} r^{-(j+1)} = 1/(r-1)$, giving $C \leq K/(r-1)$, while a single-window instance with top demand rh forces $K \geq rC_{\mathrm{win}}$. The end-to-end guarantee is therefore no better than $\frac{r}{r-1}C_{\mathrm{win}}$, which exceeds the dyadic value $2C_{\mathrm{win}}$ for every $r < 2$. The window reduction targets *existence* of a finite constant; it cannot by itself certify an optimal one.

At this point, let us summarize the state of the reduction:

$$\boxed{\begin{array}{cc} C_{\mathrm{opt}} \leq K^* \leq 2C_{\mathrm{opt}}, & K^* \geq 2C_{17} = 2.56520138025\ldots, \\ K + \beta \geq 2C_{17} \text{ for every valid relaxed pair}, & \mathrm{FT}(2m) \text{ holds on } m\text{-value bands.} \end{array}}$$

Neither the exact-mean nor the relaxed route can certify a universal constant below $2C_{17}$ on the window normalization (Corollary 7.7, Remark 6); the merger reduction's value is the *existence* of a finite constant, and the optimal constant must be pursued directly. Figure 8 shows the complete exact frontier of the KKT band flow against the instance-level wall: the flow-level obstruction is a razor edge whose relaxed minimum 2.0000000045 is dominated by the instance-level bound, confirming that further perturbation of this witness cannot improve the lower bound and that the band ladder and the record ladder are the same ladder, coupled by the exact factor two.

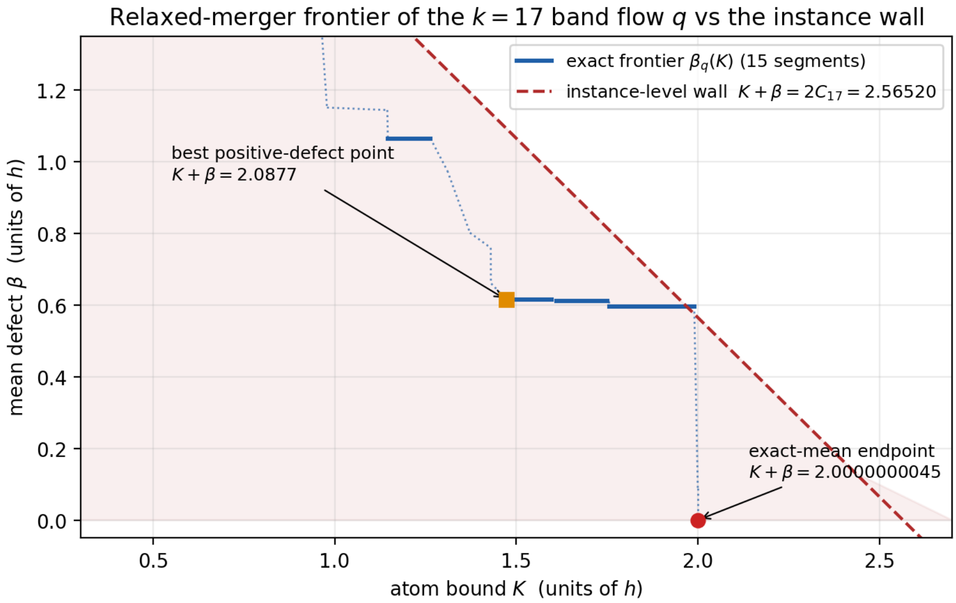

Figure 8: The complete exact relaxed-merger frontier of the $k = 17$ KKT band flow $\mathbf{q}$ (Proposition 7.6): fifteen constant-defect segments ending at the exact-mean entry $K + \beta = 2.0000000045$, with the best positive-defect point at $K + \beta = 2.0877$. The dashed line is the instance-level wall $K + \beta = 2C_{17}$ of Corollary 7.7, which dominates the flow-level frontier: perturbing this witness cannot improve the lower bound, and new band lower bounds require new record instances.

8 Graph classes and the mixture characterization

The merger reduction bounds the universal constant C through demand structure. In this section, we show that the same two-step machinery (exact-mean averaging in one direction, separation plus a potential shift in the other) can characterize the constant of every reasonable graph class, which turns known structured theorems into mixture and band statements and organizes what remains open. We call a class $\mathcal{G}$ of acyclic single-source instances *restriction-closed* if it is preserved by deleting arcs.

8.1 The characterization and the class table

Theorem 8.1 (mixture characterization)**.** *Let $\mathcal{G}$ be a restriction-closed class of graphs. Let $C(\mathcal{G})$ be the optimal cost-preserving additive constant over $\mathcal{G}$, and let*

$$T(\mathcal{G}) := \inf\big\{t : \textit{for every feasible flow } \mathbf{p} \textit{ of an instance of } \mathcal{G}, \\ \mathbf{p} \in \operatorname{conv}\{\mathbf{y} \textit{ is a support routing} : \ \mathbf{y} \le \mathbf{p} + tD\mathbf{1}\}\big\}.$$

Then $C(\mathcal{G}) = T(\mathcal{G})$.

Proof. ($C \le T$.) For $t > T(\mathcal{G})$ and any instance $(\mathbf{x}, \mathbf{c})$ of the class, take an exact-mean mixture at $\mathbf{x}$ with atoms $\le \mathbf{x} + tD\mathbf{1}$: its mean cost is $\mathbf{c}^\mathsf{T}\mathbf{x}$, so some atom is cost-preserving.

($T \le C$.) This is the separation argument of Theorem 7.3 without the band specialization. Let $C_0 > C(\mathcal{G})$, let $\mathbf{p}$ be a feasible flow of a class instance, and let $\mathcal{S} = \{\mathbf{y} \text{ is a support routing} : \mathbf{y} \le \mathbf{p} + C_0 D\mathbf{1}\}$, nonempty by the class theorem with zero cost. If $\mathbf{p} \notin \operatorname{conv}\mathcal{S}$, separate strictly by $\mathbf{w}$, shift by acyclic potentials to $\mathbf{c} \ge \mathbf{0}$ on $\operatorname{supp}(\mathbf{p})$ with $\mathbf{c}^\mathsf{T}(\mathbf{y} - \mathbf{p}) = \mathbf{w}^\mathsf{T}(\mathbf{y} - \mathbf{p})$ for every routing, and apply the class theorem to the support-restricted instance (it is still in $\mathcal{G}$

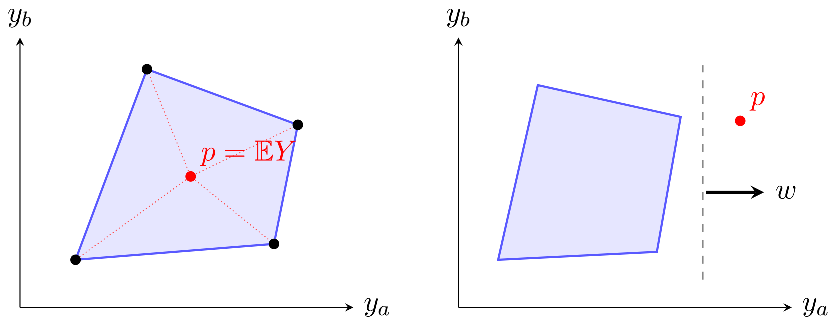

(a) exact-mean mixture (b) strict separator

Figure 9: The mixture characterization in a two-arc projection. If the fractional point lies in the hull of good routings, barycentric weights are an exact-mean decomposition. If it lies outside, a separator w exists; acyclic potentials turn w into a nonnegative arc cost without changing its value on any routing relative to the fractional flow, contradicting the cost-preserving theorem.

since it is restriction-closed) with cost $\mathbf{c}$: it returns an element of $\mathcal{S}$ on the wrong side of the separator. □

Remark 7. Cost preservation is therefore a purely polyhedral property: on every restriction-closed class, the optimal constant equals the least one-sided radius at which fractional flows become exact-mean mixtures of support routings, and no cost vector appears in the definition of $T(\mathcal{G})$. In particular $C_{\text{opt}} = T(\text{all DAGs})$: the hull-entry searches of the companion program have been measuring exactly this quantity on witnesses. Lower-box deficit computations concern a different, though related, KKT relaxation and should not be conflated with T.

Corollary 8.2 (class bounds)**.**

1. Out-trees. *If the support of the flow is an out-arborescence, every terminal has a unique support path, the unique support routing equals the flow, and the class constant is* 0.

2. Two-layer hub instances. *Suppose every arc is either a shared arc $s \to m_j$ or a private arc $m_j \to t_i$, so a routing assigns each terminal to one shared arc. Then $C \le 1$: shared-arc loads are machine loads of an assignment problem with processing times $d_i \le D$, path costs are assignment costs, and the Shmoys–Tardos rounding for generalized assignment [15] produces an integral assignment with cost at most the fractional cost and every machine load at most its fractional load plus D; private arcs never bind, because $d_i \le d_i f_{ij} + D$ always. In fact the bound is sharp: $C(hub) = T(hub) = 1$. For every $m \ge 2$, take $m+1$ unit-demand terminals and m hubs, and assign every terminal fraction $1/m$ to every hub. Each shared arc has fractional load $1 + 1/m$. Every integral assignment puts two terminals on some hub and therefore has one-sided error at least $1 - 1/m$. Conversely, the uniform mixture over all assignments with occupancy multiset $(2, 1, \ldots, 1)$ has every terminal–hub marginal exactly $1/m$ and every atom has error $1 - 1/m$. Thus its hull entry is exactly $1 - 1/m$, which tends to one. By Theorem 8.1 equality follows, and* FT *restricted to hub instances holds with $K = 2$.*

3. Planar DAGs. *$C(planar) \le 2$ [18], hence $T(planar) \le 2$ and* FT *restricted to planar instances holds with $K = 4$. Together with Proposition 8.3,*

$$\frac{58676765987259}{50000000000000} = 1.17353531974518 \;\le\; C(planar) \;\le\; 2.$$

4. Series-parallel DAGs. *Almoghrabi, Skutella, and Warode [1] prove the stronger statement that every fractional multiflow on their series-parallel class is an exact convex combination of unsplittable multiflows whose arc loads differ from it by less than D. Zero-load arcs have zero load in every positive-weight atom, so the decomposition is support-preserving. Consequently T(series-parallel) ≤ 1 and C(series-parallel) ≤ 1 directly; the factor-two transfer below is unnecessary on this class. Both bounds are sharp. For $m \geq 2$, take m internally disjoint directed length-two paths from s to v, followed in series by a directed chain through $m+1$ distinct terminal vertices. Give each terminal unit demand and split it uniformly among the m branches before following its forced chain prefix. Every routing assigns two commodities to some branch, whose two arcs have fractional load $1 + 1/m$, and hence has error at least $1 - 1/m$. Conversely, the uniform mixture over assignments with occupancy multiset $(2, 1, \ldots, 1)$ has every commodity–branch marginal $1/m$, is exact-mean on the forced tail, and every atom has error $1 - 1/m$. The graph is two-terminal series-parallel, with global terminals s and the last vertex of the chain. Letting $m \to \infty$ proves*

$$T(\textit{series-parallel}) = C(\textit{series-parallel}) = 1.$$

5. Outerplanar two-exit interval spines. *In this class each terminal has an early and a late exit from one directed spine, so its binary route difference is supported on an interval. Theorem 8.5 constructs an exact-mean mixture whose every atom has absolute arc error at most D. The one-terminal family in the same theorem is cost-sharp, and hence*

$$T(\textit{outerplanar interval spine}) = C(\textit{outerplanar interval spine}) = 1.$$

8.2 Planar records and outerplanar spines

Proposition 8.3 (a new planar record)**.** *The planar cost-preserving constant satisfies*

$$C(\text{planar}) \geq \frac{58676765987259}{50000000000000} = 1.17353531974518.$$

In fact this number is the exact one-sided hull entry of an explicit six-terminal planar flow.

Proof. The instance, shown in Fig. 10, is visibly planar. Take the directed spine $v_0 \to v_1 \to \cdots \to v_9$. For terminal i, add private exits $v_{\ell_i} \to t_i$ and $v_{r_i+1} \to t_i$, where

$$([\ell_i, r_i])_{i=0}^{5} = ([0,5],[1,6],[2,4],[3,6],[3,7],[4,7]).$$

The demands and late-path shares are, respectively,

$$\mathbf{d} = \left(\frac{224889}{250000}, \frac{39841}{50000}, \frac{199989}{200000}, \frac{249607}{250000}, \frac{823711}{10^6}, 1\right)$$

and

$$\mathbf{f} = \left(\frac{124539}{10^6}, \frac{513783001}{10^9}, \frac{50197}{500000}, \frac{556341}{10^6}, \frac{162753}{500000}, \frac{379437}{10^6}\right), \qquad \mathbf{1}^{\mathsf{T}}\mathbf{f} = 2 + 10^{-9}.$$

Here terminal i uses its late exit with share f_i. Put unit cost $1/d_i$ on its early private exit and zero cost on every other arc. Thus the fractional cost is $6 - \mathbf{1}^{\mathsf{T}}\mathbf{f} = 4 - 10^{-9}$, whereas a binary routing $z \in \{0,1\}^6$ has cost $6 - \mathbf{1}^{\mathsf{T}}z$. Consequently exactly the 42 routings with $\mathbf{1}^{\mathsf{T}}z \geq 3$ are cost-good.

Direct exact evaluation of the 64 routing load vectors gives

$$\min_{\mathbf{1}^{\mathsf{T}}z \geq 3} \max_a (y_a(z) - x_a) = \frac{58676765987259}{50000000000000};$$

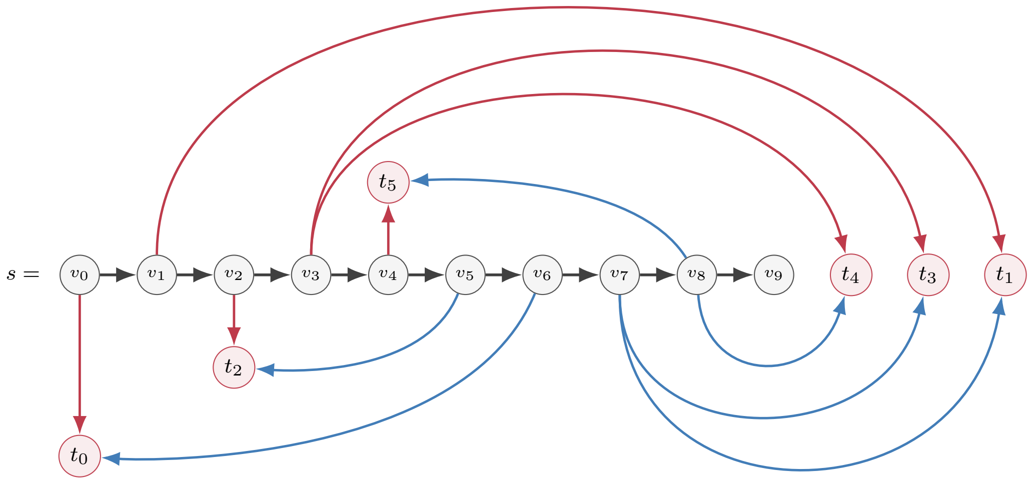

red: early/expensive exit blue: late/free exit black: directed spine

intervals: $t_0 : [0, 5]$, $t_1 : [1, 6]$, $t_2 : [2, 4]$, $t_3 : [3, 6]$, $t_4 : [3, 7]$, $t_5 : [4, 7]$

Figure 10: The explicit 16-vertex, 21-arc DAG of Proposition 8.3, drawn with its directed spine horizontal. Each terminal has an early red exit from v_{ℓ_i} and a late blue exit from v_{r_i+1}.

the unique minimizer is 100101. For completeness, the same evaluation proves the stronger hull statement without relying on the cost vector. Every one of the 18 routes strictly below the displayed level has $\mathbf{1}^\mathsf{T} z \leq 2$, which separates their hull from $\mathbf{f}$ by the gap 10^{-9}. At the displayed level, the following seven-atom mixture has barycenter exactly $\mathbf{f}$:

z	000101	000110	001100	010000	010010	010100	100101
$10^9\lambda_z$	254898000	6385999	100394000	124538999	319120001	70124001	124539000

Its coefficients sum to one, and substitution into every spine and private-exit row shows that each atom has overload at most the displayed fraction. □

Note that while the graph in Proposition 8.3 is planar, it is not outerplanar, and it is also outside the directed two-terminal series-parallel class of [1]: its orientation has seven sinks, whereas a directed graph obtained recursively from one forward arc by series and parallel composition has a unique sink. Thus it avoids two structural firewalls.

Proposition 8.4 (the common-point outerplanar wall)**.** *Every spine-and-tent graph whose intervals share one spine arc and which has at least three terminals contains $K_{2,3}$ as a minor. In particular, no strict-rank common-point cell with at least three terminals is outerplanar.*

Proof. Let the common arc be $v_r v_{r+1}$. Contract the connected left spine block $\{v_0, \ldots, v_r\}$ to one vertex L, and contract the disjoint right block $\{v_{r+1}, \ldots, v_p\}$ to one vertex R. A tent for interval $[a_i, b_i]$ is the degree-two path $v_{a_i} t_i v_{b_i+1}$. Since $a_i \leq r \leq b_i$, the singleton branch set $\{t_i\}$ is adjacent to both L and R. Retaining any three such singletons gives a $K_{2,3}$ minor. The forbidden-minor characterization of outerplanar graphs completes the proof. See Figure 11 for an illustration. □

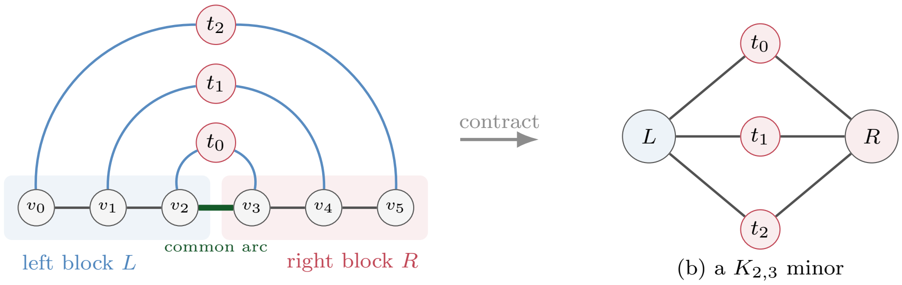

(a) three nested tents through one common arc

(b) a $K_{2,3}$ minor

Figure 11: The minor argument in Proposition 8.4.

