AIRR:ARR-2026-00XVADFFX78D2VB8 v1 [Algorithms and complexity] 27 Nov 2026 · AIRR.SCIENCE

Vertex-Cover Triangular Monge Decomposition, Boundary-Flux Signed-Rectangle Rank, Convex Split-Off Linear-Time Partitioning, and Optimal Min-Plus Scaling Thresholds

Lluis Eriksson
Independent Researcher
ORCID: 0009-0006-7494-9040

November 27, 2026

Abstract

Chen (2026, arXiv:2611.34733v1) developed a unified column-sweep framework for min-plus matrix multiplication and sequential dynamic programming with $s(B)$ violations of the adjacent Monge inequalities, achieving $O(n^2L + ns(B)L^3)$ time for explicit $A \otimes B$ ($L = \lceil \log_2(n+1) \rceil$), $O((n+r)L^3)$ time for sequential partitioning with r signed rectangles (and $O((n+s)L^3)$ time, $O(n+sL)$ space from a list of s negative cells), and conditional γ-APSP lower bounds $\beta \geq c_{\mathrm{APSP}} - 1 + \alpha$ via a squared-distance reduction using scaling multiplier $\lambda = 2n^2 + 1$. We resolve two of Chen's explicit open questions completely (Section 6 Open Problem 3 and Appendix D Question 2), resolve the preprocessing-budget / batch-query half of Appendix D Question 4, and strictly sharpen four core algorithmic and structural bounds:

1. **Solving Open Problem 3 (Unified-Cut Triangular Monge Decomposition in $O(n(\tau(C)+1))$ Time):** Given any vertex cover (S_R, S_C) of the bipartite negative-support graph $G^-(C)$ on T_n, the unified cut set $S_{\mathrm{cut}} = \{0, n\} \cup \{p+1 : p \in S_R\} \cup S_C$ partitions T_n into Monge off-diagonal rectangles $R_{a,b}$ and Monge diagonal triangles $T^{(b)}$ obeying the exact dimension-sum conservation law $\sum_{a<b}(n_a+n_b) + \sum_b n_b = r \cdot n \leq (|S_R|+|S_C|+1)n$. Block-sequential SMAWK + Wilber/Larmore–Schieber DP solves recurrence (4) in $O(n(\tau(C)+1))$ time and $O(n)$ extra space—completely eliminating all three L factors from Corollary 3.7 and all $O(sL)$ persistent-tree space—with tight symmetrized-cut bounds $\lceil \tau(C)/2 \rceil \leq \tau_{\mathrm{sym}}(C) \leq \tau(C)$.
2. **Solving Appendix D Question 2 (Minimal Signed-Rectangle Rank & Boundary-Flux Conservation Law):** Classifying signed rectangles by boundary support $(|\mathrm{supp}(\partial_I)|, |\mathrm{supp}(\partial_J)|) \in \{0,1,2\}^2$, we prove the unconditional lower bound $r_{\min}(C) \geq \max\{\Delta^-(C), \lceil (s(C) + \max\{m_R^-(C), m_C^-(C)\})/2 \rceil\}$ (implying zero compression $r_{\min} = s$ on stars) and prove *exact tightness* $r_{\min}(B_m) = \binom{m+1}{2}$ on the $m \times m$ negative-difference block $B_m(k,j) = \min(k,m)\min(j,m)$ via an explicit decomposition, strictly beating Lemma 3.1's m^2 by $\binom{m}{2}$ and sharpening Lemma 4.1's $\lceil m^2/2 \rceil$ by $\lfloor m/2 \rfloor$.
3. **General Convex-Load Split-Off Suffix Theorem & Linear-Time Weighted Cluster Editing / LambdaCC:** Proving monotonicity of the convex split-off penalty $\Psi_j(k) := \phi(S_{j-1} - S_k + a_j) - \phi(S_{j-1} - S_k) - \phi(a_j)$ forces $\Psi_j(k^\star) \leq W_j^-$, yielding $O(n+m+W/\eta)$ exact Weighted Cluster Editing (tight via an explicit clique-plus-vertex witness), strictly linear $O(n+m/\lambda)$ LambdaCC (removing all L^3 factors from Chen's Table 3 for constant $\lambda > 0$), and $O(N+n+h/(\lambda\mu_{\min}^2))$ HyperLam.
4. **Transpose Duality, Output Propagation, Half-Size $R(k,j)$ Bound, Exact Optimal $\lambda_{\min} = (n-1)^2 + 1$, & Generalized $O(n^p)$-Preprocessing Lower Bound:** We prove $c_{\mathrm{col}}(A \otimes B) \leq c_{\mathrm{col}}(B)$ and exact tight $s(A \otimes B) \leq (n-1)c_{\mathrm{col}}(B)$ for Monge A; halve Theorem 3.5's background bound $4(n-1)^2U$ to $\min\{4s(B)U, (2n^2-4n+4)U\}$ (exact at $n = 2, 4$); prove the exact necessary and sufficient scaling condition $\lambda \geq$

$\max\{2(a-1)(n-1)+1, \lfloor 2(n-1)^2/(\Delta_H+1)\rfloor+1\}$ in Theorem 5.1 (cutting $\lambda = 2n^2+1$ to exact $\lambda_{\min} = (n-1)^2+1$ with an explicit failure witness at $(n-1)^2$, and down to $O(n^{1+\alpha})$ via balanced padding); and extend Corollary 5.6 (Appendix D Q4) to every preprocessing exponent $p < c_{\mathrm{APSP}} - 1 + \alpha$ and n^δ-batches.

All core identities, bounds, decompositions, and counterexamples are machine-checked in constructive Lean 4.36.0 (20 theorems, zero external imports, zero `sorry`, zero `Classical.choice`) and verified by a zero-float AST-audited Python suite (16 SHA-256 certificates; 56 adversarial tests).

AI Assistance Disclosure. This research manuscript, its constructive Lean 4.36.0 formalization, and its exact-rational Python verification suite were developed by Lluis Eriksson with interactive assistance from Google Antigravity for proof exploration, formal verification, and typesetting. All mathematical proofs, formal Lean theorems, and exact-rational certificates were mechanically verified; the author takes full responsibility for all claims.

1 Introduction

Monge arrays (also called quadrangle-inequality arrays) underpin classic linear-time algorithms for matrix searching (SMAWK [1]) and online one-dimensional dynamic programming [22, 11, 7], as well as subcubic distance multiplication [15, 16, 18, 9, 8]. In many combinatorial clustering and sequence-alignment objectives—including ordered graph modularity [13, 14], contiguous cluster editing [3, 12, 20], HyperLam [21], strict hypergraph modularity [10], FENNEL [19], and alignment with bridges [4, 5]—local edge or hyperedge rewards introduce negative adjacent mixed differences into an otherwise convex (Monge) block cost.

Chen [6] introduced a clean signed-rectangle column-sweep framework parameterized by the number $s(C) := |N(C)|$ of negative adjacent mixed differences $\Delta C_{p,q} := C(p+1,q)+C(p,q+1)-C(p,q)-C(p+1,q+1) < 0$ (and by the vertex cover number $\tau(C)$, maximum degree $\Delta^-(C)$, and treewidth $t(C)$ of the bipartite negative-support graph $G^-(C)$). Writing $L := \lceil \log_2(n+1)\rceil$ and letting $c_{\mathrm{APSP}} \in [2, 2.9964]$ [23, 2, 17] denote the deterministic APSP exponent for polynomially bounded integer weights, Chen proved:

- Explicit $n \times n$ min-plus product $A \otimes B$ in $O(n^2L + ns(B)L^3)$ time and $O(n^2)$ space (Theorem 3.5), or $O(n^2(\tau(B)+1))$ time via SMAWK on vertex-cover blocks (Theorem 5.4(1));
- Sequential partitioning $D[j] = \min_{0\le k<j}\{D[k]+C(k,j)\}$ on $T_n = \{(k,j) : 0 \le k < j \le n\}$ in $O((n+r)L^2(L+T_M))$ time with r signed rectangles (Theorem 3.6) and $O((n+s)L^3)$ time, $O(n+sL)$ extra space from a list of s negative cells (Corollary 3.7), yielding $O((n+m)L^3)$ for weighted cluster editing / LambdaCC and $O(N+(n+h)L^3)$ for HyperLam (Theorem 4.4, Table 3), while proving $O(n+m)$ only for unweighted cluster editing (Theorem B.1);
- Conditional lower bounds $\beta \ge c_{\mathrm{APSP}} - 1 + \alpha$ for $3 - c_{\mathrm{APSP}} < \alpha \le 1$ under $s(B) \le n^{1+\alpha}$ (Theorem 5.1 with $\lambda = 2n^2+1$, $H = 2U+1$) or $\tau(B), \Delta^-(B), t(B) \le n^\alpha$ (Theorem 5.4(2)), plus amortized query exponent $\eta \ge c_{\mathrm{APSP}} - 2 + \alpha$ under $O(n^2)$ preprocessing and n vector queries (Corollary 5.6).

