Quantum Graph Collision: An Upper Bound Below $n^{2/3}$

Shiva Kintali*

September 20, 2026

Abstract

Graph collision asks whether a known graph contains an edge whose endpoints are both marked, given query access to the marking bits. We prove an unconditional bounded-error quantum query upper bound of $\widetilde{O}(n^{13/20})$ for arbitrary n-vertex graphs, improving the general $O(n^{2/3})$ bound. A full-neighborhood decomposition controls squared-codegree tails and enables weighted stochastic responses that contract both edge and marked-volume budgets. The bound includes all sampling and validation queries; computation on the known graph may take exponential time.

1 Introduction

Given a known simple graph $G = (V, E)$ on n vertices and unknown marking bits $x \in \{0,1\}^V$, graph collision asks whether

$$\mathrm{GC}_G(x) = \bigvee_{uv \in E} (x_u \wedge x_v). \tag{1.1}$$

Only quantum queries to x are charged; all graph data are known. This is a monotone 2-DNF with repeated variables: its two-vertex certificates do not immediately give a square-root algorithm in the number of distinct vertices. Read-many formula bounds help clarify that distinction [11].

Motivation. Graph collision is a primitive in triangle finding and Boolean matrix multiplication [1, 6, 7]. A star with its center fixed marked contains unstructured search, giving the $\Omega(\sqrt{n})$ benchmark [4]. Its broader significance is illustrated by the lower-bound transfer to triangle detection of Balodis and Iraids [8], and the graph-collision-based conditional bounds for path and cycle containment of Cornelissen, Gilani, and Patro [9]. Our objective is to approach the square-root scale while paying every marking-dependent operation.

Past and related work. Magniez, Santha, and Szegedy provide the general $O(n^{2/3})$ quantum-walk benchmark [1]. Structural refinements include Belovs's $O(\sqrt{n}\,\alpha(G)^{1/6})$, with $\alpha(G)$ the independence number [2, Section 4], and Gavinsky–Ito's $O(\sqrt{n} + \sqrt{\alpha^*(G)})$, where $\alpha^*(G)$ is the maximum degree sum of an independent set [3]. Ambainis et al. obtain $O(\sqrt{n}\,w^{1/6})$ for positive treewidth w and $O(\sqrt{n} + \sqrt{\alpha^{**}(G)})$, where α^{**} minimizes the independent-set degree-sum bound inside a vertex cover [4]. Further results concern random graphs [3], a specified circulant subclass [4], and pseudo-deterministic collision search, which asks for a consistent witness [10]. For bipartite graphs with r vertices per side and s missing cross-edges, Jeffery et al. give $\widetilde{O}(\sqrt{r} + \sqrt{s})$ [6, 7]; the nested-walk paper supplies the rectangular bound used here [5, Appendix A]. Our supporting tools are amplitude estimation, verified search, robust OR, and span-program/state-conversion methods

*Email: shiva.kintali@gmail.com Homepage: shivakintali.github.io

[12, 13, 14, 15, 16]. Stochastic proximal methods provide optimization background [17], but neither they nor a matrix identity supply an uncharged marking-dependent oracle.

Contributions and comparison. The full-neighborhood extraction theorem pays for the *actual* rectangular dimensions using disjoint removed edge lists, while leaving a bounded squared-codegree row tail. The response-weighted variance estimate (Proposition 3.2), separate diagonal estimate, and entire-domain parameter inequalities (Lemma 5.1) pay for its algorithmic application. We prove the three required routines: a spectral dense-epoch contraction, a clipped source-sampling response with explicit entropy initialization, and a bounded-cover-count detector obtained from a dual-adversary construction and a count guard. Their combination gives the exponent 13/20, saving 1/60 relative to 2/3. Classical enumeration and optimization use only known data; no running-time bound is claimed. The decomposition and its integration with these paid routines form the main contribution.

Proof strategy. Paid neighborhood sampling first controls marked degree moments. We partition the residual into actual-degree pieces, extract full-neighborhood modules whenever a codegree tail is too large, and test their union by complete rectangular algorithms. The remaining row-tail bound controls the offdiagonal Schur-field variance, while frozen moment constraints pay the diagonal bias. The proved response and outside routines then contract edge and marked-volume budgets. Exact exponent inequalities pay every branch, and smaller-instance aspect recursion closes the global argument without a same-size recursive call.

Main results. Our principal result is the following bound for arbitrary known graphs.

Theorem 1.1 (Unconditional arbitrary-graph bound)**.** *For every known simple graph on n vertices,*

$$Q_2(\mathrm{GC}_G) = \widetilde{O}(n^{13/20}). \tag{1.2}$$

The algorithm has a worst-case fixed cap on every input and error at most 1/3. The bound charges only marking queries.