Theorem 8.5 (forest coupling for interval spines)**.** *Consider a two-exit interval spine: terminal i has an early private exit at v_{ℓ_i} and a late private exit at v_{r_i+1}, with $\ell_i \le r_i$, and uses the late path with fractional share f_i. If the overlap graph of the intervals $I_i = [\ell_i, r_i]$ is a forest, then the fractional flow is an exact-mean mixture of support routings $\mathbf{Y}$ such that*

$$|Y_a - x_a| \le D \qquad \textit{for every raw arc } a \textit{ almost surely.}$$

Every outerplanar two-exit interval spine satisfies the hypothesis. Moreover, the resulting constant one is sharp, even with one terminal.

Proof. Terminals with $f_i \in \{0, 1\}$ are fixed and can be removed, so assume $0 < f_i < 1$. Fix an overlap edge ij. On a spine row whose nonconstant terminal set is $\{i, j\}$, the routing error is

$$E_{ij}(z_i, z_j) = d_i(z_i - f_i) + d_j(z_j - f_j).$$

The mixed states 10 and 01 always have absolute error at most D. Only the extremes can be bad:

$$\begin{aligned} 00 \text{ is bad} &\iff d_i f_i + d_j f_j > D, \\ 11 \text{ is bad} &\iff d_i(1 - f_i) + d_j(1 - f_j) > D. \end{aligned} \tag{10}$$

They cannot both be bad, because the two left sides sum to $d_i + d_j \le 2D$.

If 00 is bad, then $D(f_i + f_j) \ge d_i f_i + d_j f_j > D$, so $f_i + f_j > 1$; use the edge law

$$(p_{00}, p_{10}, p_{01}, p_{11}) = (0, 1 - f_j, 1 - f_i, f_i + f_j - 1). \tag{11}$$

It has Bernoulli marginals (f_i, f_j) and deletes the bad state. If 11 is bad, the complementary inequality gives $f_i + f_j < 1$; use

$$(p_{00}, p_{10}, p_{01}, p_{11}) = (1 - f_i - f_j, f_i, f_j, 0). \tag{12}$$

If neither state is bad, use the independent coupling. Thus every overlap edge ij has a joint law μ_{ij} with the prescribed singleton marginals and with $|E_{ij}| \le D$ throughout its support.

Root every component of the overlap forest. Draw the root bit with its Bernoulli marginal; along an oriented edge $i \to j$, draw z_j conditional on z_i according to μ_{ij}. The resulting joint distribution has vertex marginals f_i and edge marginals μ_{ij}. Every two-active spine row is therefore within the two-sided D-box. Rows with at most one active interval are immediate, and the early and late private-exit errors are $d_i(f_i - z_i)$ and $d_i(z_i - f_i)$, respectively. Since every raw load is affine in the bits, the singleton marginals also give $\mathbb{E}[\mathbf{Y}] = \mathbf{x}$ on every raw arc. Consequently $\mathbb{E}[\mathbf{c}^\mathsf{T}\mathbf{Y}] = \mathbf{c}^\mathsf{T}\mathbf{x}$ for every nonnegative cost vector, so at least one atom of this same mixture is cost-good while retaining the radius-one bound. Hence both the exact-mean and the cost-preserving constants of the class are at most one.

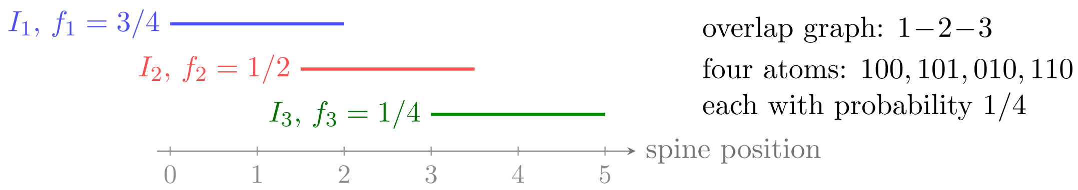

Figure 12: An explicit interval picture for the four-atom forest coupling. The first and third intervals are disjoint, so their choices need only be coordinated through terminal 2.

It remains to derive the outerplanar claim. Proposition 8.4 applied at any row shows that an outerplanar interval system has depth at most two. Its interval-overlap graph is therefore triangle-free by the Helly property. An interval graph has no induced cycle of length at least four: in such a cycle, the interval with smallest right endpoint has two cycle neighbors that both contain that endpoint and hence form a chord. Thus the overlap graph is a forest, as required.

For sharpness, use one unit-demand terminal with late share $f = 1/n$. Give its early private exit cost one and its late exit cost zero. The fractional cost is $1 - 1/n$; the early route costs one and is not cost-good, whereas the unique cost-good late route has overload exactly $1 - 1/n$. Moreover every exact-mean mixture assigns the late route its positive prescribed mass $1/n$, so that same atom is unavoidable in an exact-mean decomposition. Letting $n \to \infty$ proves sharpness for both constants. □

Let us consider a specific example with a three-interval coupling. Let $D = d_1 = d_2 = d_3 = 1$, let the overlap graph be the path 1–2–3, and take late-route shares $(f_1, f_2, f_3) = (\frac{3}{4}, \frac{1}{2}, \frac{1}{4})$. On edge 12 the state 00 is bad, so law (11) gives $(p_{00}, p_{10}, p_{01}, p_{11}) = (0, \frac{1}{2}, \frac{1}{4}, \frac{1}{4})$. On edge 23 the state 11 is bad, and (12) gives $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4}, 0)$. Rooting at terminal 2 and using the two conditional laws produces a very simple global mixture with

$$\Pr(100) = \Pr(101) = \Pr(010) = \Pr(110) = \tfrac{1}{4}.$$

Its marginals are exactly $(\frac{3}{4}, \frac{1}{2}, \frac{1}{4})$, and its two edge marginals are shown above. This four-atom example illustrates the forest-gluing proof: compatibility on shared singleton marginals is all that is needed.

The interval argument is the first member of a short interaction hierarchy. For a signed factor forest, local row laws glue exactly along the tree. If the factors have cycles but every row is pairwise, network-matrix integrality replaces gluing. If rows are ternary, a global path colouring splits the instance into two pairwise blocks. All three steps use the same interface from Definition 3.1; only the construction of the block law changes.

We now pass to the signed exact-two-path notation. For terminal i, fix its two support paths P_i^0, P_i^1, write f_i for the share on P_i^1, and let

$$g_{ai} = \mathbf{1}_{a \in P_i^1} - \mathbf{1}_{a \in P_i^0} \in \{-1, 0, 1\}.$$

Delete zero rows and identify two raw arcs when their complete signed rows are equal up to a global sign; we call each equivalence class a *row factor*. Join a factor to terminal i when its row has nonzero ith coordinate.

Remark 8 (this is not a restatement of directed series–parallel rounding)**.** Even the one-interval member need not belong to the directed two-terminal series–parallel class of [1]. After adding a common supersink and suppressing the forced last-spine path, the vertices v_0, v_1, t, T induce the forward Wheatstone orientation of K_4 minus v_0T: neither internal vertex is series-reducible and there is no parallel pair. An exhaustive series/parallel reduction search on all 3371 outerplanar topologies in the verification scope rejected every supersink completion. Theorem 8.5 is therefore an independent coupling argument rather than an application of the ASW decomposition.

8.3 One-row rounding and factor-forest couplings

Lemma 8.6 (one-row exact-mean rounding). *Let $\mathbf{w} \in [-D, D]^S$ and $\mathbf{f} \in [0,1]^S$. There is a probability law on $\mathbf{z} \in \{0,1\}^S$ with singleton marginals $\mathbb{E}\mathbf{z} = \mathbf{f}$ such that*

$$|\mathbf{w}^\mathsf{T}(\mathbf{z} - \mathbf{f})| \le D$$

in every atom.

Proof. Intersect the unit box with the equality hyperplane through $\mathbf{f}$:

$$P = \{\mathbf{u} \in [0,1]^S : \mathbf{w}^\mathsf{T}\mathbf{u} = \mathbf{w}^\mathsf{T}\mathbf{f}\}.$$

Every vertex of P has at most one fractional coordinate, since the one equality leaves at least $|S| - 1$ independent box constraints tight. Decompose $\mathbf{f}$ into vertices of P. At each vertex, round its only possible fractional coordinate j to zero or one with the probabilities preserving that coordinate. This preserves the vertex in expectation, and an atom's row value changes from $\mathbf{w}^\mathsf{T}\mathbf{f}$ by at most $|w_j| \le D$. Composing the two convex decompositions gives the required law. □

For a numerical example, let us take $\mathbf{w} = (1, \frac{3}{4})$ and $\mathbf{f} = (\frac{1}{2}, \frac{1}{3})$, and the target row value is $\frac{3}{4}$. The equality segment has endpoints $(\frac{3}{4}, 0)$ and $(0, 1)$, and $\mathbf{f} = \frac{2}{3}(\frac{3}{4}, 0) + \frac{1}{3}(0, 1)$. Rounding the first endpoint produces the explicit Boolean law

$$\Pr(10) = \tfrac{1}{2}, \qquad \Pr(00) = \tfrac{1}{6}, \qquad \Pr(01) = \tfrac{1}{3}.$$

Its marginals are $(\frac{1}{2}, \frac{1}{3})$ and its three row errors are respectively $\frac{1}{4}, -\frac{3}{4}, 0$. The construction for a factor of arbitrary arity is the same, but in a higher-dimensional box.

Theorem 8.7 (signed factor-forest coupling). *Suppose the bipartite incidence graph between row factors and terminal bits is a forest. Then the fractional flow is an exact-mean mixture of support routings $\mathbf{Y}$ satisfying $|Y_a - x_a| \le D$ on every raw arc almost surely. Consequently the exact-mean and cost-preserving constants of this class are both at most one.*

Proof. Remove deterministic bits with $f_i \in \{0, 1\}$; their raw-arc load is a common baseline and is restored unchanged afterward. Apply Lemma 8.6 to one representative of each row factor, with $w_i = d_i g_{ai}$. It supplies a factor law μ_F; the same law controls every raw arc in the factor because its signed row equals the representative or its negative. Root every component of the factor-incidence forest. Sample the root factor law (or the root terminal marginal). Whenever a sampled terminal leads to an unvisited factor, sample that factor's other bits according to μ_F conditional on the shared bit. A child factor shares exactly one previously sampled terminal with the explored component, since a second one would close a cycle. Thus the recursion is well-defined and gives every factor its prescribed law. It also preserves every singleton marginal, so every raw arc has mean exactly x_a. Every atom is within the two-sided D-box, so it remains to average the costs. □

Corollary 8.8 (balanced signed-forest coupling). *Suppose every row of g has at most two nonzero entries, the graph joining i, j when some row contains both is a forest, and for each joined pair the product $g_{ai}g_{aj}$ has one fixed sign over all rows containing both. Then the conclusion of Theorem 8.7 holds.*

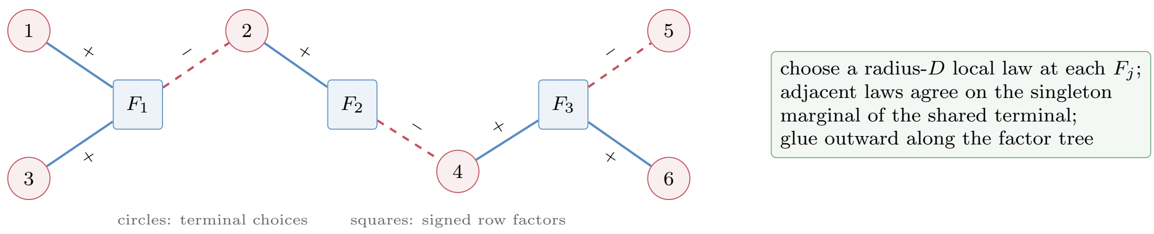

Figure 13: A signed factor forest. A square represents one complete signed row type and a circle one terminal choice; solid blue and dashed red edges indicate positive and negative incidences. Lemma 8.6 supplies a local exact-mean law at each square. Because the incidence graph is a tree, the laws can be glued successively across shared singleton marginals, which is the mechanism of Theorem 8.7.

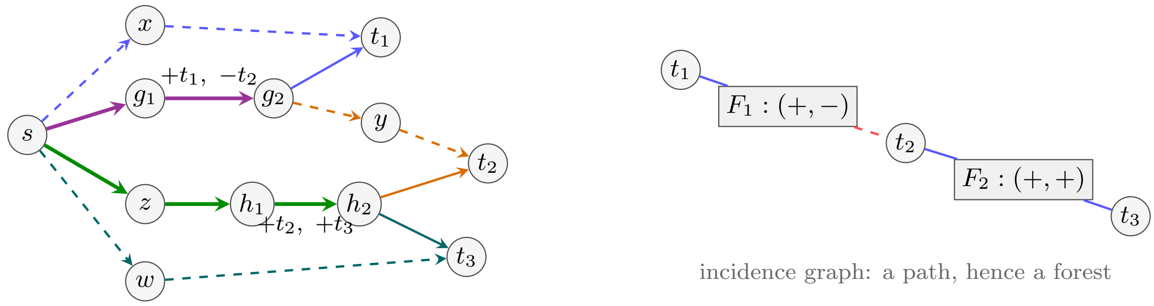

Figure 14: A DAG realization of a signed factor forest on three exact-two-path terminals. Solid arcs are selected routes, dashed arcs the alternatives. The violet segment lies on t_1's selected route and on t_2's alternative route, so its rows have support $\{1,2\}$ with signs $(+,-)$; the green segment lies on the selected routes of both t_2 and t_3, giving support $\{2,3\}$ with signs $(+,+)$. Every terminal has exactly two support paths. Right: the resulting incidence graph is the path $t_1\, F_1\, t_2\, F_2\, t_3$, a forest, so Theorem 8.7 gives exact-mean radius D.

Proof. All rows involving one pair are equal up to global sign and form one factor; the singleton rows of one variable likewise form one factor. The resulting incidence graph is obtained by subdividing the pair-interaction forest and possibly attaching singleton-factor leaves, so it is a forest. □

Remark 9 (sign consistency)**.** The assumption that the factor is acyclic is important for this gluing argument. For example, with two unit-weight variables, marginals $(\frac{3}{5}, \frac{4}{5})$, and rows $(+,+)$ and $(+,-)$, radius one on the first row forces $p_{00} = 0$, whereas radius one on the second forces $p_{10} = 0$. Writing $t = p_{11}$ gives $p_{00} = t - \frac{2}{5}$ and $p_{10} = \frac{3}{5} - t$, so the requirements demand both $t = \frac{2}{5}$ and $t = \frac{3}{5}$. The exact abstract radius is $\frac{6}{5}$, attained at $t = \frac{2}{5}$; its two row factors and two variables form a cycle of length four. Additional paths created by the DAG realization must still be considered.

8.4 The network lemma

In fact, these additional paths cannot be successfully "considered" inside the exact-two-path class: sign mixing on a pair is structurally impossible there.

Theorem 8.9 (pair-sign consistency of flows with exactly two paths)**.** *In a flow whose support has exactly two paths, for every pair of terminals $\{i, j\}$ all rows containing both have the same product sign: $g_{ai}g_{aj}$ does not depend on the arc a. Consequently the sign-consistency hypothesis of Corollary 8.8 is automatic for such flows, and the sign consistency requirement of Remark 9 admits no realization with exactly two paths.*

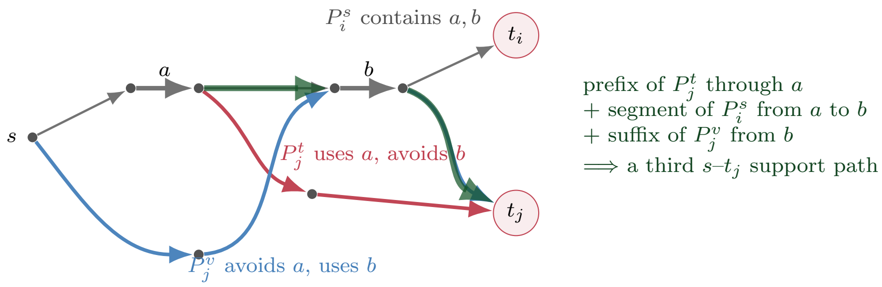

Figure 15: The forced splice in Theorem 8.9. Opposite product signs make one terminal path contain both distinguished arcs a, b, while the two paths of the other terminal contain one each. The green middle and suffix, together with the prefix through a, form a third support path to t_j, contradicting the "exactly two paths" condition.