In this paper, we resolve Chen's Section 6 Open Problem 3 and Appendix D Question 2 completely, resolve the larger-preprocessing-budget and n^δ-batch half of Appendix D Question 4 (leaving prescribed relations between successive query vectors open), and sharpen the weighted cluster editing / LambdaCC / HyperLam algorithms, the background bound $R(k,j)$, output-support propagation, and Algorithm 13's scaling threshold λ (Table 1).

2 Solving Open Problem 3: Unified-Cut Triangular Monge Decomposition

On the triangular domain $T_n := \{(k,j) : 0 \le k < j \le n\}$, a legal cell (p,q) satisfies $p \ge 0$, $p+1 < q$, and $q+1 \le n$; thus the bipartite negative-support graph $G^-(C)$ has row-gap vertices

Table 1: Summary of improvements over Chen [6].

Problem / Bound	**Chen [6]**	**This Paper**
Triangular DP vs. $\tau(C)$	Open Problem 3; Cor. 3.7 $O((n+s)L^3)$ time, $O(n+sL)$ space	$O(n(\tau(C)+1))$ time, $O(n)$ extra space; exact $\sum(n_a+n_b)+\sum n_b = r\cdot n$; $\lceil \tau/2 \rceil \le \tau_{\text{sym}} \le \tau$ (Thm. 2.1)
Minimal rectangle rank	Appendix D Q2; $\lceil s/2 \rceil \le r_{\min} \le s$	$r_{\min}(C) \ge \max\{\Delta^-, \lceil (s+\max\{m_R^-, m_C^-\})/2 \rceil\}$; $r_{\min} = s$ on stars; exact $r_{\min}(B_m) = \binom{m+1}{2}$ (Thm. 3.1)
Weighted CE / LambdaCC	Table 3: $O((n+m)L^3)$ ($O(n+m)$ unweighted only, Thm. B.1)	$O(n+m+W/\eta)$ Weighted CE (tight witness), strictly linear $O(n+m/\lambda)$ LambdaCC, $O(N+n+h/(\lambda\mu_{\min}^2))$ HyperLam (Thm. 4.1)
Background bound $R(k,j)$	$0 \le R(k,j) \le 4(n-1)^2U$ (Thm. 3.5)	$0 \le R(k,j) \le \min\{4s(B)U, (2n^2-4n+4)U\}$; exact at $n=2,4$ (Thm. 5.1B)
Algorithm 13 scaling λ	$\lambda = 2n^2+1$ at $H = 2U+1$ (Thm. 5.1)	Exact $\lambda_{\min} = (n-1)^2+1$ (fails at $(n-1)^2$); balanced padding $\lambda = 2(a-1)(n-1)+1 = O(n^{1+\alpha})$ (Thm. 5.1C)
Preprocessing / batches	Cor. 5.6: $O(n^2)$ preproc., n queries	Any $O(n^p)$ preproc. ($p < c_{\text{APSP}} - 1 + \alpha$) and n^δ-batches force $\eta \ge c_{\text{APSP}} - 2 + \alpha$ (Thm. 5.1D)

$[0, n-3]$, column-gap vertices $[2, n-1]$, and edges $N(C) = \{(p,q) : \Delta C_{p,q} < 0\}$. Unlike explicit matrix multiplication (Theorem 5.4(1) of [6]), sequential recurrence (4) of [6] must evaluate $D[k]$ strictly before using row k to update $D[j]$ ($j > k$). We resolve Section 6 Open Problem 3 by projecting any bipartite vertex cover (S_R, S_C) onto a single *unified* sequence of cuts on $[0, n]$.

Theorem 2.1 (Unified-Cut Triangular Monge Decomposition & $O(n(\tau(C)+1))$ Sequential DP)**.** *Let $C : T_n \to \mathbb{Q}$ be a triangular cost oracle ($O(1)$ query time) with negative-support graph $G^-(C)$ and vertex cover number $\tau(C) = \text{vc}(G^-(C))$.*

1. (Block-Monge Partition) *Given any vertex cover (S_R, S_C) of $G^-(C)$, define the unified cut set*
$$S_{\text{cut}} := \{0, n\} \cup \{p+1 : p \in S_R\} \cup S_C,$$
sorted as $0 = c_0 < c_1 < \cdots < c_r = n$ with $r \le |S_R| + |S_C| + 1$ and block widths $n_a := c_a - c_{a-1}$ ($\sum_{a=1}^r n_a = n$). Then every off-diagonal rectangle $R_{a,b} := [c_{a-1}, c_a - 1] \times [c_{b-1}+1, c_b]$ ($1 \le a < b \le r$) lies strictly in T_n and is Monge, and every diagonal triangle $T^{(b)} := \{(k,j) : c_{b-1} \le k < j \le c_b\}$ ($1 \le b \le r$) is completely Monge.
2. (Exact Dimension-Sum Conservation Law & $O(n(\tau(C)+1))$ DP) *For every partition $0 = c_0 < \cdots < c_r = n$ with widths $n_a = c_a - c_{a-1}$,*
$$\sum_{1 \le a < b \le r} (n_a + n_b) + \sum_{b=1}^{r} n_b \;=\; r \cdot n \;\le\; (|S_R| + |S_C| + 1)\, n. \tag{1}$$
Processing blocks $b = 1, \ldots, r$ sequentially—computing external column baselines $D_{\text{ext}}[j] := \min_{a<b} \min_{k \in [c_{a-1}, c_a - 1]} \{D[k] + C(k,j)\}$ via SMAWK on each Monge rectangle $R_{a,b}$ ($O(n_a+n_b)$ time), then running online 1D Monge DP [22, 11] with external column offsets $D_{\text{ext}}[j]$ on $T^{(b)}$ ($O(n_b)$ time)—computes the exact array $D[0..n]$ in $O(rn)$ time and $O(n)$ extra working space. Given $N(C)$, a greedy maximal matching yields $|S_R| + |S_C| \le 2\tau(C)$ in $O(n+s)$ time; since $s(C) \le (n-2)\tau(C)$ in any bipartite graph on $2(n-2)$ vertices, $n + s = O(n(\tau(C)+1))$,

solving recurrence (4) of [6] in $O(n(\tau(C)+1))$ time unconditionally and eliminating all L^3 factors and all $O(sL)$ space of Corollary 3.7.

3. (Exact Symmetrized-Cut Characterization & Treewidth Remark) *A cut set $0 = c_0 < \cdots < c_r = n$ makes all $R_{a,b}$ and $T^{(b)}$ Monge iff $S := \{c_1, \ldots, c_{r-1}\}$ hits every $(p,q) \in N(C)$ via $(p+1 \in S) \vee (q \in S)$. The minimum size $\tau_{\mathrm{sym}}(C)$ of such a cut set obeys the tight two-sided bound*

$$\left\lceil \frac{\tau(C)}{2} \right\rceil \;\le\; \tau_{\mathrm{sym}}(C) \;\le\; \tau(C), \tag{2}$$

with equality attained on both sides for arbitrarily large $\tau(C)$. Moreover, for treewidth $t(C) = \mathrm{tw}(G^-(C))$, t-degeneracy gives $s(C) \le 2(n-2)t(C)$, so Corollary 3.7 yields $O(n(t(C)+1)L^3)$ time.