The main structural ingredient is a decomposition of bipartite graphs. For source side S and target side R, write $d_u = |N(u)|$ and $c_{uv} = |N(u) \cap N(v)|$.

Theorem 1.2 (Unconditional full-neighborhood tail removal)**.** *Let the graph have m edges, and let $D \geq 1$ be a known upper bound on its source degrees. For real $H > 0$ and $1 \leq t \leq D$, graph-only preprocessing produces modules with disjoint removed edge lists. Their union has a complete collision detector with cap*

$$\widetilde{O}\left(\sqrt{m}\left[\left(\frac{D}{tH}\right)^{1/6} + t^{-1/2} + \sqrt{D/H}\right]\right). \tag{1.3}$$

The final residual satisfies, for every source u,

$$\sum_{v \neq u:\ c_{uv} \geq t} c_{uv}^2 \leq H d_u. \tag{1.4}$$

No marking, codegree, or three-walk promise is assumed. If $m = 0$, or $t > D$, no removal is necessary; for $t > D$ the tail in (1.4) is empty.

Organization. Section 2 gives the model, query primitives, and graph preparation. Section 3 proves extraction and the residual-field estimates. Section 4 explains the contraction routines and states their guarantees. Section 5 presents the algorithm, justifies its parameters, and proves the main theorem. Appendix A contains the auxiliary proofs; Appendix B proves the bounded-cover-count detector; Appendix C supplies the dense and weighted-response implementations.

2 Preliminaries and graph preparation

2.1 Query model and primitives

Throughout, graphs are finite, simple, and known. The marking oracle is

$$O_x|v, b, w\rangle = |v, b \oplus x_v, w\rangle.$$

Instances with at most one vertex are answered without queries; in asymptotic parameter formulas we assume $n \geq 2$. Any known edgeless instance is also answered without queries. Classical computation on graph data is not charged to query complexity; it need not be polynomial time. A query bound is a worst-case bound on every input and every execution branch, not an expected stopping time. Bounded error means error at most $1/3$ on every input in the declared domain; $Q_2(f)$ denotes the minimum worst-case query cap achieving that error for f. Unless otherwise specified, $0 < \varepsilon < 1/3$ is a prescribed error budget. Error reduction uses a predetermined $O(\log(2/\varepsilon))$ number of independent repetitions. The notation $\widetilde{O}$ suppresses fixed powers of $\log(2n)$ and $\log(2/\varepsilon)$ only, and $\widetilde{\Theta}$ allows multiplicative factors and reciprocal factors of these forms. All parameters in the application have polynomial range. A NO marking is one whose marked support is an independent set; no probabilistic independence among marking bits is assumed.

For an edge-list graph H and marking x, its marked degree moments are

$$W(H, x) = \sum_v x_v \deg_H(v), \qquad S_2(H, x) = \sum_v x_v \deg_H(v)^2.$$

When the graph and marking are fixed, we write W and S_2. Later constraints on response vectors use declared bounds for these moments. In particular, Σ_2 denotes a graph-known upper allowance for $S_2(H, x)$, not a queried value of that moment.

Definition 2.1 (Complete and one-sided-promised routines)**.** A complete graph-collision routine is correct on every marking. A NO-volume-promised routine is correct on every YES marking and on the NO markings satisfying its declared volume bound. It still has a fixed query cap on markings outside that promise.

An edge-list subgraph contains exactly its listed edges. Its vertex set is not an instruction to add all edges induced in the original graph. We permit copied labels with a known address map π to physical vertices, and known-zero dummy labels.

Copied labels. Reversibly computing $\pi(i)$, querying O_x, and uncomputing implements a copied-label query with $O(1)$ physical queries. Known-zero dummy queries act as the identity. A complete algorithm on independent copied bits stays correct on the restriction $y_i = x_{\pi(i)}$; this does not cache answers at a shared address.

Lemma 2.2 (Rectangular collision, published import)**.** *A known bipartite edge-list graph with side sizes $a, b \geq 1$ has a complete collision detector costing*

$$O((ab)^{1/3} + \sqrt{a} + \sqrt{b}) \tag{2.1}$$

queries at constant error, with logarithmic overhead for smaller error.

Reduction to the published bound. Transpose if necessary to have $a \le b$. If $b \le a^2$, use [5, Lemma A.1], which gives $O((ab)^{1/3})$. If $b > a^2$, read all a bits and search the known union of neighbors of marked vertices on that side. This costs $O(a + \sqrt{b}) = O(\sqrt{b})$. The latter branch handles the aspect ratios not covered by that lemma. This proof only completes its range; it does not reprove the walk. □

Lemma 2.3 (OR of known unequal clocks)**.** *Suppose $r \le \mathrm{poly}(n)$ complete Boolean routines have graph-known fixed caps $q_1, \dots, q_r \ge 1$, polynomially bounded in n, and may query shared bits. Their OR can be decided with fixed cap*

$$\widetilde{O}\left(\sqrt{\sum_{j=1}^{r} q_j^2}\right). \tag{2.2}$$

An empty list costs zero queries.

