Winner Determination by Polynomial Sparse Pursuit: Approximation Guarantees and Certified Optimality Gaps

Angshul Majumdar[a,*]

[a]*Department of Electronics and Communication Engineering, Indraprastha Institute of Information Technology Delhi, New Delhi, 110020, India*

Abstract

Winner determination in combinatorial auctions is a weighted set-packing problem and is NP-hard in general. This paper develops a deterministic polynomial-time sparse-pursuit framework that couples the global linear-programming relaxation with CoSaMP-type identify–merge–re-estimate–prune dynamics. The re-estimation step keeps only a logarithmic-size moving window integral while all other bid variables remain fractionally relaxed, so each partially integral global relaxation is polynomial-time solvable. Partial integrality itself is a standard optimization device, and the coordinate ranges of an LP-optimal face are established in the mixed-integer-programming literature as cloud intervals. We therefore make no novelty claim for either the intervals $[\ell_j, u_j]$ or the binary criterion $\ell_j > 0$, $u_j < 1$. The central contribution is instead a quantitative multi-coordinate window theorem. Let $X^\star_{\mathrm{LP}}$ be the full optimal face of the ordinary LP, let ℓ_j and u_j be the minimum and maximum of coordinate j over that face, and define the weighted cloud-interval score $\psi_j = v_j \min\{\ell_j, 1 - u_j\}$. A Hoffman-type polyhedral sharpness constant $\mu > 0$ yields, for every exactification window D,

$$U(D) \leq U_{\mathrm{LP}} - \mu \sum_{j \in D} \psi_j.$$

Consequently the identify step, which selects the largest weighted cloud-interval scores, maximizes the guaranteed amount of LP-bound tightening

*Corresponding author

Email address: angshul@iiitd.ac.in (Angshul Majumdar)

among all admissible windows of the same form and size; positive selected ambiguity forces strict tightening. If V_{C} is the returned feasible welfare, $V^\star$ is the exact WDP optimum, U_{P} is the best pursuit upper bound, and U_{LP} is the ordinary LP bound, then every instance satisfies

$$V_{\mathrm{C}} \le V^\star \le U_{\mathrm{P}} \le U_{\mathrm{LP}}.$$

Hence

$$0 \le \frac{V^\star - V_{\mathrm{C}}}{V^\star} \le 1 - \frac{V_{\mathrm{C}}}{U_{\mathrm{P}}} \le 1 - \frac{V_{\mathrm{C}}}{U_{\mathrm{LP}}}.$$

In particular, $V_{\mathrm{C}} \ge 0.97 U_{\mathrm{P}}$ certifies an optimality gap of at most 3%, while $U_{\mathrm{P}} \le 1.03 V_{\mathrm{C}}$ certifies the strictly smaller bound $1 - 1/1.03 < 2.913\%$. A protected approximation seed may be used in parallel, yielding an adaptive ratio equal to the better of its uniform guarantee and the pursuit certificate. No planted solution, probabilistic valuation model, restricted-isometry assumption, or distributional hypothesis is used.

Keywords: winner determination, combinatorial auctions, sparse pursuit, CoSaMP, cloud intervals, partial integrality, linear programming, approximation algorithms, certified optimality gap
2000 MSC: 68W25, 68T20, 90C27, 90C05, 91B26

---

1. Introduction

Winner determination is the computational core of a combinatorial auction. A set of bids is feasible when no packing resource is allocated more than once, and the objective is to choose a feasible collection of maximum total value. In the standard binary formulation this is weighted set packing. Exact WDP has therefore been studied through complete combinatorial search, decomposition, bounding, and mixed-integer optimisation [1, 2, 3]. Approximation is also classical because the exact problem is NP-hard and inherits the hardness of weighted set packing [4, 5].

The present paper asks a different question from exact search: can one retain a genuine sparse-pursuit mechanism, keep the entire procedure polynomial, and nevertheless obtain a rigorous deterministic certificate of distance from the unknown exact optimum? The construction begins from the ordinary LP relaxation of WDP. A CoSaMP-type procedure repeatedly identifies bid coordinates, merges them with a small retained support, re-estimates on the merged set, and prunes a feasible incumbent. Re-estimation does not

solve a smaller isolated auction. Instead, the merged coordinates are made integral inside the *full global relaxation*; every nonwindow bid remains in the model as a fractional variable.

Making a selected subset of binary variables integral while relaxing the rest is not, by itself, new. Branch-and-bound strengthens LP relaxations by branching on fractional integer variables, while relax-and-fix methods work with selected integral windows and relaxed complements [6, 7]. There is also substantially closer prior art. Cloud Branching introduced intervals obtained from alternative optima of the current LP relaxation and used them to identify branching candidates that remain fractional throughout the cloud [8]; subsequent work explicitly studied cloud intervals as projections of the LP-optimal face and their implications for branching under dual degeneracy [9]. Thus neither optimal-face coordinate intervals nor the binary condition that a coordinate stays strictly between zero and one over the whole optimal face is claimed as new here. The novelty question is instead whether a *multi-variable logarithmic pursuit window* admits a quantitative, additive, globally valid lower bound on certificate tightening, and whether the pursuit rule optimizes that bound. This paper supplies that link.

Let U_{LP} be the ordinary LP optimum and let $X^{\star}_{\mathrm{LP}}$ be its complete optimal face. For every bid coordinate define

$$\ell_j := \min_{x \in X^{\star}_{\mathrm{LP}}} x_j, \qquad u_j := \max_{x \in X^{\star}_{\mathrm{LP}}} x_j, \tag{1}$$

and the weighted cloud-interval ambiguity

$$\delta_j := \min\{\ell_j, 1 - u_j\}, \qquad \psi_j := v_j \delta_j. \tag{2}$$

The interval $[\ell_j, u_j]$ is the full-face version of the cloud interval of Berthold and Salvagnin [8]; Gamrath, Berthold, and Salvagnin subsequently studied these optimal-face projections directly [9]. For binary variables, $\ell_j > 0$ and $u_j < 1$ is precisely the established full-face condition that the variable is fractional throughout the optimal face and hence cannot preserve the parent LP value in either binary child. Our new quantity is the weighted distance $\psi_j = v_j \min\{\ell_j, 1 - u_j\}$ and, more importantly, its simultaneous-window consequence below. The ranges in (1) require only two additional LPs per coordinate and are therefore polynomially computable.

A Hoffman-type polyhedral error bound [10] yields an instance-dependent

sharpness modulus $\mu > 0$ such that every exactification window D satisfies

$$U(D) \le U_{\mathrm{LP}} - \mu\Psi(D), \qquad \Psi(D) := \sum_{j \in D} \psi_j. \tag{3}$$

For a fixed retained support and window cardinality, the identify step selects the largest ψ_j values and therefore maximizes $\Psi(D)$, hence maximizes the guaranteed tightening lower bound $\mu\Psi(D)$ among all admissible windows of that form. In particular,

$$\Psi(D) > 0 \quad \Longrightarrow \quad U(D) < U_{\mathrm{LP}}. \tag{4}$$

No LP uniqueness assumption is needed because the proxy is defined from the full optimal face rather than from one solver-selected optimum.