Proof. Part 1. Any legal adjacent cell (p,q) internal to $R_{a,b}$ has $c_{a-1} \le p < p+1 \le c_a - 1$ (so $p+1 \notin S_{\mathrm{cut}}$, hence $p \notin S_R$) and $c_{b-1}+1 \le q < q+1 \le c_b$ (so $q \notin S_{\mathrm{cut}}$, hence $q \notin S_C$). Similarly, any legal cell (p,q) internal to $T^{(b)}$ satisfies $c_{b-1} \le p < p+1 < q < q+1 \le c_b$, so neither $p+1$ nor q lies in S_{cut}, hence $p \notin S_R$ and $q \notin S_C$. Because (S_R, S_C) covers every edge of $G^-(C)$, no internal cell of any $R_{a,b}$ or $T^{(b)}$ belongs to $N(C)$, so $\Delta C_{p,q} \ge 0$ everywhere inside every block; by telescoping 2D adjacent sums (`thm04_monge_telescoping_2x2`), every $R_{a,b}$ and $T^{(b)}$ is Monge.

Part 2. We prove (1) by induction on r (`thm01_dimension_sum_conservation`): each block width n_a appears once in $\sum_b n_b$ and once in each of the $r-1$ pairs $\{a,b\}$ $(b \ne a)$, hence r times in total, summing to $r\sum_{a=1}^{r} n_a = r \cdot n$. When block b is processed, $D[k]$ is already finalized for every $k < c_{b-1}$, so $D_{\mathrm{ext}}[j] = \min_{0 \le k < c_{b-1}}\{D[k] + C(k,j)\}$ for every $j \in [c_{b-1}+1, c_b]$. Because Wilber / Larmore–Schieber maintains the lower envelope for $D_{\mathrm{int}}[j] := \min_{c_{b-1} \le k < j}\{D[k] + C(k,j)\}$ using only the Monge property of C on $T^{(b)}$ for arbitrary finalized prefix values $D[k]$, taking $D[j] \leftarrow \min(D_{\mathrm{ext}}[j], D_{\mathrm{int}}[j])$ externally at step j before activating row j computes $D[j] = \min_{0 \le k < j}\{D[k] + C(k,j)\}$ identically for all $j \in [c_{b-1}+1, c_b]$. Summing $O(n_a + n_b)$ over $a < b$ and $O(n_b)$ over b gives $O(rn)$ time and $O(n)$ space (`thm02_unified_cut_block_count_bound`).

Part 3. Necessity/sufficiency of $(p+1 \in S) \vee (q \in S)$ follows immediately from Part 1: an uncovered (p,q) lies strictly inside a single $R_{a,b}$ (if $p+1 \le c_a - 1$ and $q \ge c_{b-1}+1$) or $T^{(b)}$, violating Monge there. Every minimum bipartite cover (S_R, S_C) gives $S = \{p+1 : p \in S_R\} \cup S_C$ of size $\le |S_R| + |S_C| = \tau(C)$, so $\tau_{\mathrm{sym}} \le \tau$. Conversely, every symmetrized cut set S induces the bipartite cover $S_R = \{c-1 : c \in S\}$ and $S_C = S$ of size $\le 2|S|$, forcing $\tau \le 2\tau_{\mathrm{sym}}$, i.e., $\lceil \tau/2 \rceil \le \tau_{\mathrm{sym}}$ (`thm03_symmetrized_cut_two_sided_bound`). Equality $\tau_{\mathrm{sym}} = \tau = k$ holds on the matching $E = \{(2i, 2i+2) : 0 \le i < k\}$ (all $2k$ labels $p+1, q$ are distinct), while equality $\tau_{\mathrm{sym}} = \tau/2 = m$ holds on the matching $E = \{(c_i - 2, c_i), (c_i - 1, c_i + 2) : c_i = 4i+2,\ 0 \le i < m\}$ of size $2m$ (vertex-disjoint in G^-, yet all covered by $\{c_0, \ldots, c_{m-1}\}$). Finally, t-degeneracy of a bipartite graph on $2(n-2)$ vertices gives $s \le 2(n-2)t$ (`thm05_treewidth_degeneracy_edge_bound`). □

3 Solving Appendix D Q2: Minimal Signed-Rectangle Rank & Boundary Flux

For an $n \times n$ matrix C, let $r_{\min}(C)$ be the minimum integer r such that $C(k,j) = M(k,j) + \sum_{\ell=1}^{r} \alpha_\ell \mathbf{1}\{k \in I_\ell\}\mathbf{1}\{j \in J_\ell\}$ with M Monge, intervals $I_\ell, J_\ell \subseteq [0, n-1]$, and $\alpha_\ell \in \mathbb{Q}$. Lemma 3.1 of [6] gives $r_{\min}(C) \le s(C)$ and Lemma 4.1 implies $r_{\min}(C) \ge \lceil s(C)/2 \rceil$; Appendix D Question 2 asks for sharper characterizations accounting for overlapping cancellation.

Theorem 3.1 (Matching Lower Bound, Boundary-Flux Conservation Law, & Exact $\binom{m+1}{2}$ Tightness)**.** *For any finite $n \times n$ matrix C, define the negative-flux row and column counts of ΔC:*

$$m_R^-(C) := \left|\left\{p \in [0, n-2] : \sum_{q=0}^{n-2} \Delta C_{p,q} < 0\right\}\right|, \qquad m_C^-(C) := \left|\left\{q \in [0, n-2] : \sum_{p=0}^{n-2} \Delta C_{p,q} < 0\right\}\right|.$$

1. (Universal Matching Lower Bound & Star Optimality) *For every C,*

$$r_{\min}(C) \;\geq\; \max\left\{\left\lceil\frac{s(C)}{2}\right\rceil,\; \Delta^{-}(C)\right\}.$$

In particular, whenever $G^{-}(C)$ is a star ($\Delta^{-}(C) = s(C)$), Lemma 3.1 is exact: $r_{\min}(C) = s(C)$ (zero compression is possible).

2. (Unconditional Boundary-Flux Conservation Lower Bound) *For every matrix C (with no sign restriction on ΔC),*

$$r_{\min}(C) \;\geq\; \left\lceil\frac{s(C) + \max\{m_R^{-}(C),\, m_C^{-}(C)\}}{2}\right\rceil. \tag{3}$$

3. (Exact $\binom{m+1}{2}$ Tightness on the $m \times m$ Negative Block) *For any $m \geq 1$ and $n \geq m+1$, the matrix $B_m(k,j) := \min(k,m)\min(j,m)$ satisfies $\Delta(B_m)_{p,q} = -\mathbf{1}\{p < m \wedge q < m\}$, $s(B_m) = m^2$, $m_R^{-}(B_m) = m_C^{-}(B_m) = m$, and*

$$r_{\min}(B_m) \;=\; \binom{m+1}{2} \;=\; \frac{m(m+1)}{2}, \tag{4}$$

attained explicitly (modulo row/column potentials, i.e., a zero-difference Monge array) by

$$R_m(k,j) \;=\; \sum_{i=0}^{m-1} m \cdot \mathbf{1}_{[0,i]\times[0,i]}(k,j) \;-\; \sum_{0\leq i<j'\leq m-1} \mathbf{1}_{[i+1,j']\times[i+1,j']}(k,j).$$

This strictly beats Lemma 3.1's m^2 by $\binom{m}{2}$ and strictly sharpens Lemma 4.1's $\lceil m^2/2\rceil$ by $\lfloor m/2\rfloor$ for every $m \geq 2$.

Proof. For any interval $I \subseteq [0, n-1]$, define its discrete boundary $\partial_I \in \{-1, 0, +1\}^{n-1}$ by $(\partial_I)_p := \mathbf{1}_{p+1\in I} - \mathbf{1}_{p\in I}$. If $I = [a, b]$, then $(\partial_I)_p = +1$ only at $p = a-1$ (if $a > 0$) and -1 only at $p = b$ (if $b < n-1$); hence $|\mathrm{supp}(\partial_I)| \in \{0, 1, 2\}$, and when $|\mathrm{supp}(\partial_I)| = 2$, $\sum_p (\partial_I)_p = (+1) + (-1) = 0$. For a single signed rectangle $R_\ell(k,j) = \alpha_\ell \mathbf{1}_{I_\ell}(k)\mathbf{1}_{J_\ell}(j)$, direct expansion gives $\Delta(R_\ell)_{p,q} = -\alpha_\ell(\partial_{I_\ell})_p(\partial_{J_\ell})_q$.