The proof is given in Appendix A.

Lemma 2.4 (Known-cover volume terminal)**.** *Let U be a known vertex cover of H. For a declared $L \ge 0$, there is a routine with cap*

$$O\left(\sqrt{|U| + L}\log(2/\varepsilon)\right) \tag{2.3}$$

with error at most ε on all YES inputs and on NO inputs with $\sum_{u\in U} x_u \deg_H(u) \le L$. The cap is fixed on all inputs, including those outside this domain.

The proof is given in Appendix A.

Adaptive error accounting. All routines receive error budgets before execution. If call j has conditional failure probability at most ε_j on every preceding good NO transcript, the probability of a first failure is at most $\sum_j \varepsilon_j$: condition on the preceding transcript and sum the disjoint first-failure events. No independence is required. The same accounting applies to all-sign collision preservation when each call has that guarantee uniformly conditional on every preceding good history, on either input sign.

2.2 Graph preparation

We collect the preparation steps used before and during the algorithm. Their proofs are given in Appendix A.

Paid source-neighborhood cleanup.

Lemma 2.5. *Let S be a known source subset of a graph on at most N labels. For known $0 < K \le N$ and $\theta > 0$, a routine with all-input cap*

$$\widetilde{O}\left(\sqrt{N} + \frac{\sqrt{NK}}{\theta}\right) \tag{2.4}$$

returns YES or a known induced subgraph H. On every NO marking with at most K marked vertices in S, with probability $1-\varepsilon$ it returns H and every remaining vertex has at most θ marked neighbors in S. On either sign, deleting a marked vertex has probability at most ε; thus every collision survives unless a collision is reported or a charged failure occurs.

The proof is given in Appendix A.

Global degree moments and their inheritance.

Theorem 2.6. *For fixed $1/2 < p < 1$, fixed-cap $\widetilde{O}(n^p)$ preprocessing returns YES or a known induced subgraph H. On NO markings, with probability $1-\varepsilon$ it returns H and*

$$W(H,x) := \sum_v x_v \deg_H(v) \le C_p n^{2-p}, \tag{2.5}$$

$$\sum_{v:\deg_H(v)\ge d} x_v \deg_H(v) \le C_p n^{3-2p}/d \quad (d \ge 1), \tag{2.6}$$

$$S_2(H,x) := \sum_v x_v \deg_H(v)^2 \le C_p n^{3-2p}\log(2n). \tag{2.7}$$

Here C_p is a known constant depending only on p. Every collision is preserved on either sign with preservation failure probability at most ε, unless YES is reported.

The proof is given in Appendix A.

Lemma 2.7 (Moment inheritance)**.** *Edge deletion never increases W or S_2. Splitting the incidences at a vertex into repeated-bit clones preserves W and does not increase S_2. For an exact edge partition $E = \dot{\bigcup}_j E_j$,*

$$\sum_j W(E_j,x) = W(E,x), \qquad \sum_j S_2(E_j,x) \le S_2(E,x).$$

The proof is given in Appendix A.

Clones, orientation, and actual-degree partition. At an epoch with edge budget M, first discard isolated vertices. Split each remaining vertex of degree d_v into $\lceil d_v/D \rceil$ nonempty clones, $D = \lceil M/n \rceil$, distributing its incident edges into groups of size at most D. If $m \le M$, the total number of clones is at most $n + 2m/D \le 3n$. An original edge joins its two designated clones, so the clone graph is simple and has exactly the original collision predicate. Orient each edge from its endpoint of larger full clone degree, breaking ties by known labels. Dyadically bucket its source degree. Each bucket is represented bipartitely using role copies; every original edge occurs in exactly one bucket. Both bucket degrees are at most a fixed multiple of its degree scale d, and its edge count is $O(nd)$. None of these operations queries the marking.

Lemma 2.8 (Exact almost-biregular partition)**.** *Let a known bipartite graph have at most n' vertices per side and $m_0 \ge 1$ edges. Put $h = 1 + \lceil \log_2 \max(1,n') \rceil$. There is a graph-only exact edge partition into at most $1 + \lceil 2h^2 \ln m_0 \rceil$ nonempty pieces. Each piece has dyadic caps d_L, d_R, no isolated displayed vertices, and actual degrees in $[d_L/(8h^2), d_L]$ on the left and $[d_R/(8h^2), d_R]$ on the right. Writing e for the edge count of that piece, its side sizes lie in $[e/d_L, 8h^2 e/d_L]$ and $[e/d_R, 8h^2 e/d_R]$, respectively.*

The proof is given in Appendix A.