The exactification window contains only $O(\log m)$ bids. Hence its binary assignments can be enumerated in polynomial time, with each assignment leaving an ordinary LP over the remaining bids. Let $V^\star$ denote the exact winner-determination value, V_{C} the best feasible value returned by the pursuit, and U_{P} the smallest partially integral upper bound encountered along the pursuit trajectory. The resulting deterministic sandwich is

$$V_{\mathrm{C}} \le V^\star \le U_{\mathrm{P}} \le U_{\mathrm{LP}}. \tag{5}$$

Consequently,

$$0 \le \frac{V^\star - V_{\mathrm{C}}}{V^\star} \le 1 - \frac{V_{\mathrm{C}}}{U_{\mathrm{P}}} \le 1 - \frac{V_{\mathrm{C}}}{U_{\mathrm{LP}}}. \tag{6}$$

Thus a near-one pursuit certificate is already a deterministic near-exactness proof. In particular,

$$V_{\mathrm{C}} \ge 0.97 U_{\mathrm{P}} \quad \Longrightarrow \quad \frac{V^\star - V_{\mathrm{C}}}{V^\star} \le 0.03, \tag{7}$$

and

$$U_{\mathrm{P}} \le 1.03 V_{\mathrm{C}} \quad \Longrightarrow \quad \frac{V^\star - V_{\mathrm{C}}}{V^\star} \le 1 - \frac{1}{1.03} < 0.02913. \tag{8}$$

The second statement is the strict less-than-3% certificate.

A uniform worst-case approximation factor can be protected independently. If the seed routine used to initialize the incumbent has approximation ratio $\rho_0(\kappa)$ for packing rank κ, then incumbent retention and (5) give

$$\frac{V^\star}{V_{\mathrm{C}}} \le \min\left\{\rho_0(\kappa), \frac{U_{\mathrm{P}}}{V_{\mathrm{C}}}\right\}. \tag{9}$$

This deliberately separates known set-packing approximation theory from the pursuit contribution. A simple greedy seed gives the standard κ-approximation, but much stronger polynomial-time weighted set-packing algorithms are known, including the classical roughly $\kappa/2$ local-search scale and subsequent refinements [11, 12, 13]. Any such seed can be inserted without changing the pursuit theory.

The paper is entirely deterministic and theoretical. There are no planted instances, stochastic valuations, average-case assumptions, experiments, or restricted-isometry hypotheses. Section 2 states exact WDP and its sparse-support decision form. Section 3 develops the ordinary LP, the partially integral relaxation, and the optimal-face sharpness machinery. Section 4 gives the polynomial LP-CoSaMP algorithm with weighted cloud-interval identification. Section 5 proves feasibility, monotonicity, polynomial complexity, quantitative pursuit-driven tightening, and global certificates. Section 6 derives exact, absolute, relative, and 3% optimality-gap certificates. Section 7 couples those certificates to a protected approximation seed. Section 8 records the inherited hardness barrier. Section 9 distinguishes the contribution from branch-and-bound and relax-and-fix and places it relative to related optimisation work in *Theory of Computing Systems* [14, 15, 16].

2. Exact WDP and its sparse-support form

2.1. Binary packing formulation

Let $\mathcal{R} = \{1, \dots, n\}$ be the set of packing resources and let $\mathcal{B} = \{1, \dots, m\}$ index the bids. In the basic model the resources are auction items. XOR restrictions can be represented by additional bookkeeping resources shared by all bids in the same XOR class, so all constraints can be represented by one binary incidence matrix.

Bid j has nonnegative value v_j and requests a nonempty resource set $R_j \subseteq \mathcal{R}$. Zero-value bids may be deleted without changing the optimum, so after this trivial preprocessing we assume $v_j > 0$ for every retained bid. For any bid-index set $S \subseteq \mathcal{B}$, write $v(S) := \sum_{j \in S} v_j$. For any $D \subseteq \mathcal{B}$, let A_D denote the submatrix formed by the columns indexed by D, let v_D and x_D denote the corresponding subvectors, and let $D^c := \mathcal{B} \setminus D$. Define

$$A_{ij} = \begin{cases} 1, & i \in R_j, \\ 0, & i \notin R_j, \end{cases} \qquad A \in \{0,1\}^{n \times m}. \tag{10}$$

A binary decision vector $x \in \{0,1\}^m$ indicates accepted bids. Feasibility is exactly

$$Ax \leq \mathbf{1}. \tag{11}$$

The exact winner-determination value is

$$V^\star := \max_{x \in \{0,1\}^m} \left\{ v^\top x : Ax \leq \mathbf{1} \right\}. \tag{12}$$

This is the weighted set-packing formulation of WDP [1, 2].

Definition 2.1 (Packing rank)**.** The packing rank is

$$\kappa := \max_{1 \leq j \leq m} |R_j| = \max_{1 \leq j \leq m} \|a_j\|_0, \tag{13}$$

where a_j is column j of A.

2.2. Thresholded sparse-support decision problem

The sparse-support interpretation is most transparent in a thresholded decision form. For $\tau \geq 0$, define

$$\begin{aligned} \mathcal{P}(\tau): \qquad \min_{x \in \{0,1\}^m} \quad & \|x\|_0 \\ \text{s.t.} \quad & Ax \leq \mathbf{1}, \\ & v^\top x \geq \tau. \end{aligned} \tag{14}$$

The objective counts accepted bids; it does not replace the welfare objective of (12). Its role is to expose WDP as support selection under packing constraints.

Proposition 2.1 (Threshold equivalence)**.** *For every $\tau \geq 0$,*

$$\mathcal{P}(\tau) \text{ is feasible} \quad \Longleftrightarrow \quad \tau \leq V^\star. \tag{15}$$

Consequently,

$$V^\star = \sup\{\tau : \mathcal{P}(\tau) \text{ is feasible}\}. \tag{16}$$

Proof. If (14) is feasible, one of its feasible binary points has value at least τ, so $V^\star \geq \tau$. Conversely, if $\tau \leq V^\star$, every optimal binary solution of (12) is feasible for (14). Equation (16) follows. □

Remark 2.1 (Purpose of the sparse view). The paper does not claim that replacing (12) by (14) makes WDP easy. The sparse view motivates the CoSaMP-style identification and correction dynamics introduced later. Polynomial tractability comes from relaxing global integrality and restoring it only on logarithmic-size moving windows.

3. Global relaxation and logarithmic partial integrality

3.1. Ordinary LP relaxation

Relaxing $x \in \{0,1\}^m$ to $0 \le x \le 1$ gives

$$U_{\mathrm{LP}} := \max_{0 \le x \le 1} \left\{ v^\top x : Ax \le \mathbf{1} \right\}. \tag{17}$$

Let $X^\star_{\mathrm{LP}}$ denote the complete set of optimizers of (17).

Lemma 3.1 (LP upper bound)**.** *Every instance satisfies*

$$V^\star \le U_{\mathrm{LP}}. \tag{18}$$

Proof. Every feasible binary point of (12) is feasible for (17). □

The ordinary LP is therefore useful twice: its optimizer supplies a fractional guide for pursuit, and its optimum supplies a global upper bound. The weakness is that this upper bound may be loose when important coordinates are fractional. Indeed, even in the unweighted κ-uniform hypergraph-matching subclass, the standard set-packing LP has integrality gap $\kappa - 1 + 1/\kappa$ in general [17]. The next construction makes a selected set of coordinates integral without making the whole problem integer.