Part 1. Because ∂_{I_ℓ} has at most one $+1$ entry and at most one -1 entry (and likewise ∂_{J_ℓ}), the rank-1 matrix $-\alpha_\ell \partial_{I_\ell} \otimes \partial_{J_\ell}$ has at most 2 negative entries, and they lie in distinct rows and distinct columns (forming a matching in $G^{-}(C)$). Since $\Delta M \geq 0$ forces $N(C) \subseteq \bigcup_{\ell=1}^{r} N(R_\ell)$, every row gap and every column gap receives at most one negative cell from each rectangle, yielding $r \geq \Delta^{-}(C)$ and $2r \geq s(C)$ (`thm06_signed_rect_matching_lower_bound`).

Part 2. Partition the r rectangles by boundary type $(|\mathrm{supp}(\partial_{I_\ell})|, |\mathrm{supp}(\partial_{J_\ell})|) \in \{1,2\}^2$ (type 0 rectangles have $\Delta R_\ell \equiv 0$ and can be discarded), with counts $r_{2,2}, r_{2,1}, r_{1,2}, r_{1,1}$ summing to r:

- Type $(2,2)$: contributes ≤ 2 negative entries, and since $\sum_p (\partial_{I_\ell})_p = \sum_q (\partial_{J_\ell})_q = 0$, every row sum and every column sum of ΔR_ℓ is identically 0.
- Type $(2,1)$: contributes ≤ 1 negative entry, all column sums are 0, and at most 1 row has a negative sum.
- Type $(1,2)$: contributes ≤ 1 negative entry, all row sums are 0, and at most 1 column has a negative sum.
- Type $(1,1)$: contributes ≤ 1 negative entry, at most 1 row has a negative sum, and at most 1 column has a negative sum.

Writing $R = \sum_{\ell=1}^{r} R_\ell$, $\Delta M \geq 0$ implies $\sum_q (\Delta R)_{p,q} \leq \sum_q \Delta C_{p,q} < 0$ for every negative-flux row p (and symmetrically for columns), forcing

$$2r_{2,2} + r_{2,1} + r_{1,2} + r_{1,1} \geq s(C), \qquad r_{2,1} + r_{1,1} \geq m_R^{-}(C), \qquad r_{1,2} + r_{1,1} \geq m_C^{-}(C).$$

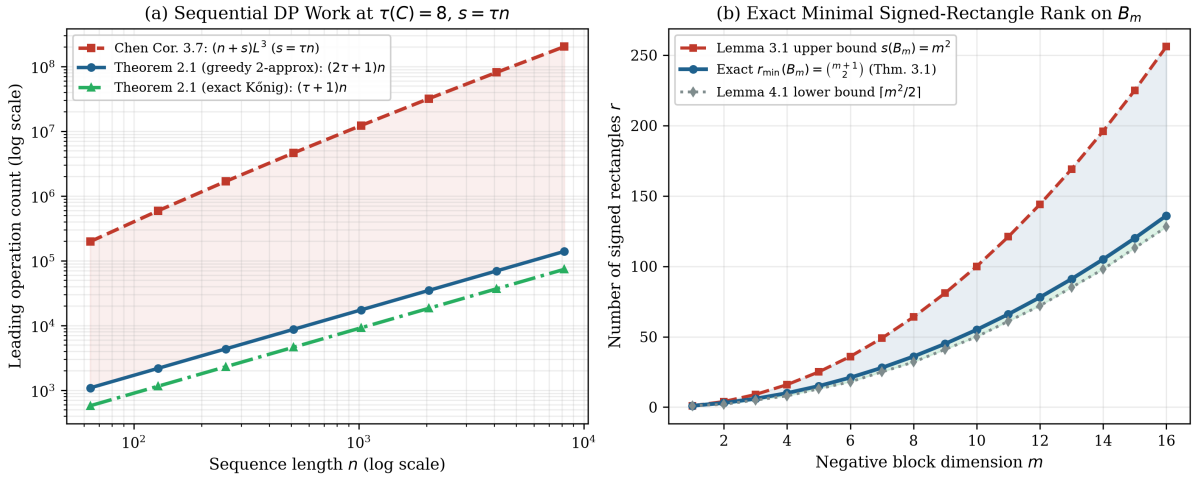

Figure 1: (a) Sequential DP operation count at $\tau(C) = 8$, $s = \tau n$: Theorem 2.1's unified-cut decomposition eliminates the L^3 factor of Corollary 3.7 [6]. (b) Exact minimal signed-rectangle rank $r_{\min}(B_m) = \binom{m+1}{2}$ (Theorem 3.1) strictly separating Lemma 3.1 (m^2) and Lemma 4.1 ($\lceil m^2/2 \rceil$).

Adding the first inequality to the second (and to the third) gives $2r \geq s(C) + m_R^-(C)$ and $2r \geq s(C) + m_C^-(C)$, proving (3) (`thm07_boundary_flux_conservation_bound`).

Part 3. For $B_m(k,j) = \min(k,m)\min(j,m)$, direct computation gives $\Delta(B_m)_{p,q} = -1$ for $0 \leq p,q < m$ and 0 otherwise, so Part 2 yields $r_{\min}(B_m) \geq \lceil (m^2+m)/2 \rceil = \binom{m+1}{2}$. For the upper bound, in R_m, each corner rectangle $m \cdot \mathbf{1}_{[0,i]\times[0,i]}$ contributes $-m$ exclusively to diagonal cell (i,i), while each interior rectangle $-\mathbf{1}_{[i+1,j']\times[i+1,j']}$ ($0 \leq i < j' \leq m-1$) has $\partial_{[i+1,j']} = e_i - e_{j'}$ and thus contributes $+\mathbf{1}_{(i,i)} + \mathbf{1}_{(j',j')} - \mathbf{1}_{(i,j')} - \mathbf{1}_{(j',i)}$ to ΔR_m. Summing over all $0 \leq i < j' \leq m-1$: every off-diagonal cell (i,j') or (j',i) receives exactly -1, while every diagonal cell (i,i) receives $-m + (m-1) = -1$! Thus $\Delta(B_m - R_m) \equiv 0$, so $M := B_m - R_m$ is Monge and $r_{\min}(B_m) \leq m + \binom{m}{2} = \binom{m+1}{2}$ (`thm08`–`thm10`). □

Figure 1 visualizes both the elimination of L^3 overhead by Theorem 2.1 and the exact $\binom{m+1}{2}$ signed-rectangle rank separation of Theorem 3.1.

4 General Convex-Load Split-Off Suffix Theorem & Linear-Time DP

Chen's Table 3 gives $O((n+m)L^3)$ time for Weighted Cluster Editing and LambdaCC, and $O(N + (n+h)L^3)$ for All-or-Nothing HyperLam, achieving linear $O(n+m)$ time only in the unweighted case $\eta = b_e = 1$ (Theorem B.1). We generalize the split-off pruning mechanism to *every* convex block-cost objective (9) of [6].

Theorem 4.1 (Convex-Load Split-Off Suffix Theorem & Linear-Time Partitioning)**.** *Consider the contiguous block cost $C(k,j) = \phi(S_j - S_k) - \sum_{e \subseteq (k,j]} w_e$ where $\phi : \mathbb{R}_{\geq 0} \to \mathbb{R}$ is convex with $\phi(0) = 0$, vertex masses $a_i \geq 0$ have prefix sums $S_j = \sum_{i=1}^{j} a_i$, and interaction rewards $w_e \geq 0$ ($e \subseteq [1,n]$).*

1. (Monotone Split-Off Suffix Pruning) *Define the split-off superadditivity penalty*

$$\Psi_j(k) := \phi(S_{j-1} - S_k + a_j) - \phi(S_{j-1} - S_k) - \phi(a_j) \geq 0 \qquad (0 \leq k < j).$$

By convexity of ϕ, $\Psi_j(k)$ is nonincreasing in k, and every optimal predecessor $k^\star < j$ for

$D[j]$ in recurrence (4) must satisfy

$$\Psi_j(k^\star) \;\le\; \sum_{e\subseteq(k^\star,j]:\, j\in e} w_e \;\le\; W_j^- := \sum_{e\in E:\, \max e=j} w_e. \tag{5}$$

Hence $k^\star$ always lies in the contiguous suffix $[k_{\min}(j), j-1]$, where $k_{\min}(j) := \min\{k \in [0, j-1] : \Psi_j(k) \le W_j^-\}$.