Proof. Suppose arcs $a \neq b$ satisfy $g_{ai}g_{aj} = -g_{bi}g_{bj} \neq 0$, with $a \in P_i^s \cap P_j^t$ and $b \in P_i^u \cap P_j^v$. The sign condition says $s \oplus t \neq u \oplus v$, so exactly one of $s = u$ and $t = v$ holds; without loss of generality $s = u$, so both arcs lie on the single path P_i^s of terminal i while $t \neq v$: since $g_{aj}, g_{bj} \neq 0$, arc a lies on P_j^t only and arc b on P_j^v only. Say a precedes b along P_i^s. Concatenate the prefix of P_j^t through a, the segment of P_i^s from the head of a to the tail of b, and the suffix of P_j^v beginning with b. The endpoints of consecutive pieces agree. In a DAG this directed walk is simple (a repeated vertex would give a directed cycle) and it stays inside the support because all three pieces do. It contains both a and b, hence differs from P_j^t, which avoids b, and from P_j^v, which avoids a. Thus, we get a third support path for terminal j, contradicting exactness. If b precedes a, swap their roles. □

The flow structure gives more than pairwise sign consistency.

Lemma 8.10 (network structure of two-path differences)**.** *After zero rows are deleted, the route-difference matrix g of a support flow with exactly two paths is a network matrix and hence is totally unimodular.*

Proof. Prune the positive-support DAG so that every arc lies on a path to a designated terminal, and let $p(v)$ be the number of paths with source v. A suffix from v to a terminal injects these paths into that terminal's two paths, so $p(v) \leq 2$. Vertices with $p = 1$ form an out-arborescence. A vertex with $p = 2$ has either one predecessor with $p = 2$, or exactly two predecessors with $p = 1$; consequently every component of the $p = 2$ subgraph is an out-arborescence whose root has two incoming arcs from the $p = 1$ arborescence. Split the head of those two incoming arcs into two leaf copies for every component. This produces one rooted tree T. The suffix inside a $p = 2$ component is common to both paths of every terminal below it and cancels. Each remaining difference column is therefore the signed incidence vector of the unique tree path between two leaves of T, equivalently the difference of two root-path vectors. These are precisely the columns of a network matrix. Zero common-suffix rows do not affect total unimodularity. □

Remark 10 (the arbitrary-base three-path extension fails)**.** Lemma 8.10 does not extend by choosing an arbitrary reference path for each terminal with three support paths. On vertices $0, \ldots, 5$, take source 0, terminals $4, 5$, and arcs

$$01, 02, 05, 12, 14, 23, 34, 35.$$

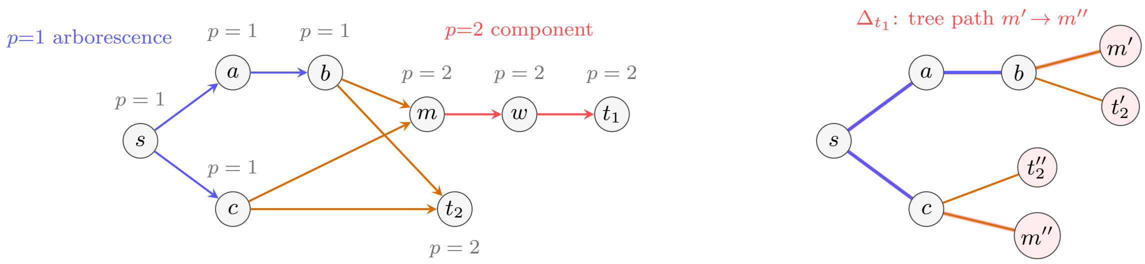

Figure 16: Lemma 8.10 on a concrete DAG. Left: the support graph with its path counts. Vertices with $p = 1$ and the blue arcs form the out-arborescence; there are two $p = 2$ components, the red chain $m \to w \to t_1$ and the single vertex t_2, each entered by exactly two orange arcs from the arborescence. The common suffix $m \to w \to t_1$ cancels in the difference of t_1's two paths. Right: splitting each component root into two leaf copies produces the tree T; the difference of t_1's paths becomes the signed tree path from m' to m'' (shaded), and the difference of t_2's paths the tree path from t'_2 to t''_2.

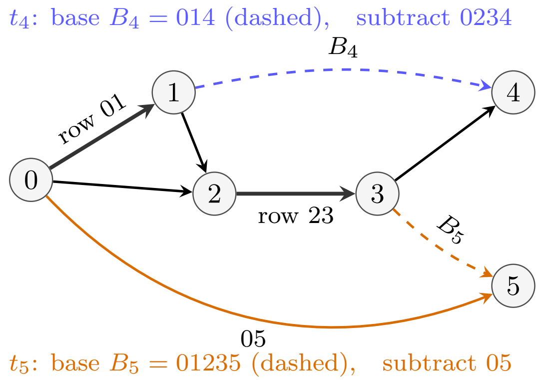

	Δ_4	Δ_5
01	-1	-1
23	1	-1

det = 2: not totally unimodular

5 of the 9 base assignments are good

Figure 17: The six vertex witness of Remark 10. Terminals 4 and 5 have three support paths each. With the base assignment $B_4 = 014$ and $B_5 = 01235$ (dashed), subtracting the bases from the paths 0234 and 05 produces difference columns whose restriction to the two highlighted arcs 01 and 23 has determinant two, so this base choice destroys total unimodularity. Note that five of the nine assignments are still good.

Each terminal has exactly three support paths. For terminal 4, subtract the path 014 from 0234; for terminal 5, subtract 01235 from 05. On rows $01, 23$ the two difference columns are

$$\begin{pmatrix} -1 & -1 \\ 1 & -1 \end{pmatrix},$$

whose determinant is two. Enumerating all topologically ordered DAGs through six vertices, we found no such minor through five vertices and found eight non-TU base assignments at six. We also found a TU base choice for every DAG in that scope. Thus the arbitrary-base extension is false, while the weaker question of whether a good base exists does go through in this example.

Theorem 8.11 (pair-row closure)**.** *Consider an exact-two-path support flow and suppose every difference row $g_{a\star}$ has at most two nonzero entries. Then the fractional flow is an exact-mean mixture of support routings $\mathbf{Y}$ satisfying*

$$|Y_a - x_a| \le D \qquad \text{on every raw arc almost surely.}$$

Consequently the exact-mean and cost-preserving constants of this class are both exactly one. No acyclicity assumption on the row-factor incidence graph is required.

Proof. The proof has one local and one global step. Locally, a pair row can make at most one Boolean corner violate the D-box, and that bad corner is cut off by one integral inequality. Globally, all these inequalities are rows of one totally unimodular network matrix, so they admit a simultaneous Boolean decomposition. We now verify the two statements.

Rows of support zero are irrelevant, and a singleton row is within D in every routing. Fix a row supported by terminals i, j. For $\ell \in \{i, j\}$ let $s_\ell = g_{a\ell}$, and consider the local coordinates

$$\xi_\ell = \begin{cases} z_\ell, & s_\ell = 1, \\ 1 - z_\ell, & s_\ell = -1, \end{cases} \qquad \phi_\ell = \begin{cases} f_\ell, & s_\ell = 1, \\ 1 - f_\ell, & s_\ell = -1. \end{cases}$$

Then the row error is

$$E_a = d_i(\xi_i - \phi_i) + d_j(\xi_j - \phi_j). \tag{13}$$

The mixed corners are automatically good; for example,

$$-d_j\phi_j \leq E_a(1, 0) \leq d_i(1 - \phi_i),$$

and both endpoint magnitudes are at most D; the other mixed corner is symmetric. Thus only 11 can violate the upper bound and only 00 can violate the lower bound. They cannot both violate, since their values differ by $d_i + d_j \leq 2D$. Hence every pair row forbids at most one of its four Boolean states.

If 11 is forbidden, then

$$d_i(1 - \phi_i) + d_j(1 - \phi_j) > D,$$

which, because $d_i, d_j \leq D$, implies $\phi_i + \phi_j < 1$. The fractional point therefore satisfies $\xi_i + \xi_j \leq 1$, the facet of the convex hull of the other three corners. If 00 is forbidden, the symmetric argument gives $\phi_i + \phi_j > 1$ and the valid facet $\xi_i + \xi_j \geq 1$. In the original z coordinates the normal of either inequality is exactly $g_{a\star}$ up to a global sign, and its right-hand side is an integer.

Let Q be the unit box together with the one such clause for every pair row that has a forbidden corner. The marginal vector $\mathbf{f}$ lies in Q. The constraint matrix of Q consists of signed rows of g and signed unit rows, so it is totally unimodular by Lemma 8.10. All right-hand sides are integral; hence Q is an integral polytope. Write $\mathbf{f}$ as a convex combination of its Boolean vertices. Each vertex avoids every forbidden corner and therefore satisfies $|E_a| \leq D$ on every difference row. Zero difference rows are common baselines, so the same statement holds on all raw arcs. The combination has singleton marginals f_i and therefore mean raw flow exactly $\mathbf{x}$. Cost averaging selects an atom of cost at most $\mathbf{c}^\mathsf{T}\mathbf{x}$.

For sharpness, one unit-demand terminal with one route share tending to zero has an unavoidable exact-mean atom, and a zero cost on that route with positive cost on the other makes it the unique cost-good atom; its normalized overload tends to one. □

For a specific example, consider one positive pair row with unit demands and $(\phi_i, \phi_j) = (\frac{1}{5}, \frac{1}{5})$, the four errors at $00, 10, 01, 11$ are $-\frac{2}{5}, \frac{3}{5}, \frac{3}{5}, \frac{8}{5}$. Only 11 is bad, and the valid clause is therefore $\xi_i + \xi_j \leq 1$. The marginal point satisfies it strictly and has the explicit good-state decomposition

$$(\tfrac{1}{5}, \tfrac{1}{5}) = \tfrac{3}{5}(0, 0) + \tfrac{1}{5}(1, 0) + \tfrac{1}{5}(0, 1),$$

as illustrated in Fig. 18. The proof of Theorem 8.11 simply imposes at once all such clauses that forbid corners; total unimodularity says their common polytope still decomposes into Boolean corners.

The theorem also identifies exactly why the abstract odd-hole idea cannot be realized in this class.

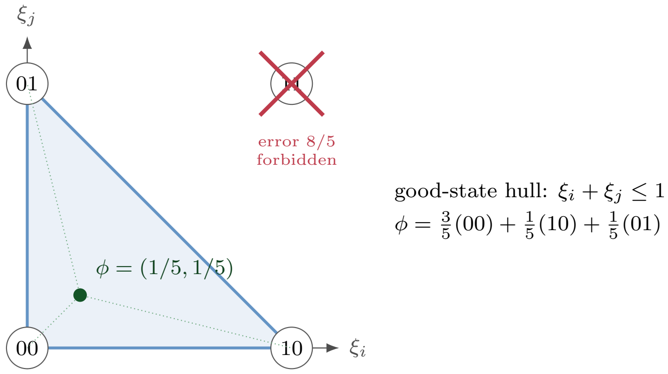

Figure 18: The numerical pair-row example as a polytope. The corner 11 is the only state with error above $D = 1$, so the good-state hull is the triangle $\xi_i + \xi_j \le 1$. The prescribed marginal $(\frac{1}{5}, \frac{1}{5})$ lies inside it with the displayed exact-mean decomposition.

Corollary 8.12 (signed cycle parity)**.** *Suppose distinct difference rows and terminal columns form a clean cycle of length n: the induced square submatrix has exactly two nonzeros in every row and column. If S is the product of the two entry signs over all cycle rows, then*

$$S = (-1)^n.$$

In particular, an all-positive odd hole is not realizable by a support flow with exactly two paths. More generally, no system consisting only of pair and singleton rows can have exact-mean radius above one; every exact-two-path obstruction above one contains a row of arity at least three.

Proof. Let B be the cycle submatrix, ordered so that row r meets columns r and $r+1$ cyclically. Its determinant has only the two cyclic matching terms, and division by the first gives

$$\frac{\det B}{\prod_r B_{rr}} = 1 + (-1)^{n-1} S.$$

The right-hand side is either zero or has absolute value two. Total unimodularity from Lemma 8.10 excludes the latter, so $S = (-1)^n$. The rest follows from Theorem 8.11. □

The unsigned interaction graph itself need not be chordal: consecutive leaf-pair paths on a star realize an unsigned hole. Its row signs have to follow Corollary 8.12. Thus the abstract all-positive C_5 value $\frac{6}{5}$ from [11] is indeed a gap between the flow and the matrix rather than a valid exact-two-path realization target. The $\frac{9}{8}$ triangle evades the theorem through its middle row of arity three, with signed row supports $110, 111, 011$.

Corollary 8.13 (the unique three-terminal high-arity core)**.** *Let a totally unimodular three-column signed matrix have one row of arity three and all remaining rows of arity at most two. Up to column switches, row signs, and a permutation of the columns, its distinct nonsingleton rows are contained in*

$$111, \qquad 110, \qquad 011. \tag{14}$$

Thus the triangle is the only maximal three-terminal matrix at the first high-arity frontier.

Proof. Switch columns to make the ternary row 111. A pair row with opposite signs on its pair would form a 2×2 minor of absolute determinant two with this row, so every pair row is positive on its support after a global row sign. Having all three pair supports would make the three pair rows the all-positive triangle matrix, again of determinant two. Hence the pair supports form a forest on three vertices and, after a permutation, are contained in the path $110, 011$. □

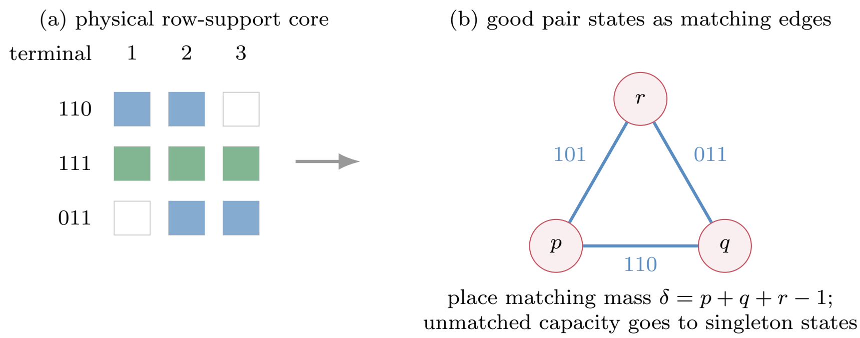

Figure 19: The local ternary core and the matching reduction in the proof of Theorem 8.14. In the middle regime $1 < p+q+r < 2$, the rank-two Boolean states are the three edges on the right. A distribution supported on the zero state, singleton states, and good pair states is equivalent to a fractional matching of mass $\delta = p+q+r-1$ under vertex capacities p, q, r.

8.5 The exact ternary envelope

The next theorem and the arity-four proof use the same two-sided view of a good-state hull. On the primal side, low-rank good states carry the marginals through a matching or an explicit convex mixture. If that hull construction fails, every crossing state is bad and therefore chooses a row with

$$E_R(z; d, f) = d(R \cap z) - \sum_{i \in R} d_i f_i > C.$$

The chosen rows are then combined on the dual side: a fractional matching, a row cover, or a Helly kernel makes their strict inequalities incompatible with the boundary equation for f. The constants $9/8$ and $C_\star$ are the thresholds at which this incompatibility first becomes universal, not isolated outputs of a cell enumeration.

Theorem 8.14 (exact local ternary envelope)**.** *Normalize $D = 1$. On three terminal columns, consider the maximal realizable core whose positive row supports are*

$$110, \qquad 111, \qquad 011. \tag{15}$$

Over all demands in $(0,1]^3$ and all marginal vectors in $[0,1]^3$, its one-sided and two-sided exact-mean radius envelopes are both exactly $\frac{9}{8}$. The two-sided upper bound is invariant under independent row signs and therefore also bounds the one-sided envelope of every realizable signing of the support topology in (15)*. Deleting rows cannot increase the envelope; the three proper maximal subtopologies have envelope one.*

Proof. The proof is a primal–dual dichotomy. Its outline goes as follows: outside the middle rank strip $1 < p+q+r < 2$, an elementary Boolean simplex already carries the marginal. Inside it, good pair states are the edges of a fractional matching. If the matching carries the excess rank, we obtain the desired mixture; otherwise all three pair states must be bad, and their three witness-row inequalities collapse to the quadratic bound (16). The two-sided argument then glues the low- and high-rank simplex laws by complementation.

Write the demands as $a, b, c \le 1$, the marginals as $(f_1, f_2, f_3) = (p, q, r)$, and let $s = p+q+r$. Private rows are always within one. We first prove the one-sided upper bound. Every Boolean state of rank zero or one is good at every radius $C > 1$, because a positive shared row then has at most one positive summand. If $s \le 1$, the usual distribution on the zero state and the three singleton states has mean (p, q, r). If $s \ge 2$, the complementary distribution on the all-one

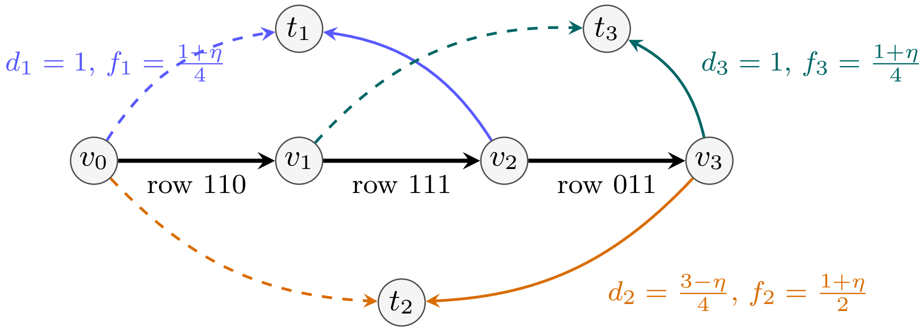

Figure 20: The ternary core as an actual flow instance: a spine with intervals $[0,1]$, $[0,2]$, $[1,2]$, whose three spine arcs carry exactly the rows 110, 111, 011. Dashed arcs are early exits, solid arcs late exits, and the annotated demands and shares are the sharp family of Theorem 8.14, whose three pair errors all equal $(3-\eta)^2/8 \to 9/8$. This is the triangle mechanism of Part I [11], now identified as the unique maximal arity-three core.

and three pair states works: the total complementary marginal is $3 - s \le 1$. Hence assume $1 < s < 2$ and set $\delta = s - 1$.