3.2. Partial exactification

Definition 3.1 (Partially integral relaxation)**.** For a bid-index set $D \subseteq \mathcal{B}$, define

$$U(D) := \max_x \left\{ v^\top x : Ax \le \mathbf{1}, 0 \le x \le 1, x_j \in \{0,1\}\ \forall j \in D \right\}. \tag{19}$$

The set D is called an *exactification window*.

Theorem 3.1 (Nested relaxation sandwich)**.** *For every $D \subseteq \mathcal{B}$,*

$$V^\star \le U(D) \le U_{\mathrm{LP}}. \tag{20}$$

If $D \subseteq D'$, then

$$U(D') \le U(D). \tag{21}$$

Proof. Every feasible integer WDP solution is feasible for (19), giving the left inequality. Dropping the integrality constraints on D gives (17), giving the right inequality. If $D \subseteq D'$, the feasible region associated with D' is contained in the feasible region associated with D. □

The following characterization will later show exactly when a pursuit window improves the global certificate.

Theorem 3.2 (Equality and strict tightening)**.** *For a fixed window D,*

$$U(D) = U_{\mathrm{LP}} \tag{22}$$

if and only if there exists $\bar{x} \in X^{\star}_{\mathrm{LP}}$ such that

$$\bar{x}_D \in \{0,1\}^{|D|}. \tag{23}$$

Equivalently,

$$X^{\star}_{\mathrm{LP}} \cap \{x : x_D \in \{0,1\}^{|D|}\} = \varnothing \quad \Longrightarrow \quad U(D) < U_{\mathrm{LP}}. \tag{24}$$

Proof. If an optimizer of (17) is integral on D, it is feasible for (19) with objective U_{LP}, and Theorem 3.1 forces equality. Conversely, if $U(D) = U_{\mathrm{LP}}$, any optimizer of (19) is also globally LP-optimal and is integral on D by construction. Statement (24) is the contrapositive. □

3.3. Optimal-face ambiguity and polyhedral sharpness

A proxy built from one arbitrarily chosen LP optimizer can be misleading when the LP has multiple optima. We therefore measure ambiguity over the complete optimal face. For every coordinate j, define

$$\ell_j := \min\{x_j : x \in X^{\star}_{\mathrm{LP}}\}, \qquad u_j := \max\{x_j : x \in X^{\star}_{\mathrm{LP}}\}. \tag{25}$$

Both values are obtained by ordinary linear programs over the optimal face, equivalently by adding $v^{\top}x = U_{\mathrm{LP}}$ to the constraints of (17). Hence all $2m$ values in (25) are polynomially computable.

Remark 3.1 (Cloud intervals are prior art). The interval $[\ell_j, u_j]$ in (25) is not introduced here as a new geometric object. Berthold and Salvagnin define a *cloud interval* from alternative LP optima and show that, for binary variables, if the cloud contains all vertices of the optimal face then variables whose interval stays strictly between the two integer values are exactly those guaranteed to improve the dual bound in both child nodes [8]. Gamrath, Berthold, and Salvagnin later study cloud intervals explicitly as projections of the LP-optimal face and analyze whether integer variables can reach integral values while staying optimal [9]. Accordingly, the conditions $\ell_j > 0$ and $u_j < 1$, and the use of full-optimal-face ranges to avoid dependence on

one arbitrary LP optimum, are treated as established prior art. The contribution developed below begins with the weighted distance $v_j \min\{\ell_j, 1-u_j\}$ and culminates in a simultaneous-window Hoffman bound and a window-selection dominance theorem; neither of those quantitative multi-coordinate statements is supplied by the cloud-branching results.

Definition 3.2 (Weighted cloud-interval ambiguity)**.** For bid j, define

$$\delta_j := \min\{\ell_j, 1-u_j\}, \qquad \psi_j := v_j \delta_j. \tag{26}$$

For a set $D \subseteq \mathcal{B}$, define

$$\Psi(D) := \sum_{j \in D} \psi_j. \tag{27}$$

Because $0 \le \ell_j \le u_j \le 1$, one has $\delta_j \ge 0$, and $\delta_j > 0$ exactly when every LP optimum keeps coordinate j strictly between zero and one.

Since every retained v_j is positive, define the weighted norm

$$\|z\|_{1,v} := \sum_{j=1}^{m} v_j |z_j|. \tag{28}$$

Write $\mathrm{dist}_{1,v}(x, X^\star_{\mathrm{LP}})$ for the corresponding distance to the LP-optimal face.

Lemma 3.2 (LP-face sharpness)**.** *There exists a constant $\mu > 0$, depending only on the instance, such that every x feasible for the ordinary relaxation satisfies*

$$U_{\mathrm{LP}} - v^\top x \ge \mu \, \mathrm{dist}_{1,v}(x, X^\star_{\mathrm{LP}}). \tag{29}$$

Proof. The set $X^\star_{\mathrm{LP}}$ is the nonempty polyhedron obtained by intersecting the LP feasible region with $v^\top x \ge U_{\mathrm{LP}}$; the reverse inequality already holds for every LP-feasible x. A Hoffman error bound therefore gives a finite constant H such that, for every LP-feasible x,

$$\mathrm{dist}_{1,v}(x, X^\star_{\mathrm{LP}}) \le H\big(U_{\mathrm{LP}} - v^\top x\big).$$

Such error bounds hold for arbitrary fixed norms; see Güler, Hoffman, and Rothblum [10]. Taking $\mu := 1/H$ proves the claim. □

Theorem 3.3 (Quantitative tightening from robust ambiguity)**.** *For every exactification window $D \subseteq \mathcal{B}$,*

$$U(D) \le U_{\mathrm{LP}} - \mu \Psi(D). \tag{30}$$

Consequently,

$$\Psi(D) > 0 \quad \Longrightarrow \quad U(D) < U_{\mathrm{LP}}. \tag{31}$$

Proof. Let y be feasible for the partially integral relaxation associated with D. For every $x^\star \in X^\star_{\mathrm{LP}}$ and every $j \in D$, $y_j \in \{0,1\}$ and $x^\star_j \in [\ell_j, u_j]$, so

$$|y_j - x^\star_j| \geq \min\{\ell_j, 1 - u_j\} = \delta_j.$$

Therefore

$$\|y - x^\star\|_{1,v} \geq \sum_{j \in D} v_j \delta_j = \Psi(D). \tag{32}$$

Taking the infimum over $x^\star \in X^\star_{\mathrm{LP}}$ and then applying Lemma 3.2 gives

$$v^\top y \leq U_{\mathrm{LP}} - \mu \Psi(D).$$

Maximizing over all feasible y of the partially integral relaxation proves (30). The strict statement follows because $\mu > 0$. □

Remark 3.2 (Sufficient, not necessary). Positive weighted cloud-interval ambiguity is a sufficient condition for strict tightening, not a necessary one. A group of coordinates can jointly eliminate the LP-optimal face even when each coordinate is individually integral at some LP optimum. Theorem 3.2 remains the exact characterization; Theorem 3.3 supplies the quantitative selection principle used by the pursuit.