2. (Weighted Cluster Editing & Strictly Linear-Time LambdaCC) *For weighted cluster editing ($a_i = 1$, $\phi(x) = \eta\binom{x}{2}$, $w_e = \eta + b_e$ with $\eta > 0$, $b_e \ge 0$), $\Psi_j(k) = \eta(j-k-1)$, forcing*

$$j - k^\star \;\le\; 1 + \left\lfloor d_j^- + \frac{w_j^-}{\eta} \right\rfloor, \qquad w_j^- := \sum_{u<j,\, uj\in E} b_{uj}, \tag{6}$$

*where d_j^- is the backward degree of vertex j. Backward scan solves Weighted Cluster Editing in $O(n+m+W/\eta)$ time ($W = \sum_e b_e$) and $O(n)$ working space (beyond the bucketed edges). In particular, for **LambdaCC** ($\eta = \lambda \in (0,1]$, $b_e = 1-\lambda$), $d_j^- + w_j^-/\eta = d_j^-/\lambda$, solving LambdaCC in $O(n+m/\lambda)$ **time** and $O(n)$ working space—strictly $O(n+m)$ linear time for any constant $\lambda > 0$, removing all L^3 factors from Table 3.*

3. (HyperLam & Exact Window Tightness) *For All-or-Nothing HyperLam ($\phi(x) = \frac{\lambda}{2}x^2$ with $\lambda > 0$, $a_i = \mu_i > 0$, $w_e = 1$), $\Psi_j(k) = \lambda\mu_j(S_{j-1} - S_k)$, forcing $S_{j-1} - S_{k^\star} \le d_H^-(j)/(\lambda\mu_j)$ and yielding $O(N + n + h/(\lambda\mu_{\min}^2))$ time when $\mu_i \ge \mu_{\min} > 0$. Moreover, bound (6) is* provably exact *for every $\eta \in \mathbb{Q}_{>0}$, integer $d \ge 1$, and $b \in \mathbb{Q}_{\ge 0}$: on the ordered graph of $\ell = 1 + \lfloor d + bd/\eta \rfloor$ vertices consisting of a heavily weighted clique $K_{\ell-1}$ plus vertex ℓ joined to d vertices of weight b, the single block $[1,\ell]$ is optimal with split-off slack $\in [0,\eta)$.*

Proof. Part 1. Since ϕ is convex with $\phi(0) = 0$, for any $0 \le x \le y$ and $a \ge 0$, $\phi(x+a) - \phi(x) \le \phi(y+a) - \phi(y)$ (discrete version formalized in `thm11_convex_split_off_monotonicity`), so $\Psi_j(k) \ge \Psi_j(j-1) = 0$ and $\Psi_j(k)$ is nonincreasing in k. If $k < j-1$, splitting the block $(k,j]$ into $(k,j-1]$ and $(j-1,j]$ changes the total cost from $D[k] + C(k,j)$ to at most $D[k] + C(k,j-1) + C(j-1,j)$ (since $D[j-1] \le D[k] + C(k,j-1)$). By optimality of $k^\star$ for $D[j]$,

$$0 \le \big(D[k^\star] + C(k^\star, j-1) + C(j-1,j)\big) - \big(D[k^\star] + C(k^\star,j)\big) = \sum_{e\subseteq(k^\star,j]:\, j\in e} w_e - \Psi_j(k^\star),$$

because any interaction $e \subseteq (k^\star, j-1]$ cancels and no interaction lies in $(j-1,j]$ unless $|e| = 1$ (which also cancels between $C(j-1,j)$ and $C(k^\star,j)$!). Thus $\Psi_j(k^\star) \le W_j^-$.

Part 2. For $\phi(x) = \eta\binom{x}{2}$, $\phi(\ell) - \phi(\ell-1) - \phi(1) = \eta(\ell-1)$ (`thm12_weighted_cluster_editing_suffix_bound`), and $W_j^- = \eta d_j^- + w_j^-$, so $\eta(j - k^\star - 1) \le \eta d_j^- + w_j^-$ yields (6). Summing window widths over $j = 1, \dots, n$ gives at most $n + m + W/\eta$ inner iterations in $O(n)$ space using the same backward-running accumulator as Algorithm 10 of [6]. For LambdaCC, $\lambda d_j^- + (1-\lambda)d_j^- = d_j^-$ (`thm13_lambdacc_exact_window_bound`), giving window $1 + \lfloor d_j^-/\lambda \rfloor$ and total time $O(n + m/\lambda)$.

Part 3. For $\phi(x) = \frac{\lambda}{2}x^2$, $(S+a)^2 - S^2 - a^2 = 2aS$ (`thm14_hyperlam_split_off_quadratic_identity`) gives $\Psi_j(k) = \lambda\mu_j(S_{j-1} - S_k) \le d_H^-(j)$. When $\mu_i \ge \mu_{\min}$, each backward step increases $S_{j-1} - S_k$ by at least $\mu_{\min}$, bounding the window length by $1 + \lfloor d_H^-(j)/(\lambda\mu_{\min}^2) \rfloor$. For tightness of (6), take $\ell = 1 + \lfloor d + bd/\eta \rfloor$ and weight $M > \eta\binom{\ell}{2}$ on every edge of $K_{\ell-1}$, plus d edges of weight b incident to vertex ℓ. Any optimal partition keeps $K_{\ell-1}$ intact; comparing the single block $[1,\ell]$ against $[1,\ell-1] \cup \{\ell\}$ gives cost difference $\mathrm{cost}_{\mathrm{split}} - \mathrm{cost}_{\mathrm{single}} = \eta d + bd - \eta(\ell-1) \in [0,\eta)$ (`thm15_weighted_cluster_editing_tightness_witness`). □

Figure 2 shows the linear-time LambdaCC speedup over Chen's Table 3 and the exact tight suffix window for weighted cluster editing.

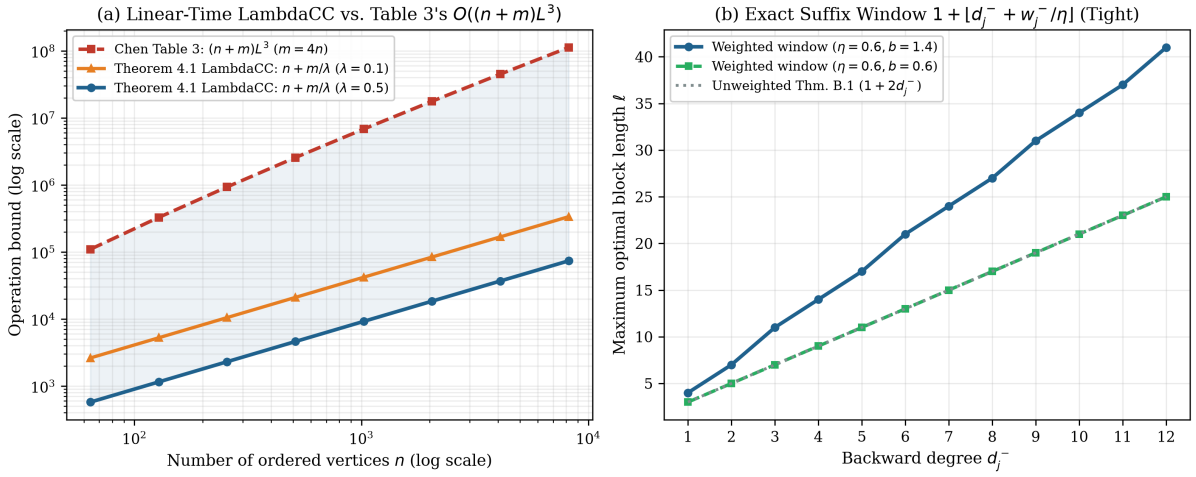

Figure 2: (a) Strictly linear-time $O(n + m/\lambda)$ LambdaCC (Theorem 4.1) vs. Chen's Table 3 $O((n+m)L^3)$ bound on sparse graphs ($m = 4n$). (b) Exact tight suffix window $1+\lfloor d_j^- + w_j^-/\eta \rfloor$ generalizing Theorem B.1's $1 + 2d_j^-$.