Regard the three pair states as the edges of a graph on the terminals. A distribution on the zero state, the singleton states, and the good pair states has marginal (p, q, r) precisely when the good-edge graph has a fractional matching of total mass δ with vertex capacities p, q, r. Unmatched capacity is assigned to singleton states. If all three edges are good, the maximum matching mass is

$$\min\{s/2,\ s - \max(p, q, r)\} \ge s - 1.$$

If two good edges form a path with centre j, their capacity is $\min\{f_j, f_i + f_k\}$. The second term is at least δ because $f_j \le 1$. If the first were smaller than δ, then $f_i + f_k > 1$; the missing pair state ik would have upper error below one on every row, and hence would be good. If only ij is good and, say, $f_i < \delta$, then $f_j + f_k > 1$, which similarly forces the missing state jk to be good. The other endpoint is symmetric. Thus every nonempty good-edge graph supplies the required matching.

Hull failure can therefore occur only if all three pair states are bad. Put

$$x = 1 - p, \qquad y = 1 - q, \qquad z = 1 - r, \qquad \sigma = x + z < 2 - y.$$

For the pair states $110, 011, 101$, respectively, the only shared upper errors that can exceed one are

$$A = ax + by, \qquad B = by + cz, \qquad F = ax + cz - b(1 - y).$$

Since $\min(ax, cz) \le \sigma/2$ and $a, c \le 1$,

$$\min(A, B, F) \le \min\left\{\frac{\sigma}{2} + by, \sigma - b(1 - y)\right\}.$$

The first affine function increases in b, the second decreases, and they meet at $b = \sigma/2$. Consequently

$$\min(A, B, F) \le \frac{\sigma(1 + y)}{2} < \frac{(2 - y)(1 + y)}{2} = \frac{9}{8} - \frac{(2y - 1)^2}{8} \le \frac{9}{8}. \tag{16}$$

At radius $9/8$ at least one pair state is good, and the matching argument places the marginal in the good-state hull.

For the two-sided statement, the outer ranges $s \le 1$ and $s \ge 2$ use the same two simplex distributions: on the low simplex every negative error is bounded by s, while on the high simplex every positive complementary error is bounded by $3-s$, and the remaining single opposite contribution is at most one. Again assume $1 < s < 2$. Let I be the indices of good singleton states e_i, and let J be the indices of good complementary pair states $\mathbf{1} - e_i$. A singleton can exceed one only on the lower side; if $i \notin I$, its two omitted marginals have sum greater than one, so $f_i < \delta$. A pair state can exceed one only on the upper side; if $i \notin J$, its two selected complementary marginals have sum greater than one, so $f_i > \delta$.

Let $h_i = f_i - \delta$, $P = \sum_i (h_i)_+$, and $N = \sum_i (-h_i)_+$. First, note that $N \le \delta$: it is immediate with one negative coordinate; with two, the third marginal is at most one and the two negative-coordinate marginals sum to at least δ; with three, $N = 3\delta - s = 2s - 3 \le \delta$. If $I \cap J$ contains ℓ, assign initial weights

$$\alpha_i = (h_i)_+ \quad \text{to } e_i, \qquad \beta_i = (-h_i)_+ \quad \text{to } \mathbf{1} - e_i,$$

and add

$$T = \delta - N = (2-s) - P \ge 0$$

to both α_ℓ and β_ℓ. The sign implications above make every state carrying positive weight good, and

$$\sum_i \alpha_i = 2 - s, \qquad \sum_i \beta_i = s - 1, \qquad \alpha_i - \beta_i = f_i - (s-1).$$

Thus these weights have total one and marginal exactly (p, q, r).

It remains to show that nonempty I, J cannot be disjoint. The sign implications show first that $I \cup J = \{1, 2, 3\}$. The possible bad upper errors of $\mathbf{1} - e_i$ are

$$\begin{aligned} U_1 &= b(1-q) + c(1-r), \\ U_2 &= a(1-p) - bq + c(1-r), \\ U_3 &= a(1-p) + b(1-q), \end{aligned}$$

and the possible bad lower magnitudes of e_i are

$$\begin{aligned} L_1 &= bq + cr, \\ L_2 &= ap - b(1-q) + cr, \\ L_3 &= ap + bq. \end{aligned}$$

Every nontrivial partition (I, J) contains one of the cross-pairs

$$\begin{aligned} U_1 + L_2 &= c + ap, & U_2 + L_1 &= a(1-p) + c, \\ U_2 + L_3 &= a + c(1-r), & U_3 + L_2 &= a + cr. \end{aligned} \tag{17}$$

Each sum is at most two, contradicting the two strict bad inequalities when $C > 1$. If J is empty, all three pair states are upper-bad, which is excluded at $9/8$ by (16). If I is empty, complement all states and marginals and use the same one-sided argument. This proves the two-sided upper bound.

For sharpness, let $\eta \to 0$ and take

$$a = c = 1, \quad b = \frac{3-\eta}{4}, \qquad p = r = \frac{1+\eta}{4}, \quad q = \frac{1+\eta}{2}.$$

Here $s = 1 + \eta$, and the three pair-state errors A, B, F all equal $(3-\eta)^2/8$. The all-one state is bad above this level as well, so below it all good states have rank at most one and cannot have mean of rank $1 + \eta$. The envelope therefore tends to $9/8$. Finally, Lemma 8.6 gives radius one for $\{111\}$, and Theorem 8.11 does so for $\{110, 011\}$.

For the remaining nested case $\{111, 110\}$, we can use a direct floating argument. As long as both first coordinates are fractional, move in the direction $(b, -a, 0)$ until one reaches zero or one, and express the current point as the corresponding convex combination of the two endpoints. This preserves $az_1 + bz_2$ exactly. Applying this recursively leaves at most one fractional coordinate z_j among the first two in every branch, while preserving all coordinate expectations. In such a branch, the integral first-two contribution plus the current fractional contribution still equals $ap + bq$. Apply the one-row exact-mean rounding of Lemma 8.6 to the at most two remaining fractional variables z_j, z_3, with weights d_j, c. Its change on the full row has absolute value at most $\max(d_j, c) \le 1$; on the prefix row only z_j changes, by at most $d_j \le 1$. Composing the branch distributions retains the original three marginals. Hence the nested two-row envelope is also at most one; private-row sharpness makes all three proper envelopes exactly one. □

Remark 11 (why the two-sided construction is self-dual)**.** Taking the complement $\mathbf{z} \mapsto \mathbf{1} - \mathbf{z}$, $\mathbf{f} \mapsto \mathbf{1} - \mathbf{f}$ negates every row error, exchanges s with $3 - s$, singletons with complementary pairs, I with J, and the weights α with β. The two-sided middle-regime mixture is therefore an interpolation between the two one-sided boundary solutions (mass $2 - s$ on good singletons at $\mathbf{f}$, mass $s - 1$ on complemented good singletons at $\mathbf{1} - \mathbf{f}$) and its feasibility identity $\alpha_i - \beta_i = f_i - (s - 1)$ is forced by that symmetry. Feasibility itself reduces to the inequality $\sum_i \big((s-1) - f_i\big)_+ \le s - 1$, which holds unconditionally for $\mathbf{f} \le \mathbf{1}$ and $1 < s < 2$: with one small coordinate it is immediate, with two it telescopes against $f_3 \le 1$, and with three it reads $2s - 3 \le s - 1$. The only genuine obstruction is $I \cap J = \varnothing$, which the proof refutes separately.

To show a specific finite point near the 9/8 wall, take $\eta = 1/100$, getting

$$(a, b, c) = \left(1, \frac{299}{400}, 1\right), \qquad (p, q, r) = \left(\frac{101}{400}, \frac{101}{200}, \frac{101}{400}\right).$$

The marginal rank is $101/100 > 1$, while each of the three pair states has the same critical error

$$\frac{89401}{80000} = 1.1175125.$$

Below this number every good state has rank at most one, whose convex hull cannot contain a point of rank 101/100. Letting η decrease turns this concrete obstruction continuously into 9/8.

8.6 The four-column support normal forms

Proposition 8.15 (four-column network-support normal forms)**.** *Consider a four-column network matrix with a row supported on all four columns, and ignore singleton rows and duplicate supports. Up to a permutation of the columns, its row supports are contained in one of the following five maximal families:*

$$\begin{array}{c|l} 0 & 01,\ 012,\ 0123,\ 013,\ 02 \\ 1 & 01,\ 012,\ 0123,\ 02,\ 13 \\ 2 & 01,\ 012,\ 0123,\ 023,\ 03 \\ 3 & 01,\ 012,\ 0123,\ 023,\ 23 \\ 4 & 01,\ 0123,\ 02,\ 13,\ 23. \end{array} \tag{18}$$

These are support normal forms: their realizing network rows may carry global signs, which remain relevant to a one-sided envelope. The four-terminal interval core from [11] is form 0 up to relabelling.

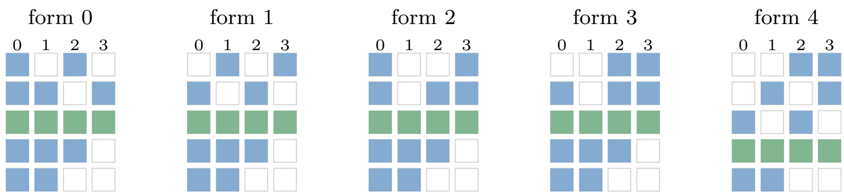

Figure 21: The five maximal four-column support normal forms of Proposition 8.15, shown side by side as support-incidence matrices. Columns are terminals $0, 1, 2, 3$; each horizontal strip is one row support, and the green strip is the row meeting all four terminals.

Proof. Represent the full-support row by a central edge of the host tree. Every other row edge lies on one of its two sides. On either rooted side, the sets of terminal paths using its edges are descendant-leaf sets and hence form a laminar family. Thus the graph whose vertices are the remaining row supports and whose edges join crossing supports is bipartite.

Conversely, a bipartition of this crossing graph splits the supports into two laminar families. The inclusion forest of each family constructs a rooted side tree with precisely those descendant-leaf sets; connect the two roots by the full-support edge and join the two copies of each terminal by its unique tree path. This realizes the supports, up to the global sign of each side row. Hence support realizability is equivalent to bipartiteness of the crossing graph.

There are ten optional nonsingleton proper supports on four columns. Testing the $2^{10} = 1024$ subsets gives 221 bipartite crossing graphs, of which 75 are inclusion-maximal. Canonicalization under the 24 column permutations gives exactly the five families in (18), with orbit counts $24, 24, 12, 12, 3$. We also verified this with a second implementation that obtained the same classification by trying all two-colourings directly and checking laminarity in each colour, without constructing a crossing graph. $\square$

Remark 12 (local versus global arity)**.** Proposition 8.15 branches the four-column parameter program into five local cells; it does not classify an arbitrary row-arity-four system with many overlapping column restrictions. The same distinction is important already at arity three. On the local four-column branch the envelope is now completely determined: Section 9 proves that the interval form attains exactly $(299 - 41\sqrt{41})/32$ and the other four forms exactly $9/8$.

8.7 From pairs to triples: Property B and the global ternary theorem

We now turn to the first unconditional global rounding theorem. Theorem 8.11 rounds any system whose rows meet at most two terminals within a single D-box, so for rows of arity at most three the natural plan is to two-colour the terminals so that no arity-three row is monochromatic. Restricted to either colour class every row has arity at most two; pair-row closure rounds each class within its own D-box, and composing the two independent mixtures preserves the mean at total radius $2D$. A hypergraph that admits a two-colouring of its vertices with no monochromatic edge is said to have *Property B*, a classical notion named after Bernstein [8]. The plan succeeds exactly when the *ternary row-support hypergraph*, whose vertices are the terminals and whose edges are the supports of the arity-three rows, has Property B.

For general three-uniform hypergraphs Property B can fail. The smallest failure is the Fano plane, seven points and seven three-point lines with any two lines sharing a point: every two-colouring of its points leaves a monochromatic line, while every three-uniform hypergraph with fewer than seven edges is two-colourable. The route through Property B therefore faces a classical wall, and the first question is whether that wall can materialize here. Proposition 8.16 shows that it cannot: the Fano incidence pattern is not realizable by paths in a tree. Proposition 8.17 then establishes Property B for a natural subclass by matching-polytope methods, and Theorem 8.18 proves it for every tree-path ternary system by path colouring.

Proposition 8.16 (the Fano ternary obstruction is not a tree-path system)**.** *There do not exist seven simple paths in a tree and seven tree edges whose path–edge incidence hypergraph is the Fano plane: each tree edge used by the three paths on one Fano line and each path using the three tree edges through one Fano point, among the seven designated row edges.*

Proof. Contract every nondesignated host edge. Because the host was a tree, the seven designated edges remain the edges of a tree, and each terminal path contracts to the three-edge simple path given by its three Fano lines. Every pair of Fano lines meets in a point. Therefore every pair of the seven remaining host edges lies together in a simple path containing exactly three edges. The contracted host tree consequently has diameter at most three. If its diameter is at most two, it is a star, where no three distinct edges form a simple path. If its diameter is three, it is a double star, and every three-edge simple path contains the unique central edge. The Fano line assigned to that edge would then have to belong to all seven point triples, whereas every Fano line belongs to exactly three. □

Proposition 8.17 (bridgeless endpoint quotients have Property B)**.** *Let a tree-path support system have row arity at most three. Contract in the host tree every edge whose row does not have arity three, replace each nonconstant terminal path by an edge joining its two endpoints in the contracted tree, and call the resulting endpoint multigraph G_3. If G_3 is bridgeless after isolated vertices and loops are discarded, then the ternary row-support hypergraph has Property B. Consequently the system has an exact-mean two-sided radius at most $2D$.*

Proof. Every remaining vertex has degree at least two. Replace each vertex v of G_3 of degree r by a cycle of r ports and attach its incident edges bijectively to those ports; for $r = 2$ use two parallel cycle edges. The resulting multigraph H is cubic. It is bridgeless: every old edge lies on a cycle lifted from G_3, and every new port edge lies on its vertex cycle.

The vector $x_e = 1/3$ belongs to the perfect-matching polytope of H. Indeed, it has degree one at every vertex. Every cubic component has even order, so an odd shore splits some component and has a nonempty odd cut. In a cubic graph $|\delta(S)| \equiv |S| \pmod 2$, so this cut has odd size; bridgelessness excludes size one and hence forces at least three cut edges. Thus every odd-cut inequality has value at least one. By the perfect-matching polytope theorem [4], write x as a convex combination of perfect matchings.

Each ternary host edge defines a three-edge cut of G_3. Taking the union of the corresponding port cycles on either shore shows that the same three old edges form a cut of H. Its shore is odd because H is cubic. Every perfect matching therefore uses an odd positive number of its three edges, while their mean incidence under the displayed decomposition is $3 \cdot (1/3) = 1$. Hence every matching of positive weight in the decomposition uses exactly one edge from every ternary cut simultaneously. Colour the old edges according to membership in any such matching. Every ternary row then has one terminal of one colour and two of the other. Loops, which correspond to columns absent from all ternary rows, may be coloured arbitrarily.

Within either colour class every restricted row has arity at most two, so Theorem 8.11 gives an exact-mean two-sided radius-D mixture. Independent product composition of the two mixtures preserves the full marginal vector and adds their error radii, proving the final claim. □

Theorem 8.18 (global ternary closure)**.** *In every tree-path support system of row arity at most three, the hypergraph formed by the rows of arity exactly three has Property B. Consequently* *every exact-two-path support flow with difference-row arity at most three is an exact-mean mixture of support routings* $\mathbf{Y}$ *satisfying*

$$|Y_a - x_a| \le 2D \qquad \text{on every raw arc almost surely.}$$

It therefore has a cost-preserving support routing with $\mathbf{Y} \le \mathbf{x} + 2D\mathbf{1}$.

More generally, maximum difference-row arity m *has exact-mean two-sided radius at most*

$$\left\lceil \frac{\lfloor 3m/2 \rfloor}{2} \right\rceil D. \tag{20}$$

Proof. The assembly is a two-line reduction once the path-colouring theorem is in place: properly colour the tree paths, group the colours into pairs, and apply pair-row closure independently inside every colour pair. The rest of the proof establishes the required colouring bound and checks that every row has restricted arity at most two.

We use the following sharp path-colouring theorem of Pritchard and Rothvoss [12, Theorem 3]: a multiset of simple paths of maximum edge load Δ in a tree has a proper colouring with at most $\lfloor 3\Delta/2 \rfloor$ colours, where paths sharing a tree edge receive different colours. For completeness, its short proof is as follows. If the host tree is a star, turn a two-edge leaf-to-leaf path into the corresponding edge of a multigraph. Turn a one-edge path at leaf u into an edge from u to a private new vertex. The multigraph has maximum degree Δ, and Shannon's multigraph edge-colouring theorem gives $\lfloor 3\Delta/2 \rfloor$ colours.