3.4. Why a logarithmic window is polynomial

Fix

$$s := \max\left\{1, \left\lceil \frac{1}{3}\log_2(m+1) \right\rceil \right\}. \tag{33}$$

The pursuit will merge at most s retained indices with at most $2s$ newly identified indices, so every exactification window satisfies $|D| \leq 3s$.

Lemma 3.3 (Polynomial solvability of a pursued window)**.** *If $|D| \leq 3s$ with s defined by* (33)*, then $U(D)$ can be computed in polynomial time. In particular,*

$$2^{|D|} \leq 8(m+1) \tag{34}$$

binary assignments are sufficient.

Proof. Enumerate $z \in \{0,1\}^{|D|}$. Discard any z with $A_D z \not\leq \mathbf{1}$. For each remaining assignment solve

$$\max_{0 \leq y \leq 1} \left\{ v_D^\top z + v_{D^c}^\top y : A_{D^c} y \leq \mathbf{1} - A_D z \right\}. \tag{35}$$

The largest completion value equals $U(D)$. Since $\lceil q \rceil \le q+1$, equation (33) implies $3s \le \log_2(m+1)+3$, hence (34). A linear number of polynomial-size LPs is polynomial in the rational input size. □

Remark 3.3 (General polynomial window family). For any fixed $\alpha > 0$ one may instead use

$$s_\alpha := \max\left\{1, \left\lceil \frac{\alpha}{3} \log_2(m+1) \right\rceil \right\}. \tag{36}$$

Then each re-estimation enumerates $O(m^\alpha)$ assignments. Thus logarithmic exactification yields a tunable polynomial hierarchy without changing any correctness theorem below.

4. LP-CoSaMP with moving partial integrality

4.1. Protected seed and support memory

Let $\mathcal{A}_0$ be any polynomial-time routine that returns a feasible winner set B_0. Denote

$$V_0 := v(B_0). \tag{37}$$

The pursuit always retains the better of its current incumbent and a newly constructed candidate, so the seed value can never be lost. For a finite index set S, an integer $r \ge 1$, and scores $w = (w_j)_{j \in S}$, write $\mathrm{Top}_r(S; w)$ for the subset of up to r indices in S having largest scores, with ties broken by increasing bid index.

The CoSaMP memory contains at most s incumbent bids. At iteration t define

$$P_t := \mathrm{Top}_s(B_t; v), \tag{38}$$

where ties are broken by bid index. Since $P_t \subseteq B_t$, it is itself feasible.

4.2. Identify: weighted cloud-interval pursuit

The primary pursuit score is the weighted cloud-interval ambiguity ψ_j from (26). The underlying interval geometry is inherited from Cloud Branching [8, 9]; the new element is to weight the full-face distance by bid value and use the resulting additive mass inside the multi-coordinate Hoffman bound of Theorem 3.3. Because the score is computed from the complete LP-optimal face, it does not depend on which particular LP optimizer a solver returns.

For CoSaMP-style dynamic tie breaking, let $\bar{x}^{(t)}$ be the most recently returned fractional guide and define

$$\chi_j^{(t)} := v_j \bar{x}_j^{(t)} \big(1 - \bar{x}_j^{(t)}\big). \tag{39}$$

The algorithm ranks eligible bids lexicographically by

$$\big(\psi_j, \chi_j^{(t)}, -j\big). \tag{40}$$

Thus robust ambiguity is primary, current fractional ambiguity breaks robust-score ties, and the index gives deterministic final tie breaking.

Let E_t be the set of bid indices that have appeared in at least one exactification window before iteration t. Among eligible indices in $\mathcal{B} \setminus (P_t \cup E_t)$, choose up to $2s$ highest-ranked bids and call the resulting set Ω_t. If all remaining primary and secondary scores vanish, use deterministic index order among the remaining eligible indices. If no eligible index remains, set $\Omega_t = \varnothing$; the retained set P_t is still exactified and is marked exposed after the iteration.

Lemma 4.1 (Pursuit window dominance for guaranteed tightening)**.** *Fix iteration t and the retained support P_t. Among all eligible sets Ω having the same cardinality as Ω_t,*

$$\Psi(P_t \cup \Omega_t) \geq \Psi(P_t \cup \Omega). \tag{41}$$

Consequently, the chosen identify batch maximizes the guaranteed tightening lower bound

$$\mu\Psi(P_t \cup \Omega) \tag{42}$$

among all admissible windows of that form and cardinality.

Proof. The retained set P_t is fixed. The primary ranking in (40) selects the largest eligible values of ψ_j, while the secondary score only resolves ties among equal primary values. Therefore Ω_t maximizes $\sum_{j\in\Omega} \psi_j$, and adding the fixed contribution $\sum_{j\in P_t} \psi_j$ gives (41). Multiplication by the common positive constant μ preserves the ordering. □

4.3. Merge, re-estimate, prune

The merged exactification window is

$$D_t := P_t \cup \Omega_t, \qquad |D_t| \leq 3s. \tag{43}$$

Solve the partially integral relaxation (19) on D_t. Let

$$U_t := U(D_t) \tag{44}$$

and let $\hat{x}^{(t)}$ be one optimizer. By Lemma 3.3, this re-estimation is polynomial.

The coordinates in D_t are binary in $\hat{x}^{(t)}$. Their selected subset is

$$Z_t := \{j \in D_t : \hat{x}_j^{(t)} = 1\}. \tag{45}$$

Start a feasible integer candidate from Z_t. Scan indices outside D_t in nonincreasing order of

$$q_j^{(t)} := v_j \hat{x}_j^{(t)}, \tag{46}$$

and add a bid whenever packing feasibility is preserved. Denote the final candidate by C_t. The incumbent is updated by

$$B_{t+1} \in \arg\max\{v(B_t), v(C_t)\}, \tag{47}$$

with deterministic tie breaking. The best known upper bound is updated by

$$\bar{U}_{t+1} := \min\{\bar{U}_t, U_t\}, \qquad \bar{U}_0 := U_{\mathrm{LP}}. \tag{48}$$

Finally set $E_{t+1} := E_t \cup D_t$ and use $\hat{x}^{(t)}$ as the next fractional guide. The process stops when all bids have been exposed or when the upper and lower bounds coincide.

5. General deterministic pursuit theory

5.1. Feasibility, monotonicity, and complexity

Lemma 5.1 (Feasible candidate generation)**.** *Every C_t and every incumbent B_t generated by Algorithm 1 is feasible for* (12)*.*

Proof. The seed is feasible by assumption. Since $\hat{x}^{(t)}$ is feasible for (19) and binary on D_t, the selected set Z_t satisfies the packing constraints. The subsequent scan accepts an outside bid only when compatibility is preserved. Hence C_t is feasible. The incumbent update chooses between feasible sets. □

Theorem 5.1 (Monotone incumbent)**.** *The incumbent values satisfy*

$$v(B_{t+1}) \geq v(B_t) \tag{49}$$

for every completed iteration.