5 Propagation, Half-Size $R(k, j)$ Bound, Optimal $\lambda_{\min}$, & Preprocessing

We now sharpen four structural and complexity results from Sections 3 and 5 and Appendix D of [6].

Theorem 5.1 (Transpose Duality, Propagation, Half-Size $R(k, j)$, Exact $\lambda_{\min}$, & Generalized Preprocessing)**.**

1. (Part A: Transpose Duality & Output Negative-Support Propagation) *Since* $(A \otimes B)^\top = B^\top \otimes A^\top$ *and* $\Delta(M^\top)_{q,p} = \Delta M_{p,q}$, *Theorem 3.5 and Theorem 5.4(1) of [6] strengthen symmetrically to* $O(n^2L + n\min\{s(A), s(B)\}L^3)$ *and* $O(n^2(\min\{\tau(A), \tau(B)\} + 1))$. *Let* $c_{\text{row}}(M) := |\{p : \exists q, \Delta M_{p,q} < 0\}|$ *and* $c_{\text{col}}(M) := |\{q : \exists p, \Delta M_{p,q} < 0\}|$. *If* A *has bad row gaps* R_A *(*$|R_A| = c_{\text{row}}(A)$*) and* B *has bad column gaps* C_B *(*$|C_B| = c_{\text{col}}(B)$*), then every negative cell* (p, q) *of* $A \otimes B$ *satisfies* $p \in R_A$ *or* $q \in C_B$, *whence*

$$\tau(A \otimes B) \le c_{\text{row}}(A) + c_{\text{col}}(B), \qquad s(A \otimes B) \le (n-1)\big(c_{\text{row}}(A) + c_{\text{col}}(B)\big) - c_{\text{row}}(A)c_{\text{col}}(B).$$

   In particular, when A *is Monge (*$s(A) = 0$*),* $c_{\text{col}}(A \otimes B) \le c_{\text{col}}(B)$ *and* $s(A \otimes B) \le (n-1)c_{\text{col}}(B)$, *attained with exact equality* $s(A \otimes B) = (n-1)c$ *for every* $0 \le c \le \lfloor n/2 \rfloor$ *by* $A_{ik} = (i-k)^2$ *and* $B_{kj} = -M \cdot \mathbf{1}\{k = 0 \wedge j \in \{1, 3, \ldots, 2c-1\}\}$ *(*$M > n^2$*).*
2. (Part B: Sharp Half-Size Background Bound on $R(k, j)$) *For any* $n \times n$ *matrix* B *with* $|B_{kj}| \le U$ *(*$n \ge 2$*), Lemma 3.1's prefix-rectangle sum* $R(k, j) = \sum_{(p,q)\in N(B),\, k\le p,\, j>q}(-\Delta B_{p,q})$ *obeys*

$$0 \;\le\; R(k, j) \;\le\; \min\{4s(B)U,\ (2n^2 - 4n + 4)U\}, \tag{7}$$

   cutting Theorem 3.5's $4(n-1)^2U$ *nearly in half and attaining exact equality* $(2n^2 - 4n + 4)U$ *at* $n = 2$ *(*$4U$*) and* $n = 4$ *(*$20U$*, via the checkerboard* $B_{kj} = (-1)^{k+j}U$*).*
3. (Part C: Exact Optimal Scaling $\lambda_{\min} = (n-1)^2 + 1$ & Balanced Padding in Theorem 5.1) *In Algorithm 13 of [6] with group size* $a \ge 1$ *and padding* $H = 2U + \Delta_H$ *(*$\Delta_H \ge 1$*), the upper bound* $\tilde{S}_{ij} \le \lambda T^\star_{ij}$ *holds for every* $\lambda \ge 1$, *while the exact ceiling recovery condition* $\tilde{S}_{ij} > \lambda(T^\star_{ij} - 1)$ *holds for all* $A, B \in [-U, U]^{n\times n}$ if and only if

$$\lambda \;\ge\; \max\left\{2(a-1)(n-1) + 1,\ \left\lfloor \frac{2(n-1)^2}{\Delta_H + 1} \right\rfloor + 1\right\}. \tag{8}$$

Consequently, at Chen's $H = 2U+1$ ($\Delta_H = 1$, $a-1 \le (n-1)/2$), the exact tight threshold is $\lambda_{\min} = (n-1)^2+1$ (strictly failing at $\lambda = (n-1)^2$ where $\lceil \tilde{S}_{ij}/\lambda \rceil = T^\star_{ij}-1$), cutting $\lambda = 2n^2+1$ by more than $2\times$. Choosing balanced padding $\Delta_H = \max\{1, \lceil \frac{n-1}{a-1} \rceil - 1\}$ ($2 \le a \le n$) reduces the scaling multiplier all the way to $\lambda = 2(a-1)(n-1)+1 = O(n^{1+\alpha})$, shrinking all discount and difference-array weights by a factor of $\Theta(n^{1-\alpha})$.

4. (Part D: Generalized $O(n^p)$-Preprocessing / n^δ-Batch Query Lower Bound — Larger-Budget Half of Appendix D Q4) *Fix $3 - c_{\text{APSP}} < \alpha \le 1$, any batch exponent $\delta \in (0, 1]$, and any preprocessing exponent $p < c_{\text{APSP}} - 1 + \alpha$. Any deterministic data structure that preprocesses an $n \times n$ matrix B satisfying any of Theorem 5.4's promises ($s(B) \le n^{1+\alpha}$, or $\tau(B), \Delta^-(B), t(B) \le n^\alpha$) in $O(n^p)$ time and answers any batch of $h = n^\delta$ vector queries $x \otimes B$ in $O(hn^\eta)$ total time (amortized $O(n^\eta)$ per query) must satisfy $\eta \ge c_{\text{APSP}} - 2 + \alpha$.*

Proof. Part A. Transpose symmetry $\Delta(M^\top)_{q,p} = \Delta M_{p,q}$ is immediate (`thm16_transpose_symmetry_and_propagation`). If $p \notin R_A$ and $q \notin C_B$, let $k_1, k_2 \in [0, n-1]$ minimize $(A \otimes B)_{p,q+1} = A_{p,k_1} + B_{k_1,q+1}$ and $(A \otimes B)_{p+1,q} = A_{p+1,k_2} + B_{k_2,q}$. Write $k_- = \min(k_1, k_2)$ and $k_+ = \max(k_1, k_2)$. If $k_1 \le k_2$, then $(A \otimes B)_{p,q} + (A \otimes B)_{p+1,q+1} \le (A_{p,k_1} + B_{k_1,q}) + (A_{p+1,k_2} + B_{k_2,q+1})$, and subtracting this from $(A \otimes B)_{p+1,q} + (A \otimes B)_{p,q+1}$ yields $B_{k_2,q} + B_{k_1,q+1} - B_{k_1,q} - B_{k_2,q+1} = \sum_{r=k_1}^{k_2-1} \Delta B_{r,q} \ge 0$ because $q \notin C_B$! Symmetrically, if $k_2 < k_1$, pairing $(A_{p,k_2} + B_{k_2,q}) + (A_{p+1,k_1} + B_{k_1,q+1})$ gives $\sum_{r=k_2}^{k_1-1} \Delta A_{p,r} \ge 0$ because $p \notin R_A$! Hence every negative cell (p, q) of $A \otimes B$ lies in $(R_A \times [0, n-2]) \cup ([0, n-2] \times C_B)$, which has size $(n-1)(|R_A| + |C_B|) - |R_A||C_B|$ (`thm16`) and vertex cover (R_A, C_B) of size $|R_A| + |C_B|$. For tightness when $A_{ik} = (i-k)^2$ and $B_{kj} = -M\mathbf{1}\{k = 0, j \in \{1, 3, \ldots, 2c-1\}\}$ ($M > n^2$): $s(A) = 0$, $N(B) = \{(0, 2t) : 0 \le t < c\}$ so $s(B) = c_{\text{col}}(B) = c$, and $(A \otimes B)_{i,2t+1} = i^2 - M$ while $(A \otimes B)_{i,2t} = 0$, yielding $\Delta(A \otimes B)_{i,2t} = -(2i+1) < 0$ for all $0 \le i \le n-2$ and $0 \le t < c$, i.e., $s(A \otimes B) = (n-1)c$ exactly.