Otherwise choose a tree edge uv with neither endpoint a leaf. Contract the component strictly beyond v to v, retaining uv, and do the symmetric contraction beyond u. Map every path to both smaller trees, discarding an image with no edge. By induction, properly colour both image multisets with the common palette. The paths containing uv are exactly the paths with nonempty images on both sides; they receive distinct colours in either colouring because they share uv. A permutation of one palette makes the two colours agree path by path on this common set, and the two colourings paste. This proves the path-colouring theorem, including parallel paths.

For $\Delta = 3$, use a four-colour palette and merge colours $1, 2$ into red and $3, 4$ into blue. Three paths on one host edge have three distinct proper colours, so they cannot all lie in one two-colour part. Thus every ternary row is bichromatic. In either colour block every restricted signed row has arity at most two. Apply Theorem 8.11 separately to the two fractional terminal subflows. Product composition preserves their sum $\mathbf{x}$ in expectation and adds their raw-arc errors, giving radius $2D$. Cost averaging selects a complete atom of cost at most $\mathbf{c}^{\mathsf{T}}\mathbf{x}$.

For general m, properly colour the tree paths with $K = \lfloor 3m/2 \rfloor$ colours, partition the palette into $\lceil K/2 \rceil$ parts of size at most two, apply pair-row closure to each part, and compose. This proves (20). □

Figure 22 illustrates why four colours become two blocks. On a ternary host edge the three incident terminal paths receive three distinct proper colours. For example, colours $1, 2, 4$ become red, red, blue after the $\{1, 2\}/\{3, 4\}$ merge, so neither block sees more than two terminals on that row. The same one global path colouring works simultaneously on every host edge; the proof then rounds the red and blue terminal sets independently and pays one D-box for each. Note that a row-by-row colouring would miss this point.

Remark 13 (a global ternary route through Property B)**.** Theorem 8.18 supersedes the bridge qualifier in Proposition 8.17; the latter remains an alternative matching-polytope proof on a natural subclass. The stronger graph-only assertion is false: all ten three-edge cuts of the five-edge star form the complete three-uniform hypergraph on five edge variables, which has no Property-B colouring. The tree-path theorem uses the selected laminar cut geometry, not cographic structure alone.

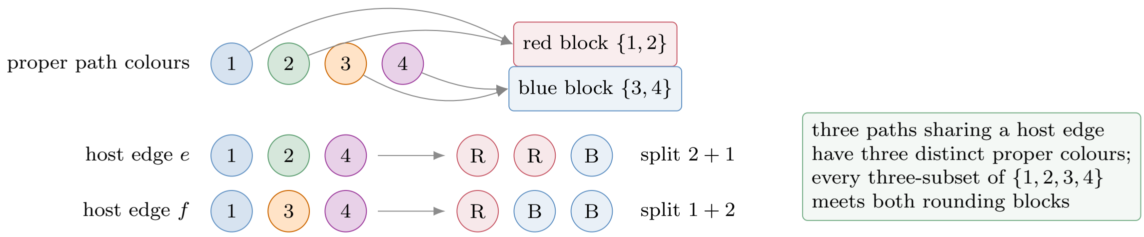

Figure 22: Why four proper path colours yield two rounding blocks in Theorem 8.18. The merge $\{1,2\}/\{3,4\}$ is fixed once for the entire host tree. Any ternary row sees three distinct proper colours, so it contains terminals from both blocks and has restricted arity at most two in either block. Different rows may realize the two possible $2+1$ orientations, as shown.

The tempting stronger exact-one rule is false: an exact fourteen-column tree-path witness with ten ternary rows has no exact transversal among its 2^{14} colourings, while 960 colourings satisfy Property B. Hence a proof must genuinely allow the two colours to occur in both 1–2 orientations on different rows. The four-colour merge does exactly that. The Shannon-tight load-three star and the fourteen-column wall both really need four proper path colours, so replacing four proper colours by three is not a hidden simplification.

8.8 Hybrid covers and the congestion transfer

The forest theorem composes with the chain machinery of Section 4 into one instance-adaptive bound.

Corollary 8.19 (hybrid block cover)**.** *Partition the terminals into blocks so that each block B is either:*

(a) a demand-divisibility chain,

(b) a set of exact-two-path terminals whose block-restricted signed rows $(d_i g_{ai})_{i\in B}$ have a factor-forest incidence graph in the sense of Theorem 8.7, or

(c) a set of exact-two-path terminals every one of whose block-restricted rows has at most two nonzero entries (with no acyclicity requirement by Theorem 8.11).

Then

$$T \;\le\; \frac{1}{D}\sum_{B}\max_{i\in B} d_i, \qquad \textit{hence} \qquad C \;\le\; \frac{1}{D}\sum_{B}\max_{i\in B} d_i.$$

Proof. A chain block carries the exact-mean two-sided decomposition of radius $(\max_{i\in B} d_i)$ from Theorem 3.3 with Lemma 3.4; a forest block carries it by Theorem 8.7 applied to the block subflow, whose support consists of the block terminals' two paths and whose signed rows are the block restrictions; a pair-row block carries it by Theorem 8.11 applied to the same block subflow, since restriction can only lower row arity. Theorem 4.1 composes the blocks. □

Remark 14 (cutting factor cycles by partitioning)**.** Restricting rows to a block deletes factor-graph edges, so a partition can *cut* every cycle of the signed factor graph: terminals of a cycle placed in different blocks stop interacting. Corollary 8.19 therefore sets a price on the cycle structure: one blockwise-maximal demand per part. Its extremes recover earlier bounds: all-singleton blocks give the trivial demand sum, equal-value blocks the scale-sum corollary, one forest block the sharp constant one. Theorem 8.11 shows, however, that cycle rank alone is not the right difficulty parameter: pair-row systems have radius one at arbitrary cycle rank. The abstract 6/5 sign wall and the all-positive C_5 both violate network-matrix minors and are unrealizable here. The first realizable cycle obstruction is instead the 9/8 triangle from [11], whose middle factor has arity three. Any quantitative ladder must therefore track high-arity excess together with factor cycle rank; the record instances are large in both parameters.

For structured theorems that provide only one suitable two-sided routing, rather than an exact-mean decomposition, the machinery of Lemma 3.5 still yields the following general transfer in six lines.

Theorem 8.20 (two-sided congestion transfers to cost at factor two)**.** *Let $\mathcal{G}$ be restriction-closed, and suppose every feasible flow* $\mathbf{p}$ *of every instance of* $\mathcal{G}$ *admits a routing* $\mathbf{y}$ *in the support of* $\mathbf{p}$ *with* $|\mathbf{y}-\mathbf{p}| \le tD\mathbf{1}$. *Then* $C(\mathcal{G}) \le 2t$.

Proof. Given an instance $(\mathbf{x}, \mathbf{c})$ of the class, let $\mathbf{q}$ minimize $\mathbf{c}^{\mathsf{T}}\mathbf{p}$ over feasible flows with $\mathbf{p} \le \mathbf{x} + tD\mathbf{1}$, with multipliers $\boldsymbol{\lambda} \ge \mathbf{0}$ as in Lemma 3.5 at $\rho = t$. The support restriction at $\mathbf{q}$ is in $\mathcal{G}$ by closure, so the hypothesis yields a support routing $\mathbf{y}$ of $\mathbf{q}$ with $|\mathbf{y}-\mathbf{q}| \le tD\mathbf{1}$. Support paths are $(\mathbf{c}+\boldsymbol{\lambda})$-shortest by Lemma 3.5(1), hence

$$\mathbf{c}^{\mathsf{T}}\mathbf{y} = \mathbf{c}^{\mathsf{T}}\mathbf{q} + \boldsymbol{\lambda}^{\mathsf{T}}(\mathbf{q}-\mathbf{y}) \le \mathbf{c}^{\mathsf{T}}\mathbf{q} + tD\|\boldsymbol{\lambda}\|_1 \le \mathbf{c}^{\mathsf{T}}\mathbf{x},$$

where the middle inequality uses the *lower* half of two-sidedness against $\boldsymbol{\lambda} \ge \mathbf{0}$, and the last is Lemma 3.5(2). The load side is $\mathbf{y} \le \mathbf{q} + tD\mathbf{1} \le \mathbf{x} + 2tD\mathbf{1}$. ☐

Corollary 8.21. *If the additive two-sided conjecture of [9] holds, that is, every instance admits a routing with* $|\mathbf{y}-\mathbf{x}| \le D\mathbf{1}$, *then* $C_{\mathrm{opt}} \le 2$. *This derives, inside the present framework, the implication stated in the introduction. A general theory of such transfers is the subject of [17]; we make no priority claim for the special case recorded here.*

Finally, we checked computationally where the record ladder leaves the planar world. The $k = 4$ family (1.1397...) from [11] is planar (though not outerplanar), while every later certified member is non-planar: the $k = 5$ instance (1.16798...), the $k = 6$ and $k = 7$ Pell families, the $k = 8$ stationary cell, the 1.28249... record of [10], and its standing refinement [2]. The mutually crossing chords of their interval systems obstruct every embedding. The ladder therefore exits the planar class immediately after its first rung. A separate planar search, however, gave Proposition 8.3, and its six-interval topology is not a member of the general-record chain. Thus the general gap now lives strictly outside planar, while inside the planar class the problem is to close the remaining gap [1.17353531974518, 2].

This leads to four frontier classes. Two-layer instances are settled at constant 1, series-parallel instances are settled at 1, planar instances are capped at 2; series-parallel sharpness is witnessed by a distinct-terminal family of parallel branches followed by a common terminal chain. Moreover, three common-point tents already contract to $K_{2,3}$, so the common-point strict-rank program cannot enter the outerplanar class. The complementary outerplanar interval regime is now settled exactly at one: its interval-overlap graph is a forest, and compatible two-variable couplings glue along that forest without losing either the exact mean or the two-sided D-box. This does not settle arbitrary multi-path outerplanar DAGs. Finally, the smallest natural restriction-closed class containing the record geometry—DAGs of subdivided crossing chords over a directed spine, or equivalently the two-order prefix systems of [10]—has no known ceiling at all. By Theorem 8.1, an upper bound for that class would be a mixture-radius theorem for two-order prefix systems, which places the class program of [10] and the upper-bound program of this note on the same object.

9 The exact local arity-four envelope

This section closes the fourth rung of the local ladder. Together with the pair-row theorem ($E(2) = 1$) and the local ternary theorem ($E(3) = 9/8$), it completes the identification of the

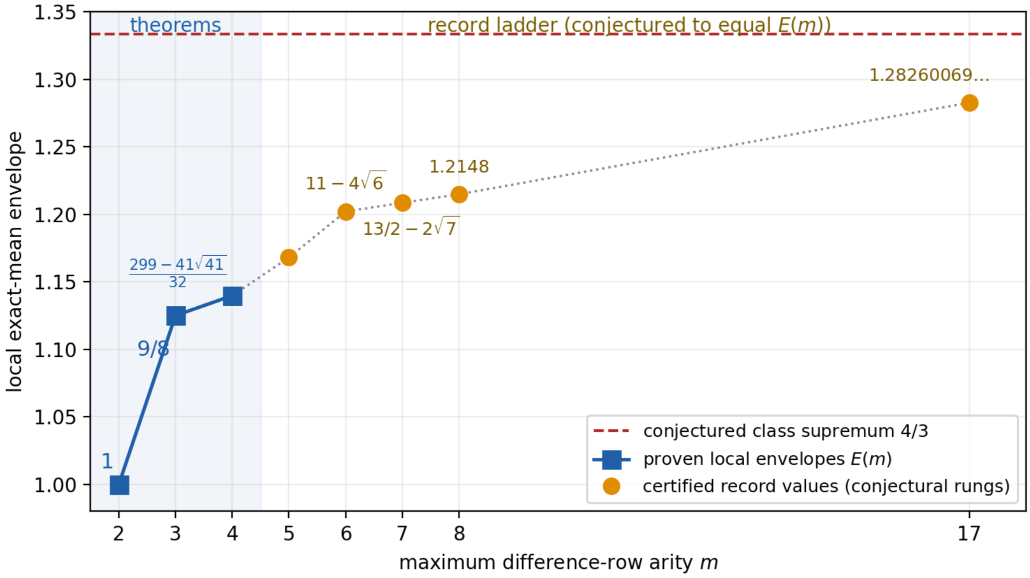

Figure 23: The local envelope ladder. Squares are theorems of this paper: $E(2) = 1$ at any cycle rank, $E(3) = 9/8$, and $E(4) = (299 - 41\sqrt{41})/32$, the last two being the constants first found as lower bound counterexamples [11, 10]. Circles are the certified record values at higher arity, conjectured to be the higher rungs of the same ladder; the dashed line is the conjectured class supremum 4/3 of Part II.

first three nontrivial local envelopes with the constants discovered as lower-bound records in [11, 10]. The conceptual novelty here is the same obstruction duality already visible at arity three: push a separator to a Boolean boundary, let each bad crossing state choose its worst row, and combine those rows into a cover of the boundary normal. The rank-one, high-row, and anchored arguments are three realizations of that scheme. The proof does not rely on exhaustive enumeration or solvers. Figure 23 shows the ladder of all our results so far.

9.1 Convention and theorem

Normalize the maximum demand to one. A positive local support system on four columns is a family $\mathcal{R} \subseteq 2^{[4]}$. Terminal i has demand $0 < d_i \leq 1$ and selected-path share $0 \leq f_i \leq 1$. For $\mathbf{z} \in \{0,1\}^4$ let

$$E_R(\mathbf{z}; \mathbf{d}, \mathbf{f}) = \sum_{i \in R} d_i (z_i - f_i), \qquad M(\mathbf{z}; \mathbf{d}, \mathbf{f}) = \max_{R \in \mathcal{R}} E_R(\mathbf{z}; \mathbf{d}, \mathbf{f}),$$

and $G_C(\mathbf{d}, \mathbf{f}) = \{\mathbf{z} : M(\mathbf{z}; \mathbf{d}, \mathbf{f}) \leq C\}$. The *positive one-sided exact-mean radius* of $\mathcal{R}$, denoted $E^+_{4,\mathrm{loc}}(\mathcal{R})$, is the least C such that

$$\mathbf{f} \in \operatorname{conv} G_C(\mathbf{d}, \mathbf{f}) \tag{E1}$$

for every $\mathbf{d}, \mathbf{f}$. The subscript records that the four columns are fixed once and for all: this is the local envelope at arity four, the quantity denoted $E(4)$ in the overview, in its positive one-sided convention. Singleton and private rows are automatic for $C > 1$.

The five maximal realizable positive row forms are

$$\begin{aligned}
\mathsf{F}_0 &: \ 01, 012, 0123, 013, 02,\\
\mathsf{F}_1 &: \ 01, 012, 0123, 02, 13,\\
\mathsf{F}_2 &: \ 01, 012, 0123, 023, 03,\\
\mathsf{F}_3 &: \ 01, 012, 0123, 023, 23,\\
\mathsf{F}_4 &: \ 01, 0123, 02, 13, 23.
\end{aligned}$$

Theorem 9.1 (exact local arity-four envelope)**.** *In the notation above,*

$$E^+_{4,\mathrm{loc}}(\mathsf{F}_0) = \frac{299 - 41\sqrt{41}}{32}, \qquad E^+_{4,\mathrm{loc}}(\mathsf{F}_j) = \frac{9}{8} \quad (1 \le j \le 4). \tag{E2}$$

Thus, if $E^+_{4,\mathrm{loc}}$ denotes the maximum over the five realizable normal forms, then

$$E^+_{4,\mathrm{loc}} = C_\star.$$

Proof map. The calculation below has a simple architecture even though its boundary algebra is unfortunately quite involved. First, strict separation reduces a failed exact-mean decomposition to one upper facet of a Boolean downset. Second, binary and ternary projection close every facet that forgets a coordinate. Third, only three genuinely four-dimensional facet families remain: rank one is handled by fractional covers, rank two by a high-row Helly argument, and the skew facets by statewise domination and two master covers. One cover-certificate lemma is the common engine: a bad crossing state chooses a witness row, and a cheap fractional cover of the facet normal makes all those witness inequalities arithmetically incompatible. The table in the next subsection lists these reductions.

The lower bound for F_0 is the sharp four-interval family. Each $\mathsf{F}_1, \dots, \mathsf{F}_4$ contains the sharp ternary core 110, 111, 011: delete coordinate 0 in F_1, coordinate 1 in F_2, coordinate 1 in F_3, and coordinate 0 in F_4 (then delete singleton and duplicate rows). Fixing the deleted share at zero (with any positive demand) therefore gives the lower bound 9/8. The rest of this section proves the matching upper bounds.

9.2 From strict separators to equality facets

In this section, we show that if the marginal point were outside the good-state hull, some nonnegative facet would separate them. Lowering the marginals onto that facet's equality face only makes bad states worse, so the whole proof happens on finitely many equality faces, nine families in all.

For $C > 1$ the good states form a downset containing all states of rank at most one. The exact downset/facet census reduces failure of (E1) to an upper facet

$$\mathbf{n}^\mathsf{T}\mathbf{z} \le b \quad (\mathbf{z} \in G_C(\mathbf{d}, \mathbf{f})), \qquad \mathbf{n} \ge \mathbf{0}, \qquad \mathbf{n}^\mathsf{T}\mathbf{f} > b. \tag{E3}$$

Lemma 9.2 (positive facet push)**.** *It suffices to exclude all-bad crossing states on the equality face* $\mathbf{n}^\mathsf{T}\mathbf{f} = b$.