Algorithm 1 LP-CoSaMP with logarithmic partial integrality

Require: Incidence matrix A, values v, polynomial feasible seed routine $\mathcal{A}_0$.
Ensure: Feasible winner set B_{C}, value V_{C}, pursuit bound U_{P}.

Compute s from (33).
Solve (17); obtain U_{LP} and $\bar{x}^{(0)}$.
Compute $\ell_j, u_j, \delta_j, \psi_j$ for all bids from (25)–(26).
$B_0 \leftarrow \mathcal{A}_0(A, v)$, $E_0 \leftarrow \varnothing$, $\bar{U}_0 \leftarrow U_{\mathrm{LP}}$, $t \leftarrow 0$.
while $E_t \neq \mathcal{B}$ and $\bar{U}_t > v(B_t)$ **do**
    $P_t \leftarrow \mathrm{Top}_s(B_t; v)$.
    Compute secondary scores $\chi^{(t)}$ from (39).
    Identify up to $2s$ eligible bids by (40); call them Ω_t.
    if $\Omega_t = \varnothing$ **then**
        Take up to $2s$ remaining unexposed indices by index order.
    end if
    Merge $D_t \leftarrow P_t \cup \Omega_t$.
    Solve (19) on D_t; obtain U_t and $\hat{x}^{(t)}$.
    Construct C_t from (45) and the compatibility scan using (46).
    Update B_{t+1} by (47) and $\bar{U}_{t+1}$ by (48).
    $E_{t+1} \leftarrow E_t \cup D_t$, $\bar{x}^{(t+1)} \leftarrow \hat{x}^{(t)}$, $t \leftarrow t + 1$.
end while
return $B_{\mathrm{C}} \leftarrow B_t$, $V_{\mathrm{C}} \leftarrow v(B_t)$, $U_{\mathrm{P}} \leftarrow \bar{U}_t$.

Proof. Equation (47) retains the higher-value feasible set. □

Theorem 5.2 (Polynomial running time)**.** *For rational input encoded with bit length L, Algorithm 1 runs in time polynomial in m, n, and L.*

Proof. The initial global LP is polynomial. Computing the optimal-face ranges in (25) requires $2m$ additional ordinary LPs, hence polynomial work. Every completed iteration places at least one previously unexposed bid in D_t unless all bids have already been exposed, so there are at most m completed iterations; when full $2s$ identify batches are available, the usual count is $O(m/s)$. Each merged window has at most $3s$ coordinates. Lemma 3.3 reduces one re-estimation to at most $8(m+1)$ polynomial-size LPs. Robust-score lookup, dynamic tie breaking, sorting, compatibility tests, and candidate pruning are polynomial. A polynomial number of polynomial-time operations gives the claim. □

5.2. Pursuit-specific global certificates

Theorem 5.3 (Pursuit-selected quantitative tightening)**.** *For every completed iteration t,*

$$U(D_t) \le U_{\mathrm{LP}} - \mu\Psi(D_t). \tag{50}$$

Moreover, for every admissible competing identify batch Ω of the same cardinality,

$$\mu\Psi(D_t) \ge \mu\Psi(P_t \cup \Omega). \tag{51}$$

Hence, if $\Psi(D_t) \ge \gamma$ for some $\gamma > 0$, then

$$U(D_t) \le U_{\mathrm{LP}} - \mu\gamma < U_{\mathrm{LP}}. \tag{52}$$

Proof. Equation (50) is Theorem 3.3 applied to the actual merged window D_t. Equation (51) follows from Lemma 4.1. If $\Psi(D_t) \ge \gamma > 0$, substitute into (50) and use $\mu > 0$. □

Let T be the number of completed pursuit iterations and define

$$U_{\mathrm{P}} := \min\left\{U_{\mathrm{LP}}, U(D_0), \ldots, U(D_{T-1})\right\}. \tag{53}$$

Theorem 5.4 (Pursuit-certified sandwich)**.** *Algorithm 1 returns*

$$V_{\mathrm{C}} \le V^\star \le U_{\mathrm{P}} \le U_{\mathrm{LP}}. \tag{54}$$

If $V_{\mathrm{C}} > 0$, then

$$1 \le \frac{V^\star}{V_{\mathrm{C}}} \le \frac{U_{\mathrm{P}}}{V_{\mathrm{C}}} \le \frac{U_{\mathrm{LP}}}{V_{\mathrm{C}}}. \tag{55}$$

Proof. Lemma 5.1 gives $V_{\mathrm{C}} \le V^\star$. Theorem 3.1 gives $V^\star \le U(D_t) \le U_{\mathrm{LP}}$ for every pursued window. Taking the minimum gives (54); division by $V_{\mathrm{C}} > 0$ gives (55). □

Theorem 5.5 (Strict certificate improvement caused by pursuit)**.** *Strict improvement occurs whenever $\Psi(D_t) > 0$ for some pursued window. More generally, it also occurs whenever a pursued window D_t satisfies*

$$X^\star_{\mathrm{LP}} \cap \{x : x_{D_t} \in \{0,1\}^{|D_t|}\} = \varnothing. \tag{56}$$

In either case,

$$U_{\mathrm{P}} < U_{\mathrm{LP}}. \tag{57}$$

If $V_{\mathrm{C}} > 0$, then

$$\frac{U_{\mathrm{P}}}{V_{\mathrm{C}}} < \frac{U_{\mathrm{LP}}}{V_{\mathrm{C}}}. \tag{58}$$

Proof. If $\Psi(D_t) > 0$, Theorem 5.3 gives $U(D_t) < U_{\mathrm{LP}}$. Under the more general condition (56), Theorem 3.2 gives the same conclusion. Equation (53) then gives $U_{\mathrm{P}} \le U(D_t) < U_{\mathrm{LP}}$. Division by $V_{\mathrm{C}} > 0$ gives (58). □

Remark 5.1 (Why the pursuit is mathematically indispensable). Partial integrality alone is not the claimed novelty: arbitrary windows, relax-and-fix procedures, and branch-and-bound can also impose integrality on selected variables. The pursuit-specific content is Theorem 5.3: the face-robust identify rule maximizes a rigorous lower bound on certificate tightening among all admissible windows of the same form and size. If the identify step is deleted or replaced by an arbitrary window, the dominance statement (51) disappears. This is the formal distinction between LP-CoSaMP and "choose an arbitrary logarithmic window and partially exactify it."

6. Certified optimality gaps

The pursuit sandwich yields more than an approximation ratio. It directly bounds the unknown difference between the polynomial solution and the exact WDP optimum.

Theorem 6.1 (Absolute optimality-gap certificate)**.** *Algorithm 1 satisfies*

$$0 \le V^\star - V_{\mathrm{C}} \le U_{\mathrm{P}} - V_{\mathrm{C}} \le U_{\mathrm{LP}} - V_{\mathrm{C}}. \tag{59}$$

Proof. Subtract V_{C} from every term of (54). □

Theorem 6.2 (Relative optimality-gap certificate)**.** *Assume $V_{\mathrm{C}} > 0$. Then the unknown exact relative gap satisfies*

$$0 \le \frac{V^\star - V_{\mathrm{C}}}{V^\star} \le 1 - \frac{V_{\mathrm{C}}}{U_{\mathrm{P}}} \le 1 - \frac{V_{\mathrm{C}}}{U_{\mathrm{LP}}}. \tag{60}$$

Thus

$$\delta_{\mathrm{P}} := 1 - \frac{V_{\mathrm{C}}}{U_{\mathrm{P}}} \tag{61}$$

is a computable pursuit-specific upper bound on the relative distance from the exact optimum.