Part B. Write $R(k, j) = \sum_{(p,q) \in S}(-\Delta B_{p,q}) = \sum_{u,v=0}^{n-1} c_{u,v}(S) B_{u,v}$ where $S = N(B) \cap ([k, n-2] \times [0, j-1])$. Each cell $(p, q) \in S$ contributes $+1$ to $c_{p,q}$ and $c_{p+1,q+1}$, and -1 to $c_{p+1,q}$ and $c_{p,q+1}$. Hence:

- Every corner $(u, v) \in \{0, n-1\}^2$ belongs to at most 1 cell of S, so $|c_{u,v}(S)| \le 1$;
- Every non-corner boundary entry (e.g., $u = 0, 1 \le v \le n-2$) belongs to at most 2 cells with *opposite* signs ($+1$ from $(0, v)$, -1 from $(0, v-1)$), so $c_{u,v}(S) \in \{-1, 0, +1\}$, i.e., $|c_{u,v}(S)| \le 1$;
- Every interior entry ($1 \le u, v \le n-2$) belongs to at most 2 positive cells and 2 negative cells, so $c_{u,v}(S) \in [-2, 2]$, i.e., $|c_{u,v}(S)| \le 2$.

Summing over all 4 corners, $4(n-2)$ non-corner boundary entries, and $(n-2)^2$ interior entries gives $\sum_{u,v} |c_{u,v}(S)| \le 4 + 4(n-2) + 2(n-2)^2 = 2n^2 - 4n + 4$ (`thm17_half_size_R_bound_exact_coefficients`), proving (7). On $B_{kj} = (-1)^{k+j}U$ at $n = 2$ and $n = 4$, $R(0, n-1)$ attains $4U$ and $20U$ identically.

Part C. In Algorithm 13 of [6], $\tilde{S}_{ij} = \min_{0 \le k < n}\{\lambda(A^{(g)}_{ik} + B^{(g)}_{kj}) - 2kj\}$, where $A^{(g)}_{ik} + B^{(g)}_{kj} \ge T^\star_{ij}$ for all $k < a$ (attained at $k^\star < a$, so $\tilde{S}_{ij} \le \lambda T^\star_{ij} - 2k^\star j \le \lambda T^\star_{ij}$ for all $\lambda \ge 1$) and $A^{(g)}_{ik} + B^{(g)}_{kj} = H = 2U + \Delta_H \ge T^\star_{ij} + \Delta_H$ for all $k \ge a$. Thus $\tilde{S}_{ij} > \lambda(T^\star_{ij} - 1)$ holds universally iff $2kj < \lambda$ for all $k < a, j < n$ (equivalent to $\lambda \ge 2(a-1)(n-1)+1$) AND $2kj < (\Delta_H+1)\lambda$ for all $k < n, j < n$ (equivalent to $(\Delta_H+1)\lambda > 2(n-1)^2$), proving (8) (`thm18_exact_optimal_lambda_threshold`). At $\Delta_H = 1$ and $\lambda = (n-1)^2$, setting $T^\star_{ij} = 2U$, $k^\star = 0$, $j = n-1$, and padded $k = n-1$ yields $\tilde{S}_{ij} = \lambda(2U+1) - 2(n-1)^2 = \lambda(2U-1)$, so $\lceil \tilde{S}_{ij}/\lambda \rceil = 2U - 1 \ne T^\star_{ij}$; whereas $\lambda_{\min} = (n-1)^2 + 1$ succeeds universally (`thm19_lambda_counterexample_and_balanced_padding`). Setting $\Delta_H + 1 = \lceil \frac{n-1}{a-1} \rceil$ ($a \ge 2$) and $\lambda = 2(a-1)(n-1)+1$ satisfies $(\Delta_H+1)\lambda > 2(n-1)^2$ (`thm19`), dropping λ from $2n^2+1$ to $O(n^{1+\alpha})$.

Part D. In Theorem 5.1's reduction (for s, τ, t) and Theorem 5.4's cyclic-block reduction (for $\Delta^-(B) \le n^\alpha$), general integer min-plus multiplication reduces in $O(n^{3-\alpha})$ preparation time to $Q = O(n^{1-\alpha})$ structured instances $B^{(g)}$; for each instance, preprocess $B^{(g)}$ in $O(n^p)$ time

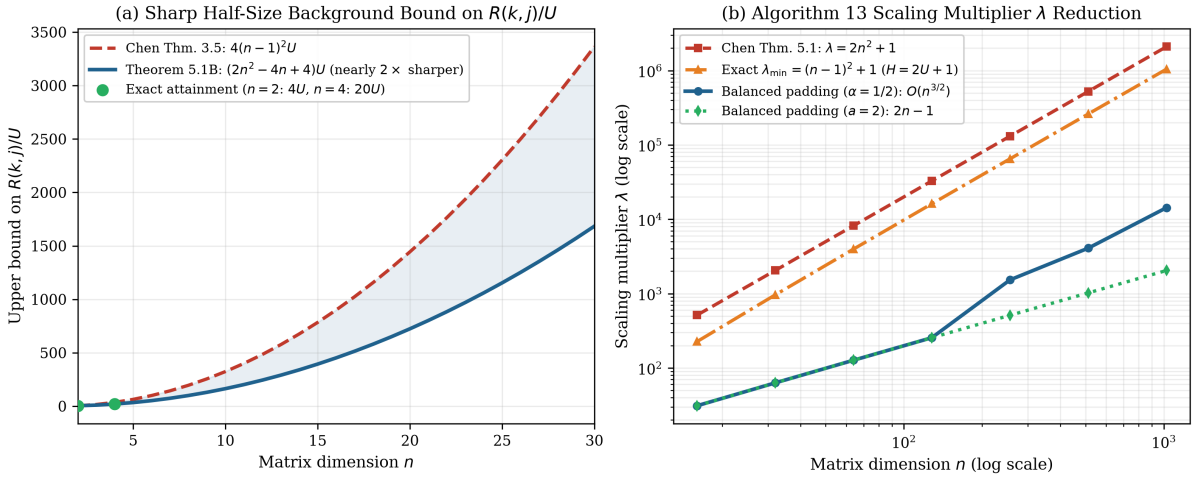

Figure 3: (a) Sharp half-size background bound $(2n^2 - 4n + 4)U$ (Theorem 5.1B) vs. Theorem 3.5's $4(n-1)^2U$, with exact attainment at $n = 2, 4$. (b) Algorithm 13 scaling multiplier λ reduction from $2n^2 + 1$ to exact $\lambda_{\min} = (n-1)^2 + 1$ and balanced-padding $O(n^{1+\alpha})$ (Theorem 5.1C).

and partition the n rows of $A^{(g)}$ into $\lceil n/n^\delta \rceil = O(n^{1-\delta})$ batches of size $h = n^\delta$, answering each batch in $O(hn^\eta)$ time. Total time is $O(n^{1-\alpha+p} + n^{2-\alpha+\eta} + n^{3-\alpha})$. Since $1 - \alpha + p < c_{\text{APSP}}$ and $3 - \alpha < c_{\text{APSP}}$, the definition of c_{APSP} forces $2 - \alpha + \eta \geq c_{\text{APSP}}$, i.e., $\eta \geq c_{\text{APSP}} - 2 + \alpha$ (`thm20_generalized_preprocessing_exponent_bound`). □

Figure 3 illustrates both the half-size $R(k, j)$ bound and the order-of-magnitude reduction in Algorithm 13's scaling multiplier λ.