Proof. Choose $\mathbf{f}' \le \mathbf{f}$ coordinatewise with $\mathbf{n}^\mathsf{T}\mathbf{f}' = b$. Such a point exists by continuity between $\mathbf{f}$ and $\mathbf{0}$. For every row and state,

$$E_R(\mathbf{z}; \mathbf{d}, \mathbf{f}') = E_R(\mathbf{z}; \mathbf{d}, \mathbf{f}) + \sum_{i \in R} d_i (f_i - f'_i) \ge E_R(\mathbf{z}; \mathbf{d}, \mathbf{f}).$$

Thus every state crossing (E3) that is bad at $\mathbf{f}$ remains bad at $\mathbf{f}'$. Notice that previously good noncrossing states need not remain good. Our anchored proof below consequently uses the full raw-row maximum M, not an inherited good-state assumption. $\square$

The cover arguments below are all instances of one elementary duality principle. Isolating it makes clear that the facet normal, rather than the particular row family, determines the target to be covered.

Lemma 9.3 (bad-state cover certificate). *Put $p_i = d_i f_i$ and suppose $\mathbf{n}^{\mathsf{T}}\mathbf{f} = b$ with $\mathbf{n} \geq \mathbf{0}$. For crossing states $\mathbf{z}$ choose witness rows $R_{\mathbf{z}}$ satisfying*

$$p(R_{\mathbf{z}}) < d(R_{\mathbf{z}} \cap \mathbf{z}) - C.$$

If there are multipliers $y_{\mathbf{z}} \geq 0$ such that

$$\sum_{\mathbf{z}: i \in R_{\mathbf{z}}} y_{\mathbf{z}} \geq \frac{n_i}{d_i} \quad (i \in [4]), \qquad \sum_{\mathbf{z}} y_{\mathbf{z}}\big(d(R_{\mathbf{z}} \cap \mathbf{z}) - C\big) \leq b,$$

then the chosen witness inequalities cannot all hold.

Proof. The face equation and the coordinate cover give

$$b = \sum_i \frac{n_i p_i}{d_i} \leq \sum_{\mathbf{z}} y_{\mathbf{z}} p(R_{\mathbf{z}}) < \sum_{\mathbf{z}} y_{\mathbf{z}}\big(d(R_{\mathbf{z}} \cap \mathbf{z}) - C\big) \leq b. \tag{E4}$$

a contradiction. □

The complete exact census has 114 downsets and 403 upper-facet cases per form. Up to coordinate permutation its nine facet families are:

facet $(n; b)$	cases over five forms	closing lemma
$(1000; 1)$	720	box constraint
$(1100; 1)$	810	binary projection
$(1110; 1)$	140	ternary projection
$(1110; 2)$	200	ternary projection
$(1111; 1)$	5	rank-one theorems
$(1111; 2)$	95	high-row Helly theorem
$(1111; 3)$	5	all-selected state
$(2111; 2)$	20	statewise domination and two master covers
$(2111; 3)$	20	anchored high-row theorem

The counts sum to $5 \cdot 403 = 2{,}015$. Section 10 describes the two independent reconstructions. An earlier numerical pass had handled 149 of these cells with solver decisions; the present proof re-derives every cell and uses none of them.

9.3 Projection and the four easy families

Lemma 9.4 (zero-section projection). *Let $J \subseteq [4]$. If the row system restricted to J has radius at most C, then the point equal to $\mathbf{f}$ on J and zero outside J belongs to* $\operatorname{conv} G_C(\mathbf{d}, \mathbf{f})$.

Proof. For a state $\mathbf{z}$ supported on J,

$$E_R(\mathbf{z}; \mathbf{d}, \mathbf{f}) = \sum_{i \in R \cap J} d_i(z_i - f_i) - \sum_{i \in R \setminus J} d_i f_i \leq E_{R \cap J}(\mathbf{z}_J; \mathbf{d}_J, \mathbf{f}_J).$$

Lift an exact-mean mixture for the restricted system by zeros. □

If a permutation of $(1110; r)$, $r \in \{1,2\}$, separates $\mathbf{f}$, the retained three-coordinate sum exceeds r. Restricting any F_j to those coordinates, deleting singletons and duplicates, gives a subtopology of the exact ternary core. Lemma 9.4 supplies a good mixture on the three coordinates; its expected rank exceeds r, so one good atom crosses the facet. The identical binary argument, using the radius-one two-column theorem, closes $(1100; 1)$. A facet $(1000; 1)$ cannot be strictly violated because $f_i \leq 1$. Finally, if $\sum_i f_i > 3$, the all-selected state satisfies for every row

$$E_R(1111; \mathbf{d}, \mathbf{f}) = \sum_{i \in R} d_i(1 - f_i) \leq \sum_i (1 - f_i) < 1,$$

and crosses $(1111; 3)$. These arguments close the four easy families without a solver.

9.4 Rank one: two fractional-cover lemmas

In this section, we show that on the face where the shares sum to one, a bad pair state pins a capacity inequality on the row that witnesses it. If a fractional cover of the four terminals by rows costs at most one, those capacities are jointly impossible: linear-programming duality turns "all pairs bad" into an arithmetic contradiction. Two explicit covers suffice for all four nontrivial forms.

On the face $\sum_i f_i = 1$, put $p_i = d_i f_i$. If a pair state ij is bad above $C > 1$, a witnessing row contains both selected terminals and gives

$$p(R) < d_i + d_j - C. \tag{E5}$$

Lemma 9.3 with $\mathbf{n} = \mathbf{1}$ and $b = 1$ says exactly that a fractional row cover of $(1/d_i)_i$ whose capacity cost is at most one contradicts an all-bad pair layer.

The following two elementary covers are the only nontrivial cases. Put $C = 9/8$ and assume $a \leq b, c, e \leq 1$.

Lemma 9.5 (three-cap cover)**.** *The three strict capacity inequalities $p(A) < a + c - C$, $p(B) < a + e - C$, and $p(F) < b + e - C$, on*

$$\begin{array}{r|l} A : \{0,1,2\} & a + c - C \\ B : \{0,3\} & a + e - C \\ F : \{0,1,2,3\} & b + e - C \end{array} \tag{E6}$$

cannot coexist with $p_0/a + p_1/b + p_2/c + p_3/e = 1$.

Proof. Let $\theta = 1/b + 1/e - 1/a$ and $\theta' = 1/c + 1/e - 1/a$. In the four regions below use the displayed weights (y_A, y_B, y_F):

$$\begin{array}{r|c} b \leq c,\ \theta \leq 0 & (1/b,\ 1/a - 1/b,\ 0) \\ b \leq c,\ \theta \geq 0 & (1/a - 1/e,\ 1/a - 1/b,\ \theta) \\ c \leq b,\ \theta' \leq 0 & (1/c,\ 1/a - 1/c,\ 0) \\ c \leq b,\ \theta' \geq 0 & (1/a - 1/e,\ 1/a - 1/c,\ \theta'). \end{array} \tag{E7}$$

They are nonnegative and cover $(1/a, 1/b, 1/c, 1/e)$; the possible slack is exactly the demand-order or θ inequality naming the row.

For completeness we give the cost check. In the first and third rows the cost increases in the free demands; setting them to one leaves respectively $1 - 1/(8a)$ and a gap $(8a(1-c)+c)/(8ac)$ below one. In the second row set first $c = 1$ and then $e = 1$. The derivative in b is

$$\frac{a}{b^2} + 1 - \frac{1}{a},$$

so the maximum is at an endpoint or at $b = a/\sqrt{1-a}$. The endpoint gaps are

$$\frac{1}{8a}, \quad \frac{(4a-3)^2}{8a}, \quad \frac{1}{8a};$$

at the stationary point the numerator of the gap is

$$8a^2 + 16a\sqrt{1-a} - 16a + 1.$$

Here $a \le \rho = (\sqrt{5}-1)/2$; concavity gives $\sqrt{1-a} \ge 1 - \rho a$, leaving $5(9-4\sqrt{5}) > 0$ at the worst endpoint. In the fourth row first set $b = 1$. The derivative in e is $(a+c-1)/e^2$. If $a + c \le 1$, set $e = a$ and then $c = a$; if $a + c \ge 1$, set $e = 1$ and use

$$\frac{\partial}{\partial c} = \frac{(a-1)(a-c^2)}{ac^2}.$$

The endpoint gaps are $1/(8a)$, $(4a-3)^2/(8a)$, or $(9-16a)/(8a)$ in the last case, where necessarily $a \le 1/2$. Every cover therefore has cost at most one, contradicting Lemma 9.3. $\square$

Lemma 9.6 (four-cap skew cover)**.** *The strict capacity inequalities on*

$$B = \{0,2,3\}, \quad F = \{0,1,2,3\}, \quad G = \{0,1,2\}$$

with respective right sides $a+e-C, b+e-C, a+c-C$ cannot coexist with $p_0/a + p_1/b + p_2/c + p_3/e = 1$.

Proof. Identify the rows (G, B, F) here with (A, B, F) in Lemma 9.5. Their capacities agree row by row $(a+c-C,\ a+e-C,\ b+e-C)$, and the supports of Lemma 9.6 contain those of Lemma 9.5: $\{0,2,3\} \supseteq \{0,3\}$ and the other two are equal. Each of the four regional weight vectors of (E7) is nonnegative and, in Lemma 9.6's system, covers every item at least as well as in Lemma 9.5's: item 2 now receives the additional weight of the middle row, and the coverage of items $0, 1, 3$ uses only memberships common to both systems. The cost bound depends only on the capacities and is therefore the computation already performed in Lemma 9.5. Lemma 9.3 gives the same contradiction. Thus the two apparent cover lemmas are one lemma applied twice. $\square$

Proposition 9.7 (the four rank-one forms)**.** *For each $\mathsf{F}_1, \ldots, \mathsf{F}_4$ and every $\mathbf{d}, \mathbf{f}$ with $\sum f_i = 1$, some pair state has error at most $9/8$.*

Proof. We list the reductions; each is literal substitution into Lemmas 9.5–9.6.

For F_1, if the minimum demand is d_0, d_2, or d_3, respectively a full-container pair among $03, 23, 03$ has error at most one because $P = \sum p_i$ is at least the minimum demand. If d_1 is minimum, use Lemma 9.5 with $(a,b,c,e) = (d_1, d_0, d_2, d_3)$ and states $12, 13, 03$.

For F_2, a minimum at 1 or 3 makes full pair 13 good. If d_0 is minimum and $p_1 \le p_3$, apply Lemma 9.5 in the original coordinate order $(0,1,2,3)$ to the bad states $02, 03, 13$: their witnessing sets are $012, 03, 0123$. The branch $p_3 \le p_1$ follows by the symmetry $1 \leftrightarrow 3$. If d_2 is minimum, put

$$(g_1, g_2, g_3) = (f_1, f_2 + f_0, f_3).$$

The ternary core on $1, 2, 3$ has a good pair at $9/8$. For each of $12, 23, 13$, its original error is no larger than the corresponding ternary error because $d_0 \ge d_2$.

For F_3, a minimum at 1 or 3 again makes full pair 13 good. For d_0 minimum and $p_1 \le p_3$, Lemma 9.6 uses the bad states $03, 13, 02$. The first two bound the full set; the 03 bound also bounds its subset 023, giving exactly the B, F, G rows of Lemma 9.6. For d_2 minimum and $p_1 \le p_3$, Lemma 9.5 uses $02, 23, 13$ after ordering its coordinates as $(2,1,0,3)$. The involution $(0\ 2)(1\ 3)$ covers the two reverse inequalities.

Finally, in F_4 both diagonal pairs 03 and 12 have the full row as their minimal container. The diagonal containing a minimum-demand coordinate has demand sum at most $1 + \min_i d_i$, whereas $P \geq \min_i d_i$; its error is at most one. $\square$

The remaining rank-one form is F_0. Its already certified sharp pair theorem gives $C_\star$: after monotonicity put $\mathbf{d} = (a, b, 1, 1)$ and $\mathbf{x} = \mathbf{1} - \mathbf{f}$, $\sum_i x_i = 3$. The six pair expressions split into the three regions $a \leq b/(1+b)$, $b/(1+b) \leq a \leq b$, and $a \geq b$. Only the middle region exceeds 9/8, with envelope

$$F(a,b) = 4a - a^2 - ab - a^2/b, \qquad \frac{b(b-4)^2}{4(b+1)} - F(a,b) = \frac{(2ab + 2a + b^2 - 4b)^2}{4b(b+1)}.$$

The one-variable maximum is attained at $b = (\sqrt{41} - 3)/4$, and the family $a = b^2$ gives $C_\star$.

9.5 Rank three: the high-row theorem

In this section, we show that on the face where the shares sum to two, only rows carrying more than 9/8 of complementary weight can hurt, and any two such rows must intersect; a Helly check gives them a common kernel. If no triple state were good, every kernel demand would have to be tiny, which inflates the complementary weights past their total of two, yielding four exact inequalities, one per intersection pattern.

Let $x_i = 1 - f_i$, $u_i = d_i x_i$, and $C = 9/8$. On $\sum f_i = 2$ we have $\sum x_i = 2$ and $u([4]) \leq 2$. A triple state omitting j has row error

$$u(R) - d_j \mathbf{1}_{j \in R}. \tag{E9}$$

Call R high if $u(R) > C$. Two high rows cannot be disjoint. In each of the five forms every pairwise-intersecting subfamily of rows has nonempty total intersection K. This is a 31-subset check per form; its exact counts are

F_j	0	1	2	3	4
intersecting families	31	23	31	23	17.

(E10)

If there is no high row, every triple state is already good. Hence suppose the high family is nonempty.

Let R^* maximize $u(R) = C + t$ over the high rows. Then $0 < t \leq 2 - C = 7/8$. If some $j \in K$ has $d_j \geq t$, omitting j makes every high row good, hence every row good. Thus an all-bad triple layer would imply $d_j < t$, and therefore $x_j > u_j/t$, for all $j \in K$.

The finite row geometry puts every pair $(\mathcal{H}, R^*)$ in one of four cases. The exact counts, over all five forms, are

case	single	direct	two-cover	staggered
count	25	226	36	10.

(E11)

We now give the complete inequalities.

If R^* is the sole high row, then $K = R^*$ and $\sum x_i > (C+t)/t > 2$. In the direct case a high row S satisfies $R^* \cap S = K$, whence

$$u(K) > 2C + t - 2 = t + \tfrac{1}{4}.$$

Using only $R^* \setminus K$ gives

$$\sum x_i > C + t + u(K)(1/t - 1) > C + \tfrac{3}{4} + \frac{1}{4t} > 2. \tag{E12}$$

In the two-cover case, high S, T satisfy $S \cap T = K$ and $S \cup T \subseteq R^*$, so $u(K) > C - t$. Consequently

$$\sum x_i > C + t + (C - t)(1/t - 1), \qquad \text{right side} - 2 = \frac{(4t-3)^2}{8t}. \tag{E13}$$

In the staggered case choose high S, T with $S \subseteq R^*$, $S \cap T = K$, and $S \cup (R^* \cap T) = R^*$. Write

$$a = u(K), \quad p = u(S \setminus K), \quad q = u(R^* \setminus S), \quad w = u(T \setminus R^*).$$

Then $a + p + q = C + t$, $q < t$, and $w > p - t$. If $a > C - t$, use (E13). Otherwise $p > t$ and

$$\sum x_i > 2C - 2a + a/t.$$

For $t \leq 1/2$ this exceeds two directly; for $t > 1/2$ its minimum at $a = C - t$ is again the right side of (E13). This contradiction proves:

Proposition 9.8 (standard rank-three faces)**.** *For every* F_j*, every* $\mathbf{d}$*, and every* $\mathbf{f}$ *with* $\sum f_i = 2$*, some triple state has error at most* $9/8$*.*

The anchored face is just as clean. Let coordinate h have coefficient two and suppose $2f_h + \sum_{i \neq h} f_i = 3$. Then

$$2x_h + \sum_{i \neq h} x_i = 2. \tag{E14}$$

The crossing triples omit $j \neq h$. If none is good, every $j \in K \setminus \{h\}$ has $d_j < t$ as above. Coordinate h, if common, has effective coefficient $2x_h \geq 2u_h$. For $t \geq 1/2$, this coefficient is at least $1/t$, so the preceding proof applies verbatim. For $t < 1/2$ every common coordinate has effective coefficient at least two. The four cases give respectively

$$2(C + t), \qquad u(R^*) + u(S), \qquad u(R^*) + u(K),$$

all greater than two. In the staggered case, if $p \leq t$ then $a > C - t$ and the cost exceeds $u(R^*) + a > 2C$; if $p > t$, using $w > p - t$ gives a cost greater than $C + a + p > 2C$.

Proposition 9.9 (anchored rank-three faces)**.** *For every form, every choice of* h*, and every* $\mathbf{f}$ *satisfying (E14), one crossing triple has error at most* $9/8$*.*

9.6 The twenty anchored rank-two placements

In this section, we show that a row that becomes larger on the selected side, or smaller on the unselected side, can only become more dangerous. This statewise order discards the irrelevant rows without making the invalid global boundary-pushing inference. The twenty realizable placements then collapse to two four-row cover systems: a nested chain and one crossed pattern. We can cover the chain greedily and use five explicit covers to settle the crossed pattern.

Here the facet is $2f_h + \sum_{i \neq h} f_i = 2$, or in complement variables

$$2x_h + \sum_{i \neq h} x_i = 3. \tag{E15}$$

The four inclusion-minimal crossing states are the three pairs hi and the triple $[4] \setminus \{h\}$. Unlike a tempting simplified proof, we must evaluate the triple on every raw row; Lemma 9.2 does not preserve goodness of noncrossing pair states.