Proof. For fixed $V_{\mathrm{C}} > 0$, define $f(z) := 1 - V_{\mathrm{C}}/z$ on $z \ge V_{\mathrm{C}}$. Then $f'(z) = V_{\mathrm{C}}/z^2 > 0$, so f is increasing. Theorem 5.4 gives $V_{\mathrm{C}} \le V^\star \le U_{\mathrm{P}} \le U_{\mathrm{LP}}$. Therefore

$$f(V^\star) \le f(U_{\mathrm{P}}) \le f(U_{\mathrm{LP}}),$$

which is exactly (60). □

Corollary 6.1 (General ε-near-exactness certificate)**.** *For any* $\varepsilon \in [0, 1)$*,*

$$V_{\mathrm{C}} \geq (1 - \varepsilon) U_{\mathrm{P}} \tag{62}$$

implies

$$\frac{V^\star - V_{\mathrm{C}}}{V^\star} \leq \varepsilon. \tag{63}$$

Proof. Condition (62) is equivalent to $1 - V_{\mathrm{C}}/U_{\mathrm{P}} \leq \varepsilon$. Apply Theorem 6.2. □

Corollary 6.2 (Three-percent certificate)**.** *If*

$$V_{\mathrm{C}} \geq 0.97 U_{\mathrm{P}}, \tag{64}$$

then

$$\frac{V^\star - V_{\mathrm{C}}}{V^\star} \leq 0.03. \tag{65}$$

If instead

$$U_{\mathrm{P}} \leq 1.03 V_{\mathrm{C}}, \tag{66}$$

then the stronger strict bound holds:

$$\frac{V^\star - V_{\mathrm{C}}}{V^\star} \leq 1 - \frac{1}{1.03} \approx 0.029126 < 0.03. \tag{67}$$

Proof. The first claim is Corollary 6.1 with $\varepsilon = 0.03$. For the second, (66) gives $V_{\mathrm{C}}/U_{\mathrm{P}} \geq 1/1.03$. Substitute into (60). □

Corollary 6.3 (Polynomial exactness certificate)**.** *If*

$$U_{\mathrm{P}} = V_{\mathrm{C}}, \tag{68}$$

then

$$V_{\mathrm{C}} = V^\star. \tag{69}$$

Proof. Theorem 5.4 places $V^\star$ between equal lower and upper bounds. □

Remark 6.1 (Why the 3% statement is not empirical). Corollary 6.2 does not assert that every instance will satisfy one of its hypotheses. It states that whenever the polynomial pursuit produces the corresponding certificate, the conclusion about the unknown exact optimum is mathematically forced. No exact solver, sample distribution, or empirical comparison is needed to validate the gap.

7. Adaptive approximation with a protected seed

A pursuit certificate is instance-specific. To retain a uniform worst-case approximation theorem as well, the algorithm protects the output of an arbitrary polynomial-time seed routine.

Definition 7.1 (Seed approximation ratio)**.** A polynomial seed routine $\mathcal{A}_0$ has approximation ratio $\rho_0(\kappa) \geq 1$ if its feasible output value V_0 satisfies

$$V_0 \geq \frac{V^\star}{\rho_0(\kappa)} \tag{70}$$

for every rank-κ instance.

Theorem 7.1 (Adaptive deterministic approximation guarantee)**.** *If the protected seed has ratio $\rho_0(\kappa)$, Algorithm 1 satisfies*

$$\frac{V^\star}{V_{\mathrm{C}}} \leq \min\left\{\rho_0(\kappa), \frac{U_{\mathrm{P}}}{V_{\mathrm{C}}}\right\}. \tag{71}$$

Equivalently, the final approximation factor is never worse than the seed guarantee and may be certified substantially better by the pursuit.

Proof. Theorem 5.1 gives $V_{\mathrm{C}} \geq V_0$, so Definition 7.1 implies $V^\star/V_{\mathrm{C}} \leq \rho_0(\kappa)$. Theorem 5.4 gives $V^\star/V_{\mathrm{C}} \leq U_{\mathrm{P}}/V_{\mathrm{C}}$. Take the smaller upper bound. □

Proposition 7.1 (Self-contained greedy fallback)**.** *If $\mathcal{A}_0$ scans bids in nonincreasing value and accepts each bid whenever it remains feasible, then*

$$\rho_0(\kappa) = \kappa \tag{72}$$

is valid.

Proof. Let G be the greedy winner set and O an optimal winner set. Charge each $o \in O$ either to itself when $o \in G$, or to the first greedy bid $g \in G$ whose prior acceptance prevented o. Because greedy scans by nonincreasing value, $v_g \geq v_o$. Since bids in O are mutually resource-disjoint, distinct optimal bids charged to g must intersect R_g through distinct resources. Hence at most $|R_g| \leq \kappa$ optimal bids are charged to one greedy bid. Summing charges gives $V^\star \leq \kappa v(G)$. □

Remark 7.1 (Known stronger seeds). Proposition 7.1 is included only to make the framework self-contained. It is not presented as a new approximation result. The weighted κ-set-packing literature contains substantially stronger polynomial-time guarantees. Berman established the classical approximately $\kappa/2$ local-search scale [11]; Neuwohner subsequently sharpened that scale, established limits of logarithmically bounded pure local search, and developed improvements beyond the $\kappa/2$ threshold [12, 13]. Any such algorithm may be substituted for the default seed in Theorem 7.1. The pursuit contribution is the second, independently computed term in (71) and the gap theorems in Section 6.

8. Complexity-theoretic limits

The general WDP contains weighted bounded-rank set packing as a subclass. Consequently, no polynomial sparse-pursuit construction can promise a parameter-independent near-exact ratio on every instance unless standard complexity assumptions fail.

Theorem 8.1 (Inherited inapproximability)**.** *Assume* $\mathrm{P} \neq \mathrm{NP}$*. There exists a constant* $c > 0$ *such that, for all sufficiently large* κ*, no polynomial-time algorithm can guarantee approximation ratio at most*

$$c\,\frac{\kappa}{\log \kappa} \tag{73}$$

on every rank-κ *WDP instance.*

Proof. Weighted κ-set packing is the restriction of WDP in which every bid uses at most κ resources and feasibility is pairwise disjointness. Hazan, Safra, and Schwartz establish an $\Omega(\kappa/\log \kappa)$ hardness of approximation for this problem unless $\mathrm{P} = \mathrm{NP}$ [5]. Any uniformly stronger polynomial WDP guarantee would contradict that lower bound on this subclass. □

Corollary 8.1 (No uniform constant independent of rank)**.** *Unless* $\mathrm{P} = \mathrm{NP}$*, no fixed constant* C *independent of* κ *can be a valid polynomial-time approximation ratio for all rank-*κ *WDP instances.*

Proof. A fixed C is $o(\kappa/\log \kappa)$ as $\kappa \to \infty$, contradicting Theorem 8.1. □

Remark 8.1 (Uniform guarantees versus certificates). Theorem 8.1 does not conflict with a certificate such as (60). A worst-case approximation factor

must hold before seeing the instance-specific relaxation geometry and therefore deteriorates with unrestricted rank. By contrast, $U_{\mathrm{P}}/V_{\mathrm{C}}$ is computed after polynomial processing of the actual instance. It may be close to one on a particular instance without implying a rank-independent uniform approximation theorem.