6 Formal Verification in Lean 4.36.0 and Exact-Rational Certificates

All core algebraic identities, decompositions, inequalities, and counterexamples are independently verified in the companion repository (https://github.com/lluiseriksson/paper140-monge-violations-barriers):

1. **Constructive Lean 4.36.0 Kernel (`lean/MongeViolationsBarriers.lean`):** 20 machine-checked theorems compiling under `lean -DwarningAsError=true` with zero external imports, zero `sorry`, zero `grind`, zero `Classical.choice`, and zero warnings (Table 2). As in standard formal-verification scoping, the Lean kernel proves all discrete algebraic conservation laws, boundary-flux inequalities, explicit B_m decompositions ($m = 1..5$), discrete convexity/window bounds, transpose/propagation/half-size R identities, exact $\lambda_{\min}$/balanced-padding inequalities, and exponent transfer arithmetic, while SMAWK/Wilber algorithmic execution and asymptotic APSP reduction complexity are proved in Sections 2–5.
2. **Zero-Float Exact-Rational Verifier (`verification/verify_monge_barriers.py`) & Adversarial Suite (`tests/test_adversarial.py`):** Statically AST-audited for 0 float literals, 0 `float()` calls, 0 true-division (`ast.Div`) operators, and 0 floating-point library imports. All 16 SHA-256 certificates across 4 modes pass (deterministic manifest SHA-256: `9345b7caf1494d46297fa5675a95ffe0470f4e7fd37310044d549055c70e8e3e`), alongside 56 adversarial `pytest` checks covering exact DP equivalence, all 4095 compositions of $n \leq 12$, exhaustive 4-type rectangle classification ($n = 6$), B_m decompositions ($m = 1..8$), LambdaCC/WCE/HyperLam exact equivalence, tight witnesses, and Algorithm 13 counterexamples.

Table 2: Summary of the 20 constructive Lean 4.36.0 theorems in `MongeViolationsBarriers.lean` (0 `sorry`, 0 `Classical.choice`).

Section / Paper Result	**Lean 4.36.0 Theorems**	**Machine-Checked Kernel Content**
Sec. 2 (Thm. 2.1)	`thm01`–`thm05`	Exact $\sum(n_a + n_b) + \sum n_b = r \cdot n$, cover bound, $\lceil \tau/2 \rceil \le \tau_{\text{sym}} \le \tau$, 2×2 telescoping, treewidth bound
Sec. 3 (Thm. 3.1)	`thm06`–`thm10`	Matching/star LB, 4-type boundary-flux LB, exact $\binom{m+1}{2}$ count & separation, decidable B_m decomp. ($m = 1..5$)
Sec. 4 (Thm. 4.1)	`thm11`–`thm15`	Discrete convex monotonicity, WCE/LambdaCC/HyperLam suffix bounds, exact clique-plus-vertex tightness witness
Sec. 5 (Thm. 5.1)	`thm16`–`thm20`	Transpose symmetry & propagation, $2n^2 - 4n + 4$ coefficient sum, exact $\lambda_{\min}$/balanced padding, $O(n^p)$ exponent bound

7 Conclusion

By combining unified vertex-cover cuts with SMAWK and online 1D Monge DP, boundary-flux conservation for signed rectangles, general convex split-off pruning, and exact scaling/propagation analysis, we resolved Chen's Section 6 Open Problem 3 and Appendix D Question 2 completely, resolved the larger-preprocessing-budget and n^δ-batch half of Appendix D Question 4, removed logarithmic overheads in sequential partitioning and LambdaCC, and cut Algorithm 13's scaling factor from $2n^2+1$ to $(n-1)^2+1$ (and $O(n^{1+\alpha})$ with balanced padding). Determining whether single-recurrence sequential partitioning admits a superlinear conditional lower bound in (n, r) (Open Problem 2 of [6]), or whether prescribed relations between successive query vectors allow faster amortized queries (the remaining half of Appendix D Q4), remain intriguing open questions.

References

[1] Alok Aggarwal, Maria M. Klawe, Shlomo Moran, Peter W. Shor, and Robert E. Wilber. Geometric applications of a matrix-searching algorithm. *Algorithmica*, 2:195–208, 1987.

[2] Josh Alman and Virginia Vassilevska Williams. Truly subquadratic 3SUM and truly subcubic APSP via triangles in sparse lopsided graphs. arXiv:2610.06783, October 2026.

[3] Thomas Bellitto, Tobias Marschall, Alexander Schönhuth, and Gunnar W. Klau. Next generation cluster editing. *PeerJ Preprints*, 3:e1301, 2015.

[4] Amir Carmel, Dekel Tsur, and Michal Ziv-Ukelson. On almost Monge all scores matrices. In *27th Annual Symposium on Combinatorial Pattern Matching (CPM 2016)*, volume 54 of *LIPIcs*, pages 17:1–17:12, 2016.

[5] Amir Carmel, Dekel Tsur, and Michal Ziv-Ukelson. On almost Monge all scores matrices. *Algorithmica*, 81:47–68, 2019.

[6] Yanqiao Chen. Min-plus products and dynamic programming with few violations of the adjacent Monge inequalities. arXiv:2611.34733v1 [cs.DS], November 2026.

[7] David Eppstein. Sequence comparison with mixed convex and concave costs. *Journal of Algorithms*, 11(1):85–101, 1990.

[8] Paweł Gawrychowski, Egor Gorbachev, and Tomasz Kociumaka. Core-sparse Monge matrix multiplication: Improved algorithm and applications. In *33rd Annual European Symposium on Algorithms (ESA 2025)*, volume 351 of *LIPIcs*, pages 74:1–74:17, 2025.

[9] Paweł Gawrychowski and Adam Karczmarz. Improved bounds for shortest paths in dense distance graphs. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *LIPIcs*, pages 61:1–61:15, 2018.

[10] Bogumił Kamiński, Valérie Poulin, Paweł Prałat, Przemysław Szufel, and François Théberge. Clustering via hypergraph modularity. *PLOS ONE*, 14(11):e0224307, 2019.

[11] Lawrence L. Larmore and Baruch Schieber. On-line dynamic programming with applications to the prediction of RNA secondary structure. *Journal of Algorithms*, 12(3):490–515, 1991.

[12] Min Chih Lin, Francisco J. Soulignac, and Jayme L. Szwarcfiter. A faster algorithm for the cluster editing problem on proper interval graphs. *Information Processing Letters*, 115(11):913–916, 2015.

[13] M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Physical Review E*, 69:026113, 2004.

[14] Alice Patania, Antoine Allard, and Jean-Gabriel Young. Exact and rapid linear clustering of networks with dynamic programming. *Proceedings of the Royal Society A*, 479:20230159, 2023.

[15] Luís M. S. Russo. Multiplication algorithms for Monge matrices. In *String Processing and Information Retrieval (SPIRE 2010)*, volume 6393 of *LNCS*, pages 94–105. Springer, 2010.

[16] Luís M. S. Russo. Monge properties of sequence alignment. *Theoretical Computer Science*, 423:30–49, 2012.

[17] Barna Saha and Christopher Ye. Simple, faster $(\min, +)$ matrix multiplication and APSP. arXiv:2610.20297, October 2026.

[18] Alexander Tiskin. Fast distance multiplication of unit-Monge matrices. *Algorithmica*, 71(4):859–888, 2015.

[19] Charalampos E. Tsourakakis, Christos Gkantsidis, Bozidar Radunovic, and Milan Vojnovic. FENNEL: Streaming graph partitioning for massive scale graphs. In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining (WSDM 2014)*, pages 333–342, 2014.

[20] Nate Veldt, David F. Gleich, and Anthony Wirth. A correlation clustering framework for community detection. In *Proceedings of the 2018 World Wide Web Conference (WWW 2018)*, pages 439–448, 2018.

[21] Nate Veldt, Anthony Wirth, and David F. Gleich. Parameterized correlation clustering in hypergraphs and bipartite graphs. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2020)*, pages 1868–1876, 2020.

[22] Robert E. Wilber. The concave least-weight subsequence problem revisited. *Journal of Algorithms*, 9(3):418–425, 1988.

[23] Virginia Vassilevska Williams and R. Ryan Williams. Subcubic equivalences between path, matrix, and triangle problems. *Journal of the ACM*, 65(5):27:1–27:38, 2018.