Lemma 9.10 (statewise row domination)**.** *Fix a state* $\mathbf{z}$*. If rows* R, S *satisfy*

$$R \cap \mathbf{z} \subseteq S \cap \mathbf{z}, \qquad S \setminus \mathbf{z} \subseteq R \setminus \mathbf{z}, \tag{E16}$$

then $E_S(\mathbf{z}; \mathbf{d}, \mathbf{f}) \geq E_R(\mathbf{z}; \mathbf{d}, \mathbf{f})$ *for every* $\mathbf{d}, \mathbf{f}$*. Consequently, for this fixed state, every row dominated in the order (E16) may be deleted when computing* $M(\mathbf{z}; \mathbf{d}, \mathbf{f})$*.*

Proof. The difference is a sum of nonnegative terms:

$$E_S(\mathbf{z}; \mathbf{d}, \mathbf{f}) - E_R(\mathbf{z}; \mathbf{d}, \mathbf{f}) = \sum_{i \in (S \setminus R) \cap \mathbf{z}} d_i(1 - f_i) + \sum_{i \in (R \setminus S) \setminus \mathbf{z}} d_i f_i.$$

This comparison is made separately for each crossing state, so it says nothing about which noncrossing states remain good; this is precisely why it does not conflict with the warning after Lemma 9.2. □

Move the anchor to coordinate 0, write $(a, b, c, e) = (d_0, d_1, d_2, d_3)$, and put $p_i = d_i f_i$. A row witnessing that a crossing state $\mathbf{z}$ is bad above $C = 9/8$ must meet $\mathbf{z}$ at least twice, since one selected coordinate contributes at most one. Such a witness gives

$$p(R) < d(R \cap \mathbf{z}) - C. \tag{E17}$$

For one witness per crossing state, consider the fractional row-cover LP with target

$$(2/a,\ 1/b,\ 1/c,\ 1/e). \tag{E18}$$

Lemma 9.3 with $\mathbf{n} = (2, 1, 1, 1)$ and $b = 2$ says that a cover of this target with cost at most two contradicts the face equation; explicitly,

$$2 = 2p_0/a + p_1/b + p_2/c + p_3/e \leq \sum_R y_R p(R) < \sum_R y_R\big(d(R \cap \mathbf{z}_R) - C\big) \leq 2. \tag{E19}$$

The following lemma supplies exactly that cover. We use the notation R/I for a witness row R whose selected intersection is I.

Lemma 9.11 (two master covers)**.** *After applying Lemma 9.10 and permuting coordinates* $1, 2, 3$*, every choice of four witnesses from the five realizable forms is stronger, row by row, than one of*

$$\begin{array}{c|cccc} \mathsf{L} & 01/01 & 012/02 & 0123/03 & 123/123 \\ \mathsf{H} & 01/01 & 012/02 & 03/03 & 0123/123 \\ \mathsf{X} & 01/01 & 012/02 & 03/03 & 12/12. \end{array} \tag{E20}$$

There are respectively $3, 7, 9$ *canonical pruned choices of these three types. Each system admits a cover of (E18) of cost strictly below two.*

Proof. The first assertion is a direct inspection of the five row families at the start of this section. A constraint becomes stronger when its row support is enlarged at the same selected intersection. Pruning and then sorting the three leaves gives nineteen canonical choices: three strengthen L, seven strengthen H, and nine strengthen X. This is only a containment check on four subsets, not a numerical cell enumeration; here is the complete list.

	01/01, 012/02, 0123/03, 123/123
L	01/01, 0123/02, 0123/03, 123/123
	012/01, 012/02, 0123/03, 123/123
	01/01, 012/02, 0123/03, 0123/123
	01/01, 012/02, 013/03, 0123/123
	01/01, 012/02, 023/03, 0123/123
H	01/01, 012/02, 03/03, 0123/123
	01/01, 0123/02, 03/03, 0123/123
	01/01, 023/02, 023/03, 0123/123
	012/01, 012/02, 013/03, 0123/123
	01/01, 012/02, 0123/03, 12/12
	01/01, 012/02, 0123/03, 23/23
	01/01, 012/02, 013/03, 12/12
	01/01, 012/02, 023/03, 23/23
X	01/01, 012/02, 03/03, 12/12
	01/01, 0123/02, 03/03, 12/12
	01/01, 023/02, 023/03, 23/23
	012/01, 012/02, 013/03, 12/12
	012/01, 012/02, 013/03, 13/13

The packing dual of the cover LP has variables $v_i \geq 0$, objective $2v_0/a + v_1/b + v_2/c + v_3/e$, and one constraint $v(R) \leq d(I) - C$ for each entry R/I. Every X packing is an H packing: its last constraint and the 03/03 constraint give

$$v_0 + v_1 + v_2 + v_3 \leq (a + e - C) + (b + c - C) \leq b + c + e - C,$$

because $a \leq 1 < C$. It therefore remains to cover L and H.

For L, let $m = \min\{b, c, e\}$. If $a \leq 2m$, weight the full row 0123 by $2/a$; it covers (E18) and costs $2(a + e - C)/a \leq 2$. Suppose $a \geq 2m$. On the nested rows $01 \subset 012 \subset 0123$ use the greedy weights

$$y_{0123} = 1/e, \quad y_{012} = (1/c - 1/e)_+, \quad y_{01} = \big(1/b - \max\{1/c, 1/e\}\big)_+. \tag{E21}$$

Their sum is $1/m \geq 2/a$, so they cover the anchor as well as all three leaves. The cost check has only the order of b, c, e in it. If $m = e$, the cost is $(a + e - C)/e \leq 1$. If $m = c$, it is

$$2 - (C - a)/c - c/e \leq 2.$$

If $m = b$ and $e \leq c$, it is $2 - b/e - (C - a)/b \leq 2$. Finally, if $m = b$ and $c \leq e$, it is

$$3 - b/c - c/e - (C - a)/b \leq 3 - \left(2\sqrt{b} + \frac{1}{8b}\right) \leq \frac{3}{2},$$

where $e \leq 1$, $C - a \geq 1/8$, and the last one-variable expression is minimized at $b = 1/4$.

For H, call its rows $A = 01$, $B = 012$, $Q = 03$, and $F = 0123$. The following five covers are exhaustive; the entries are (y_A, y_B, y_Q, y_F).

$$(\alpha, \beta, \gamma, \delta) := (1/b - 1/c,\ 2/a + 1/c - 1/b - 1/e,\ 2/a - 1/b,\ 1/b + 1/e - 2/a)\,.$$

region	cover
$b \geq c, \quad 1/c + 1/e \geq 2/a$	$(0, 1/c, 1/e, 0)$
$2/a \geq 1/b + 1/e, \quad 2/a \geq 1/c + 1/e$	$(0, 2/a - 1/e, 1/e, 0)$
$b \leq c,\ e \leq c, \quad 1/b + 1/e - 1/c \geq 2/a$	$(1/b - 1/c, 0, 1/e - 1/c, 1/c)$
$b \leq c \leq e, \quad a \geq 2b$	$(1/b - 1/c, 0, 0, 1/c)$
$b \leq c,\ a \leq 2b, \quad 1/b + 1/e \geq 2/a, \quad 2/a + 1/c \geq 1/b + 1/e$	$(\alpha, \beta, \gamma, \delta)$.

(E22)

Indeed, if $b \geq c$, the first row applies or its failure implies the second. If $b \leq c$ and $a \geq 2b$, use the fourth row when $c \leq e$ and the third when $e \leq c$. If $b \leq c$ and $a \leq 2b$, failure of $1/b + 1/e \geq 2/a$ gives the second row; if that inequality holds, use the fifth row unless its last inequality fails, in which case necessarily $e < c$ and the third row applies.

For the first four covers the costs are at most $7/4$. Here are the short checks, included to make the cover lemma self-contained. The first cost is $2+(a-C)(1/c+1/e)$. For the second, its condition gives $a \leq 2ce/(c+e)$, and substitution leaves the nonnegative gap $9(c+e-2ce)/(8ce)$ below $7/4$. In the third region the cost increases with a; at $a = 1$ its gap is

$$\frac{1}{8}\,(1/b + 7/c + 1/e - 10) \geq 0,$$

because the regional inequality gives $1/b + 1/e - 1/c \geq 2$ and $c \leq 1$. The fourth cost increases with a, e and is at most $2 - 1/(8b) \leq 7/4$, since $a \geq 2b$ forces $b \leq 1/2$.

For the fifth cover the cost is

$$V = 8 + b/e - b/c - a/b - a/e - \frac{8b+9}{4a}. \tag{E23}$$

If $a \leq b$, put $r = 1/c$ and $s = 1/e$. The last regional inequality says $s \leq 2/a + r - 1/b$; maximizing first in s and then in $r \geq 1$ gives

$$V \leq \frac{(9-2a)(2a-1)}{4a} \leq \frac{7}{4}.$$

If $a \geq b$, decreasing $1/c$ and $1/e$ can only increase (E23), so set $c = e = 1$. For $b \leq 1/2$ the resulting expression increases on $a \leq 2b$ and is at most $7/4$ at $a = 2b$, $b = 1/2$. For $b \geq 1/2$ it increases on $a \leq 1$, and hence

$$V \leq \frac{19}{4} - 2b - \frac{1}{b} \leq \frac{19}{4} - 2\sqrt{2} < 2. \tag{E24}$$

This proves the cover bound for both masters. □

Proposition 9.12 (anchored rank-two faces)**.** *For all twenty realizable form/anchor placements in the exact facet census, some crossing state has one-sided error at most* $9/8$*.*

Proof. If all four crossing states were bad, choose an undominated witness row for each. Lemma 9.11 supplies a cover of cost below two, and Lemma 9.3 gives the desired contradiction. □

9.7 Completion of the proof

The exact downset enumeration is exhaustive. The easy and projection arguments close $1,875$ cases; the rank-one theorems close five; Proposition 9.12 closes twenty; Proposition 9.8 closes ninety-five; and Proposition 9.9 closes twenty. These disjoint classes total $2,015$, so no separating facet remains. At threshold $9/8$ this proves the upper bound for $\mathsf{F}_1, \ldots, \mathsf{F}_4$. For F_0, the only boundary capable of exceeding $9/8$ is the already solved rank-one facet, whose exact value is $C_\star$; all new lemmas hold at $9/8 < C_\star$. The sharp lower families noted after Theorem 9.1 give equality in (E2).

Remark 15 (scope)**.** Theorem 9.1 identifies the first three positive local rungs: $E(2) = 1$, $E(3) = 9/8$, and $E(4) = C_\star$. It does not assert that arbitrary interacting row-arity-four systems have radius $C_\star$, nor does it settle mixed row signs or the two-sided convention.

10 Exact verification

Every explicit decomposition construction in this note has been executed in exact rational arithmetic. The representation matches the claim: graph certificates rebuild the graph from its arc list and rediscover all simple source–terminal paths by depth-first search, whereas abstract matrix, interval, and cell certificates rebuild those native objects directly. Each certified finite computation also has a separately written second implementation. The symbolic proofs provide the unrestricted quantifiers; the finite computations are regressions and exact witnesses, not substitutes for those proofs. Appendix A lists the certificates as released in the companion repository[3].

Scale-sum product on the Part II $k = 17$ record. The refined record of [10, 2] has nine distinct rational demand values with exact sum $R = \frac{99759003}{12500000} = 7.98072024$. The generator partitions seventeen terminals into nine equal-value classes, builds the exact class mixtures of Lemma 3.2, and forms the independent product: 3456 atoms of total weight one whose mean equals the fractional flow exactly on every graph arc. Every atom satisfies the two-sided bound of Corollary 4.3; the attained extreme deviations are $\max_{y,a}(y_a - x_a) = 6.6906\ldots$ and $\max_{y,a}(x_a - y_a) = 1.9429\ldots$, both well inside R.

Round-down calibration on the $16/15$ instance. On the seven-vertex counterexample [11], at $r = 2$ all three demands round down to a single value; expensive-path stripping, the box-eligible rounded mixture (four atoms), and the restoration were executed exactly: the chosen routing has true cost $30 < 58 = \mathbf{c}^{\mathsf{T}}\mathbf{x}$ and satisfies $\mathbf{y} \leq 2\mathbf{x} + D\mathbf{1}$, calibrating Theorem 6.2 end to end. The same program verifies the labelled-marginal wall of Proposition 5.1 on its six-vertex graph.

The one-band lower bound and relaxed wall. For Proposition 7.5, a new generator re-enumerates all $2^{17} = 131072$ raw-DAG routings of the standing refinement, checks the 8763 strict-sublevel routings and the 18-atom entry mixture, and derives $K = 2C_{17}$ from the band endpoint $h = 1/2$. A second implementation rebuilds every spine and private-row load directly from the interval description. For Proposition 7.6, both implementations substitute the original certificate from [10], including its 15-atom primal and 11-arc dual, exactly. They also enumerate the complete KKT support-routing set, substitute the new 16-atom exact-mean mixture, and check the predecessor separator route by route. A further generator certifies the full 15-segment relaxed frontier through all 95 relevant routing levels; its second implementation reconstructs the interval rows and private exits directly and substitutes every segment's primal and dual against all eligible routings.

Recoupling evidence. Admitting every class routing within one class demand of its class flow and intersecting with the record's cost facet, exact convex-hull entry for the $k = 17$ instance [10] occurs at $1.282494799234\ldots$, only $1.25 \cdot 10^{-9}$ above that instance's certified constant: recoupling the demand classes recovers essentially the entire gap on this instance. This is evidence about one instance, not a merger theorem.

Hub tightness and the series-parallel regression. For $m = 2, \ldots, 6$, the hub-family generator reconstructs the complete raw DAG, rediscovers all m paths of every terminal, enumerates all m^{m+1} assignments, and verifies the exact entries $\frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \frac{4}{5}, \frac{5}{6}$. A second implementation separately enumerates the balanced mixtures and their terminal–hub marginals. A

[3]`https://github.com/snikolenko/unsplittable-flows`

separate series-parallel generator performs the same census on the distinct-terminal parallel-branch–plus–chain family, verifies exact mean on every branch and tail arc, and checks the explicit two-terminal series/parallel reduction. Its second implementation rediscovers every raw path and all m^{m+1} routings without importing generator helpers. On the support of the $\frac{16}{15}$ instance of [14], one structural certificate displays a K_4 subdivision with branch vertices s, u, v, w; its independent audit suppresses the three degree-two terminals and obtains K_4. Thus neither the support nor its supersink completion is series-parallel, as required by the exact series-parallel theorem.

The planar record. The planar-record generator rebuilds the 21-arc DAG of Proposition 8.3, rediscovers exactly two simple paths for each terminal, enumerates all 64 routings, and substitutes the exact seven-atom mixture and rank separator. A separately written implementation reconstructs the interval rows and private exits and obtains the same hull entry. Planarity is certified independently by a complete rotation system: the generator obtains seven faces and Euler characteristic two, while a solver-free audit retraces every dart and recomputes the genus. The same pair stores and independently suppresses an augmented $K_{3,3}$ Kuratowski subdivision, proving non-outerplanarity, and counts the seven directed sinks that exclude the ASW two-terminal orientation.

The outerplanar wall. The structural generator contracts the two sides of a common spine arc and checks the resulting $K_{2,3}$ branch sets on three raw examples, including the certified planar record. A second implementation rebuilds connectivity and all six branch-set adjacencies directly from each raw edge list. The symbolic proof in Proposition 8.4 supplies the unrestricted quantifier; the examples are executable convention calibrations, not a finite substitute for that proof.

The outerplanar interval theorem. The forest-coupling generator exhausts 33544 interval topologies with up to six spine positions and five terminals. All 3371 outerplanar topologies have depth at most two. On the 5401 depth-two topologies it builds 16203 rational flows and exact joint laws containing 203363 positive atoms in total. Every atom is checked on every spine and private row with rational arithmetic; all three local regimes (delete 00, delete 11, or use the independent law) occur thousands of times. A separately written implementation uses a different 9934-topology census and a recursively ordered state representation, and independently verifies 3604 exact laws with 40065 positive atoms. Both programs also rebuild the raw two-path sharpness family through $n = 30$. These finite sweeps are regressions for the symbolic proof of Theorem 8.5, not a replacement for its unrestricted quantifier. A separate pair of series/parallel reducers, using opposite elimination orders, checks all 3371 outerplanar topologies in the same scope and finds that none of their supersink completions is directed two-terminal series–parallel.

The signed factor-forest theorem. The exact generator builds 3600 signed factor forests with 8807 factors, exercising every factor arity from one through five. It constructs the local equality-slice laws and the global conditional law using only rational arithmetic, then verifies 136397 positive global atoms and 931530 row–atom inequalities. A separately written implementation uses the reverse local pivot order and the junction-tree product formula on 1800 fresh instances, checking 41845 global atoms and 141733 row–atom inequalities. The pairwise gauge specialization has a further exact generator on 9915 rational instances and 3722592 row–atom inequalities, plus an independent audit on 2400 instances and 1709625 more inequalities. These finite checks exercise the constructions; the unrestricted quantifier follows from the symbolic equality-slice and forest-gluing proof of Theorem 8.7.