9. Relation to exact WDP, approximation, and mathematical AI

Exact WDP has a substantial specialized literature. Sandholm developed exact winner-determination algorithms and complexity analysis [1]; BOB improved complete search through stronger bounding and structural methods [2]; CABOB further combined decomposition, bid ordering, and strong upper and lower bounds [3]. These methods pursue exact optimality and therefore retain exponential worst-case behavior in the general problem. The present framework takes the opposite design decision: it caps the number of simultaneously integral coordinates at logarithmic size and therefore guarantees polynomial running time.

Approximate allocation in combinatorial auctions is also longstanding [4]. Because WDP is weighted set packing, the relevant benchmark is not the elementary greedy factor but the much stronger bounded-rank approximation literature. Berman established a local-search approximation on the roughly $\kappa/2$ scale [11]. Neuwohner sharpened the analysis of weighted κ-set packing, characterized limitations of bounded local search, and developed methods that pass the pure-local-search barrier [12, 13]. LP-CoSaMP is not another attempt to improve that universal approximation constant. Instead it is designed to compute a global, pursuit-dependent a posteriori certificate such as (60), which can certify near-exactness on an individual instance even when the best uniform guarantee is necessarily much weaker.

The sparse-pursuit architecture comes from CoSaMP [18]: identify candidate coordinates, merge them with retained support, re-estimate on the merged set, and prune. The analogy stops at the algorithmic architecture. In the present setting, the re-estimation operation is a partially integral WDP relaxation rather than least squares, and the proofs use polyhedral geometry rather than restricted isometry or probabilistic signal models.

Partial integrality is itself established optimization machinery. Branch-and-bound selects fractional variables, imposes binary branches, and resolves LP relaxations; strong branching explicitly estimates bound improvement before choosing a branching variable [6]. Relax-and-fix methods likewise retain

selected integer variables while relaxing others [7]. Even the optimal-face interval geometry used by the pursuit has close prior art. Cloud Branching defines coordinate intervals over a cloud of alternative LP optima and partitions branching variables according to how those intervals intersect the branching interval [8]. For binary variables, when the cloud contains the full optimal face, its F_2 class consists exactly of coordinates that remain fractional throughout the face and are therefore guaranteed to improve the dual bound in both binary children. Gamrath, Berthold, and Salvagnin later analyze cloud intervals explicitly as projections of the optimal face and their implications for branching under dual degeneracy [9]. Thus we make no novelty claim for $[\ell_j, u_j]$, for the criterion $\ell_j > 0, u_j < 1$, or for the operation $x_D \in \{0,1\}^{|D|}$ by itself.

The contribution starts where those ingredients are combined into a different object: simultaneous exactification of a logarithmic multi-variable window while every outside variable remains globally relaxed. Theorem 3.3 converts the full-face distances into the quantitative additive bound

$$U(D) \leq U_{\mathrm{LP}} - \mu \sum_{j \in D} v_j \delta_j,$$

and Lemma 4.1 proves that, for a fixed retained support and window size, the pursuit-selected batch maximizes this rigorous guaranteed decrease among all admissible competing batches. Cloud Branching uses interval information to filter or rank *individual branching variables* and to reason about child-node bound improvement; the present theorem instead gives a globally valid simultaneous-window bound, preserves a single polynomial-time relaxation rather than growing a branch-and-bound tree, and couples that bound to the CoSaMP identify–merge mechanism. This quantitative additive window theorem, not the cloud intervals themselves, is the novelty claim.

The partially integral relaxation also differs from solving a smaller WDP. Every nonwindow bid remains in the optimization problem as a relaxed variable. Therefore $U(D_t)$ remains a valid upper bound on the *full* WDP optimum. This lets a logarithmic sparse window provide global information while retaining polynomial complexity.

The theoretical-computing positioning is supported by related optimisation work in *Theory of Computing Systems*. Elbassioni et al. develop linear-programming-based techniques for algorithmic mechanism design and explicitly treat combinatorial auctions through packing relaxations and approximation [14]; Ammann and Stiller study near-optimal auctions over inde-

pendence systems, including settings whose allocation problems are NP-hard and therefore require approximation [15]; and Fekete and Krupke derive approximation algorithms for NP-hard combinatorial optimization problems using LP/IP formulations and relaxations [16]. The present paper contributes a different synthesis: a fractional mathematical model guides sparse local exactification, while the same local exactification produces certified global information about the unknown integer optimum.

10. Discussion

The theory separates three logically different quantities. First, $\rho_0(\kappa)$ is a uniform worst-case ratio inherited from the protected seed. Second, $U_{\mathrm{LP}}/V_{\mathrm{C}}$ is the familiar generic LP certificate. Third,

$$\rho_{\mathrm{P}} := \frac{U_{\mathrm{P}}}{V_{\mathrm{C}}} \tag{74}$$

is created by the pursuit trajectory itself. Theorem 5.4 guarantees $\rho_{\mathrm{P}} \leq U_{\mathrm{LP}}/V_{\mathrm{C}}$, and Theorem 5.5 gives a concrete condition for strict inequality.

This distinction is essential. Merely making a logarithmic set of variables integral would still leave the central construction open to the objection that an arbitrary partial-integrality window, a relax-and-fix window, or a shallow branch-and-bound step could generate the same sandwich. Likewise, merely using the full-optimal-face intervals would reproduce the cloud-interval geometry already developed for MIP branching [8, 9]. Theorem 5.3 closes the remaining gap. Starting from that established interval geometry, it converts weighted distances from the binary values into an additive simultaneous-window lower bound on certificate improvement, and the identify rule maximizes the resulting mass $\mu\Psi(D)$ for fixed support memory and window cardinality. Thus the novelty lies in the quantitative multi-coordinate certificate theorem and its pursuit-optimal window selection, not in rediscovering cloud intervals or partial integrality.

The optimality-gap results in Section 6 are intentionally a posteriori. They do not assert that every instance has a 3% gap. They assert something stronger and more precise: once the polynomial algorithm returns a certificate satisfying (64) or (66), the unknown exact optimum is mathematically forced to lie within the stated relative distance. No subsequent exact computation is required to validate that conclusion.

The choice of logarithmic window size is a complexity device, not a statistical sparsity assumption. A linear-size exactification window would return to mixed-integer exponential worst-case behavior. An $O(\log m)$ window permits explicit enumeration at polynomial cost. Remark 3.3 shows that one can enlarge the window polynomially by increasing the constant multiplying $\log m$, trading computation for potentially tighter bounds while preserving worst-case polynomiality.

Finally, the method does not require LP uniqueness. This advantage is inherited from the cloud-interval viewpoint: both the exact characterization in Theorem 3.2 and the pursuit proxy in Definition 3.2 are defined on the full optimizer set $X^\star_{\mathrm{LP}}$, consistent with prior work that exploits alternative optima under dual degeneracy [8, 9]. The coordinate ranges $[\ell_j, u_j]$ therefore detect ambiguity robust to solver tie breaking. What is new in the present analysis is that these established intervals enter the weighted additive mass $\Psi(D)$ and hence the quantitative simultaneous-window bound of Theorem 3.3; a proxy built from one arbitrary fractional optimum would not support that conclusion.

11. Conclusion

This paper developed a deterministic polynomial-time sparse-pursuit framework for winner determination. Exact WDP is first relaxed globally. LP-CoSaMP then uses a CoSaMP-type identify–merge–re-estimate–prune cycle in which re-estimation restores integrality only on a logarithmic moving window. Partial integrality itself is standard, and the optimal-face coordinate ranges used here are the cloud intervals of the MIP branching literature [8, 9]. The central result is therefore not either ingredient in isolation, but their coupling to a weighted Hoffman bound and sparse-pursuit window selection. The weighted cloud-interval score is computed from the complete LP-optimal face, and Theorem 5.3 proves that the pursued logarithmic window maximizes a guaranteed additive tightening lower bound among all admissible windows of the same form and size. The logarithmic window keeps this exactification polynomial-time solvable.

The principal guarantee is the pursuit-specific interval

$$V_{\mathrm{C}} \le V^\star \le U_{\mathrm{P}} \le U_{\mathrm{LP}}.$$

It yields both absolute and relative optimality-gap certificates. In particular, the computable quantity $1-V_{\mathrm{C}}/U_{\mathrm{P}}$ upper-bounds the unknown exact relative

gap. A returned solution satisfying $V_C \geq 0.97U_P$ is therefore certified within 3% of exact optimality, while $U_P \leq 1.03V_C$ certifies the stricter bound below 2.913%. Exactness itself is certified when the two pursuit bounds coincide.

A protected approximation seed can be retained simultaneously, giving the adaptive ratio in Theorem 7.1. This cleanly separates known bounded-rank set-packing approximation theory from the pursuit contribution. The inherited hardness theorem explains why the appropriate general target is not a parameter-independent near-one worst-case factor, but a polynomially computable instance-specific certificate that can become near-exact when the actual relaxation geometry permits it.

Declarations

Ethics. Not applicable.
Data availability. Not available.
Code availability. Not applicable.
Funding. Not applicable.
Conflict of interest. No conflict of interest.
Consent to publish. Yes.
Author contribution. Angshul Majumdar: Conceptualization, methodology, formal analysis, investigation, writing–original draft, and writing–review and editing.

References

[1] T. Sandholm, Algorithm for optimal winner determination in combinatorial auctions, Artificial Intelligence 135 (1–2) (2002) 1–54. doi:10.1016/S0004-3702(01)00159-X.

[2] T. Sandholm, S. Suri, BOB: Improved winner determination in combinatorial auctions and generalizations, Artificial Intelligence 145 (1–2) (2003) 33–58. doi:10.1016/S0004-3702(03)00015-8.

[3] T. Sandholm, S. Suri, A. Gilpin, D. Levine, CABOB: A fast optimal algorithm for winner determination in combinatorial auctions, Management Science 51 (3) (2005) 374–390. doi:10.1287/mnsc.1040.0336.

[4] D. Lehmann, L. I. O'Callaghan, Y. Shoham, Truth revelation in approximately efficient combinatorial auctions, Journal of the ACM 49 (5) (2002) 577–602. doi:10.1145/585265.585266.

[5] E. Hazan, S. Safra, O. Schwartz, On the complexity of approximating k-set packing, Computational Complexity 15 (1) (2006) 20–39. doi:10.1007/s00037-006-0205-6.

[6] D. R. Morrison, S. H. Jacobson, J. J. Sauppe, E. C. Sewell, Branch-and-bound algorithms: A survey of recent advances in searching, branching, and pruning, Discrete Optimization 19 (2016) 79–102. doi:10.1016/j.disopt.2016.01.005.

[7] N. Absi, W. van den Heuvel, Worst case analysis of relax and fix heuristics for lot-sizing problems, European Journal of Operational Research 279 (2) (2019) 449–458. doi:10.1016/j.ejor.2019.06.010.

[8] T. Berthold, D. Salvagnin, Cloud branching, in: C. Gomes, M. Sellmann (Eds.), Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, Vol. 7874 of Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2013, pp. 28–43. doi:10.1007/978-3-642-38171-3_3.

[9] G. Gamrath, T. Berthold, D. Salvagnin, An exploratory computational analysis of dual degeneracy in mixed-integer programming, EURO Journal on Computational Optimization 8 (2020) 241–261. doi:10.1007/s13675-020-00130-z.

[10] O. Güler, A. J. Hoffman, U. G. Rothblum, Approximations to solutions to systems of linear inequalities, SIAM Journal on Matrix Analysis and Applications 16 (2) (1995) 688–696. doi:10.1137/S0895479892237744.

[11] P. Berman, A $d/2$ approximation for maximum weight independent set in d-claw free graphs, in: M. M. Halldórsson (Ed.), Algorithm Theory – SWAT 2000, Vol. 1851 of Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2000, pp. 214–219. doi:10.1007/3-540-44985-X_19.

[12] M. Neuwohner, The limits of local search for weighted k-set packing, Mathematical Programming 206 (2024) 389–427. doi:10.1007/s10107-023-02026-3.

[13] M. Neuwohner, Passing the limits of pure local search for weighted k-set packing, in: Proceedings of the 2023 Annual ACM-SIAM Symposium

on Discrete Algorithms (SODA), SIAM, Philadelphia, PA, 2023, pp. 1090–1137. doi:10.1137/1.9781611977554.ch41.

[14] K. Elbassioni, K. Mehlhorn, F. Ramezani, Towards more practical linear programming-based techniques for algorithmic mechanism design, Theory of Computing Systems 59 (2016) 641–663. doi:10.1007/s00224-016-9704-2.

[15] S. C. L. Ammann, S. Stiller, Near-optimal auctions on independence systems, Theory of Computing Systems 68 (2024) 1160–1179. doi:10.1007/s00224-024-10189-5.

[16] S. P. Fekete, D. Krupke, What goes around comes around: Covering tours and cycle covers with turn costs, Theory of Computing Systems 68 (2024) 1207–1238. doi:10.1007/s00224-024-10178-8.

[17] Y. H. Chan, L. C. Lau, On linear and semidefinite programming relaxations for hypergraph matching, Mathematical Programming 135 (1–2) (2012) 123–148. doi:10.1007/s10107-011-0451-5.

[18] D. Needell, J. A. Tropp, CoSaMP: Iterative signal recovery from incomplete and inaccurate samples, Applied and Computational Harmonic Analysis 26 (3) (2009) 301–321. doi:10.1016/j.acha.2008.07.002.