For Theorem 8.11, a generator exhausts all 33856 topologically ordered DAGs through six vertices. It checks 7070 pairs of terminals having exactly two raw support paths and 12292 common difference rows, with no pair-sign violation. A further 1680 exact rational root-compatible pair-row instances exercise the forbidden-corner construction and extract 2233 clauses. The separate implementation rebuilds 28 pinned canonical DAGs spanning $k = 3, \dots, 9$, discovers all 336 source–terminal paths from their raw arc lists, and checks 2563 positive mixture atoms on all 504 raw arcs. It also recomputes the all-positive cycle determinants $2, 0, 2, 0, \dots$ through length eleven. For Corollary 8.13, an exact minor census checks all 256 signed three-column candidates, finds 28 totally unimodular systems ($4, 12, 12$ with zero, one, or two pair rows), and reduces them to the three displayed normal forms; a separately written determinant audit rechecks all 28. For Theorem 8.14, the generator enumerates all 102 Boolean threshold separators of the three-cube for each of four row topologies and both sign conventions, giving 816 exact QF–NRA infeasibility decisions at radius 9/8. The final proof is solver-free: a symbolic certificate checks the matching identities, the quadratic envelope, all four two-sided cross sums, and the sharp family. A separate symbolic certificate checks that the nested-case floating direction preserves both row sums and records the terminal one-factor rounding step; all 102 separators are independently infeasible at radius one in both nested sign conventions. The independent rational-grid implementation additionally checks 26973 instances per convention directly at radius one. A separate implementation uses no semialgebraic solver; on a denominator- eight share grid and a four-value demand grid it directly enumerates good Boolean states and tests hull membership by rational row reduction. All 215784 instances pass.

For Proposition 8.15, the crossing-graph generator tests all 1024 support subsets and obtains 221 compatible subsets, 75 labelled maxima, and five normal forms. The second implementation never constructs a crossing graph: it tries every two-colouring of each support family and checks laminarity separately on both sides, reproducing all three counts and the five canonical forms. For Remark 10, the raw-DAG generator checks 33864 DAGs, 11652 three-path terminal occurrences, and 38892 base assignments; exact minors find eight non-TU assignments, first at six vertices. An independent iterative path finder and Ghouila–Houri test verifies all 630 assignments through five vertices and the six-vertex witness, where four of nine base choices are non-TU and five remain TU. For Proposition 8.16, one implementation checks all 23 unlabelled eight-vertex trees and all 115920 assignments of the seven Fano lines to host edges; no assignment realizes more than three of the seven point paths. The second implementation does not enumerate assignments: it checks the host-tree diameter trichotomy and the central-edge argument in the proof. The Property-B sampling search in Remark 13 is deliberately retained at the discovery tier. Its exact-one wall is rebuilt from the rooted-tree parent array and all fourteen endpoint pairs; exhaustive enumeration checks all 16384 colourings, finding zero exact transversals and 960 Property-B colourings. A second implementation proves exact-one infeasibility with pseudo-Boolean constraints and independently enumerates the same 960 Property-B models. Finally, a structural ledger independently recomputes row arity and factor cycle rank on the ten certified interval records: every record above one has a row of arity at least three, as required by the theorem.

The complete annotated list of certificate artifacts, with the companion repository and a one-line description of what each certifies, is in Appendix A. Every theorem family has both a generator and a separately written audit; all artifacts are at generator plus second-implementation tier and have not yet been labelled outside-audited.

11 Discussion and open problems

The theorems of this work move the unrestricted universal-constant question to one quantitative statement about a factor-two window of demand values. In the language of the paper, two kinds of irreducibility remain:

- *arithmetic irreducibility* means many incomparable demands inside one band; it controls whether any universal constant exists;
- *interaction irreducibility* means overlapping high-arity row cores; it controls the sharp constants of structured two-path classes.

The first is isolated by the merger reduction, the second by the local envelope ladder.

The resulting map is more informative than individual constants. Whenever a system can be split into blocks admitting local exact-mean laws, the problem is solved by composition. Whenever it cannot, separation exposes a minimal core and bad-state/cover duality puts a price on that core. The unrestricted problem is therefore no longer an abstract "cost versus congestion" question: it asks whether every dense one-band core has bounded exact-mean radius. The structured problem asks how local high-arity cores compose. These are the two major open questions below; the remaining problems refine one or the other.

Open Problem 11.1 (the band constant). *Determine whether $K^* < \infty$; alternatively, prove* $\mathrm{FT}(K, \beta)$ *for any explicit finite pair. Either result gives a universal cost-preserving additive constant. Currently* $K^* \geq 2.565201380250176725440299636336\ldots$, $\mathrm{FT}(2m)$ *holds on m-value bands, and no upper bound on K^* is known for unbounded m.*

Open Problem 11.2 (the band ladder). *Inside a single band, mirror the lower-bound methodology of Parts I–II [11, 10] beyond the inherited threshold $2C_{17}$. Exact-mean extraction from an existing one-band record merely doubles its D-normalized hull entry, so a genuinely larger lower bound on K^* requires a new record instance. For the relaxed program, compute the exact Pareto function $\beta_{\min}(K)$ and target a class theorem minimizing $K + \beta_{\min}(K)$. The complete* 15*-segment KKT frontier is the first mandatory calibration. On that witness exact mean is optimal, so a persistent positive-defect improvement requires genuinely different geometry rather than a finer sweep of the same point.*

Open Problem 11.3 (label-correct rounding). *Proposition 5.1 identifies the mechanism obstructing all round-up routes: aggregate decompositions lose terminal labels, and the resulting mean error is the within-class covariance between demands and route marginals. Quantify it: does every one-band flow admit an aggregate decomposition whose relabelled mean error is $O(h)$? A positive answer is* $\mathrm{FT}(O(1))$.

Open Problem 11.4 (sharper shifted chains). *Is the family* (1) *optimal among round-down schemes with level-partition structure? Mixed bases $b > 2$, instance-adaptive grids (via κ_{div} on the rounded values), and combining stripping with the two-sided chain bound are all unexplored.*

Open Problem 11.5 (outerplanar beyond two exits). *Determine the exact constant for arbitrary single-source outerplanar DAGs. Theorem 8.5 settles two-exit interval spines at one, but multi-path outerplanar supports need not reduce to a forest of binary overlap constraints. This is the first topological extension in which the new forest coupling can be tested without encountering the planar lower-bound geometry of Proposition 8.3.*

Open Problem 11.6 (interacting ternary cores). *Theorem 8.11 closes the entire pair-row exact-two-path class at one, regardless of cycle rank, and Theorem 8.14 determines the exact* $9/8$ *envelope on three columns. Determine the sharp constant when arbitrarily many columns* *and overlapping rows of arity at most three are allowed. The three-column classification does not reduce such a system to one local core: several ternary rows can couple through shared columns. Thus the global arity-three envelope lies in* $[9/8, 2]$ *by Theorems 8.14 and 8.18. Interacting floating interval searches remain below* $9/8$*, but it is unknown whether interaction can beat the local value.*

Two further remarks. First, all product mixtures here are two-sided (Remark 3); the lower-bound side has not been exploited at all and connects to the two-sided conjecture of [9]. Second, the endpoint statements of the asymmetric-rounding program (radius-one UBP endpoints imply $C \leq 2$; capped LBP endpoints imply $C \leq 3$) remain the sharpest sufficient conditions of augmenting type; the present results make their remaining content precise, since instances outside a dense band are already covered unconditionally.

Looking back, all proofs in this work revolve around one principle applied twice. If we round a fractional flow to a probability law whose mean is the flow exactly, cost preservation comes for free; on every restriction-closed class the optimal constant is precisely the least radius at which such laws exist (Theorem 8.1). Along the arithmetic axis this turns divisibility into a resource: chains cost one maximal demand, optimal chain covers are a matching computation, rounding down converts the remaining gap into a multiplicative factor, and the unrestricted question collapses, within a factor of two, to a single dyadic band. Along the geometric axis it turns interaction structure into a resource: pair interactions of exact-two-path flows cost a single D-box at every cycle rank because their difference matrices are network matrices, the first genuine obstructions are the three- and four-column cores with exact envelopes $9/8$ and $(299 - 41\sqrt{41})/32$, and path colouring converts bounded row arity into the first finite constants for unbounded classes.

The lower-bound program of [11, 10] and the present upper-bound theory now meet in the middle. The constants $9/8$ and $(299 - 41\sqrt{41})/32$, first encountered as counterexample families, are identified here as the exact local envelopes $E(3)$ and $E(4)$, computed by the same bad-state/cover duality that proves the upper bounds; the certified records at higher arity are conjectured to be the next rungs of the same ladder, with the $4/3$ supremum conjectured in [10] as the horizon. What remains open is also far more concentrated now: we have one quantitative question about dense one-band cores (the band constant K^*) and one compositional question about interacting high-arity cores. Every finite claim along the way carries an exact rational certificate verified by at least two independently written programs, so both frontiers can be attacked, by people or by machines, from a fully verifiable base.

References

[1] Muhammed M. Almoghrabi, Martin Skutella, and Philipp Warode. Integer and unsplittable multiflows in series-parallel digraphs. arXiv:2412.05182, 2024.

[2] Alfredo De la Fuente. Exact 17-terminal common-point certificate for the lower bound 1.28260069 . . .. Merged pull request #1 to the companion repository `snikolenko/unsplittable-flows`, `https://github.com/snikolenko/unsplittable-flows/pull/1`, 2026. Certificate `certificates/commonpoint_k17_refined_exact.json`.

[3] Yefim Dinitz, Naveen Garg, and Michel X. Goemans. On the single-source unsplittable flow problem. *Combinatorica*, 19(1):17–41, 1999.

[4] Jack Edmonds. Maximum matching and a polyhedron with 0, 1-vertices. *Journal of Research of the National Bureau of Standards, Section B*, 69B(1–2):125–130, 1965.

[5] Stavros G. Kolliopoulos and Clifford Stein. Approximation algorithms for single-source unsplittable flow. *SIAM Journal on Computing*, 31(3):919–946, 2002.

[6] Siyue Liu and Victor Reis. Weighted chairman assignment and flow-time scheduling. In *Proceedings of the 17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, volume 362 of *LIPIcs*, pages 98:1–98:20, 2026. arXiv:2511.18546.

[7] Maren Martens, Fernanda Salazar, and Martin Skutella. Convex combinations of single source unsplittable flows. In *Algorithms – ESA 2007*, volume 4698 of *Lecture Notes in Computer Science*, pages 395–406, 2007.

[8] Edward W. Miller. On a property of families of sets. *Comptes Rendus des Séances de la Société des Sciences et des Lettres de Varsovie, Classe III*, 30:31–38, 1937.

[9] Sarah Morell and Martin Skutella. Single source unsplittable flows with arc-wise lower and upper bounds. *Mathematical Programming*, 192(1–2):477–496, 2022.

[10] Sergey Nikolenko. Deletion-star ceilings and a 1.2825... lower bound for cost-preserving single-source unsplittable flows. Zenodo preprint, `https://doi.org/10.5281/zenodo.21716713`, 2026.

[11] Sergey Nikolenko. A planar lower bound $1.1397\ldots > 9/8$ for cost-preserving single-source unsplittable flows. Zenodo preprint, `https://doi.org/10.5281/zenodo.21701162`, 2026.

[12] David Pritchard and Thomas Rothvoß. Edge-colouring hypergraphs properly (covering with matchings) or polychromatically (packing covers). arXiv:1009.6144v1, 2010. Theorem 3 is the tree-path colouring bound used in Part III.

[13] Matthew Protti. Four-terminal planar case of the dinitz-garg-goemans cost conjecture. VibeMathed `https://vibemathed.com/problem/planar-four-terminal-dgg`, GitHub `https://github.com/matthewprotti/planar-ssuf-four-terminal-bound/releases/tag/v0.1.0`, 2026.

[14] Dmitry Rybin. A counterexample to Goemans' cost-preserving unsplittable flow conjecture. Announcement, 22 July 2026; record and independent restatement at `https://mathlab.drummerduck.com/p/goemans-unsplittable-flow` and `https://vibemathed.com/problem/dinitz-garg-goemans-unsplittable-flow`, 2026. Seven-vertex instance with $c^{\mathsf{T}}x = 58$ and every congestion-good routing of cost at least 60; not peer-reviewed.

[15] David B. Shmoys and Éva Tardos. An approximation algorithm for the generalized assignment problem. *Mathematical Programming*, 62(1–3):461–474, 1993.

[16] Martin Skutella. Approximating the single source unsplittable min-cost flow problem. *Mathematical Programming*, 91(3):493–514, 2002.

[17] Chaitanya Swamy, Vera Traub, Laura Vargas Koch, and Rico Zenklusen. Unsplittable cost flows from unweighted error-bounded variants. In *Proceedings of the 2026 Symposium on Simplicity in Algorithms (SOSA)*, pages 512–523, 2026. arXiv:2510.21287.

[18] Vera Traub, Laura Vargas Koch, and Rico Zenklusen. Single-source unsplittable flows in planar graphs. arXiv:2308.02651, 2023.

A Numerical certificates

All certificates live under `certificates/` in the companion repository

`https://github.com/snikolenko/unsplittable-flows`

together with the verifier programs under `verify/` and `src/`. Every certificate is exact rational, rebuilt from raw arc lists with paths rediscovered by depth-first search, and paired with a separately written second implementation; paired filenames use several historical suffix conventions (`_second`, `_independent`, `_audit`). The list below groups the artifacts by the section they support and states what each one certifies.

Scale-sum and round-down (Sections 4–6).

- `goalA_scale_sum_k17_exact_20260815`: the nine-class mean-exact product on the $k = 17$ record; 3456 atoms, mean equal to the fractional flow on every raw arc, scale sum $R = 99759003/12500000$.
- `goalA_scale_merge_k17_exact_20260815`: the class-recoupling experiment on the same record; hull entry within $1.25 \cdot 10^{-9}$ of the record constant.
- `goalA_scale_sum_revision_exact_20260808`: the round-down stripping calibration on the public 16/15 instance and the two-terminal labelled-marginal wall.

Band frontier and merger reduction (Section 7).

- `part3_w1_refined_k17_hull_exact_20260809`: the exact hull entry of the standing record equals its critical constant; lifts the band wall to $K^* \geq 2.565201380\ldots$.
- `goalA_band_threshold_q_exact_20260809`: the exact-mean band threshold of the $k = 17$ KKT flow (16-atom mixture and terminal-marginal separator).
- `goalA_band_pareto_k17_full_exact_20260809`: the complete 15-segment relaxed-merger frontier of that flow.
- `goalA_fixedbox_k17_deficit_exact_20260813`: the radius-one deficit $0.297915757\ldots$ of the one-band witness; refutes FT(2).

Graph classes (Section 8).

- `part3_hub_tight_family_exact_20260809`: the $(m+1)$-jobs-on-m-hubs family with exact hull entry $1 - 1/m$; closes $C(\text{hub}) = 1$.
- `part3_v2_sp_tight_family_exact_20260810`: the series-parallel tightness family search.
- `part3_sp_16_15_regression_exact_20260809`: the 16/15 gadget and its supersink completion are not series-parallel (K_4 subdivisions).
- `part3_v2_planar_k6_edit_r1_exact_20260810` and `part3_v2_planar_k6_edit_r1_topology_exact_20260810`: the planar record instance 58676765987259/50000000000000 and its rotation-system planarity certificate.
- `part3_v2_commonpoint_outerplanar_wall_exact_20260810`: the $K_{2,3}$ outerplanar wall for common-point spine-and-tent graphs.
- `part3_v3_outerplanar_interval_exact_20260810` and `part3_v3_outerplanar_interval_ttsp_scope_20260810`: the outerplanar two-exit theorem regressions and the census showing the class is not covered by two-terminal series-parallel rounding.

Forests, total unimodularity, and the pair-sign theorem (Sections on two-path structure).

- `part3_v3_signed_factor_forest_exact_20260810` and `part3_v3_balanced_signed_forest_exact_20260810`: the factor-forest coupling regressions (3600 generator and 1800 independent instances).

- `part3_v4_pairrow_tu_theorem_exact_20260810`: the pair-row closure census (33,856 ordered DAGs, 12,292 common rows, no pair-sign violation).
- `part3_v4_three_column_tu_classification_exact_20260810`: the ternary sign classification (256 candidates, 28 compatible, three normal forms).
- `part3_v4_record_factor_rank_exact_20260810`: row-arity and cycle-rank ledger of the certified record geometries.

The local ternary envelope (Section on $E(3)$).

- `part3_v5_ternary_one_sided_symbolic_exact_20260810`: the 816 exact threshold decisions and symbolic identities behind $E(3) = 9/8$.
- `part3_v5_ternary_envelope_second_20260810`: the independent rational-grid rebuild (215,784 cases).
- `part3_v5_nested_floating_symbolic_exact_20260810` and `part3_v5_nested_radius1_second_20260810`: the nested core $\{111, 110\}$ has envelope exactly one.

Arity-four classification and boundary walls.

- `part3_v5_four_column_network_classification_exact_20260810` and `..._second_20260810`: the five realizable four-column normal forms (1024 subsets, 221 compatible, 75 labelled maxima).
- `part3_v5_exact_three_path_tu_census_full_20260810` and `..._second_20260810`: the three-path base-choice census and the six-vertex determinant-two wall.
- `part3_v5_fano_path_wall_exact_20260810` and `..._second_20260810`: the Fano plane is not tree-path realizable.
- `part3_v5_ternary_exact_transversal_wall_exact_20260810` and `..._second_20260810`: the fourteen-column witness with Property B but no exact transversal.

The arity-four envelope (Section 9).

- `part3_v8_e4_all_cells_solverfree_ledger_exact_20260811`: all 2,015 downset/facet cells mapped to closing lemmas with no nonlinear-solver decision.
- `part3_v8_e4_complete_occurrence_ledger_exact_20260811`: the occurrence-to-lemma audit of the same cells.
- `part3_v8_ternary_projection_closure_exact_20260811` and `..._second_20260811`: the projection-family closures.
- `part3_v8_e4_closure_manifest_20260811`: the hash-locked manifest of the closure round.