3 Extraction and variance control

3.1 Full-neighborhood extraction

For a bipartite graph with source side S and target side R, let $d_u = |N(u)|$ and $c_{uv} = |N(u) \cap N(v)|$ for $u, v \in S$. In the proof of Theorem 1.2, degrees and codegrees are always measured in the *current* graph. The symbol D remains a fixed upper bound throughout. Vertex-set overlaps among modules are allowed.

Proof of Theorem 1.2. The edgeless and $t > D$ cases are immediate. We may replace D by $\min\{D, \max(1, |R|)\}$: this can only improve the claimed cap, and if t then exceeds the new bound the tail is empty. Hence assume $1 \le t \le D \le \max(1, |R|)$. Let $J = 1 + \lceil \log_2 D \rceil$ and $\eta = H/(4J)$. For a source u of positive current degree d and a band $c = 2^k t \le D$, where k is a nonnegative integer, let

$$A = \{v \ne u : c \le c_{uv} < 2c\}, \qquad a = |A|.$$

If $ac^2 > \eta d$, record the current edge-list graph $F = G[A, N(u)]$ and remove exactly its edges. Write $e = |E(F)|$. Every vertex $v \in A$ has between c and $2c$ neighbors in $N(u)$, so

$$ac \le e < 2ac, \qquad e > \eta d/c, \qquad c \le d \le D. \tag{3.1}$$

The last inequality uses $a > 0$ at a triggering band. In particular, $e > \eta > 0$; an integer edge is deleted. Recompute the whole residual after each extraction. A known deterministic tie-breaking rule makes this a finite graph-only procedure with at most m modules.

At termination every eligible band has $ac^2 \le \eta d$. Its squared-codegree contribution is less than $4ac^2 \le 4\eta d$. The dyadic bands cover all codegrees at least t and there are at most J of them. This proves (1.4) using the final degree and final codegrees. The exclusion of $v = u$ is essential.

For the selected modules, index their quantities by j. Their removed edge lists are disjoint, hence $\sum_j e_j \le m$. Equation (3.1) gives

$$\sum_j a_j \le m/t, \qquad \sum_j d_j \le mD/\eta. \tag{3.2}$$

Indeed $d_j < e_j c_j/\eta \le e_j D/\eta$. Also

$$\begin{aligned}(a_j d_j)^{2/3} &\le (e_j d_j/c_j)^{2/3} \\ &= e_j (d_j/c_j)^{2/3} e_j^{-1/3} \\ &\le e_j \eta^{-1/3} (d_j/c_j)^{1/3} \le e_j (D/(t\eta))^{1/3}.\end{aligned} \tag{3.3}$$

Thus

$$\sum_j (a_j d_j)^{2/3} \le m(D/(t\eta))^{1/3}. \tag{3.4}$$

For module j, use the cheaper of the rectangular detector and reading its $a_j + d_j$ displayed bits. Its graph-known cap is polynomially bounded in the original vertex count, and by Lemma 2.2 its square is $\widetilde{O}(a_j + d_j + (a_j d_j)^{2/3})$. Lemma 2.3, (3.2), and (3.4) bound their OR by

$$\widetilde{O}\left(\sqrt{m/t + mD/\eta + m(D/(t\eta))^{1/3}}\right).$$

Since $H/\eta = 4J$ is logarithmic, and $\sqrt{r+s+z} \le \sqrt{r} + \sqrt{s} + \sqrt{z}$, this is (1.3). There are at most m modules, and their capped clocks have logarithmically many dyadic bins even for arbitrary real $H > 0$. Every tested module is an actual edge-list subgraph. Complete detectors remain complete on shared marking addresses, so no removed marked edge is silently discarded. The residual and removed lists form an exact edge partition. □

Remark 3.1. Using maximum possible rectangle dimensions instead of (3.2)–(3.4) loses the density charge. Conversely, overlap of vertex sets does not violate the proof: only disjointness of removed *edges* is used for $\sum e_j \le m$. This theorem does not use the additional algorithmic routines below.

3.2 Residual field estimates

Let B be the target-by-source incidence matrix of a residual, so $(B^\top B)_{uv} = c_{uv}$ and $c_{uu} = d_u$. Source vectors are denoted by x_s, a_s. In this section $\tau > 0$, $q \geq 1$, and all count, load, and volume bounds are nonnegative. All quantities are known matrices, declared bounds, or algebraic consequences of explicit inequalities. The quadratic estimate below supplies the variance bound in Lemma 4.2. Appendix C constructs its stochastic estimator and proves the conditional noise guarantee.

Proposition 3.2 (Response-weighted offdiagonal variance)**.** *Suppose $x_s \in \{0,1\}^S$, $0 \leq a_s \leq 1$, $Bx_s \leq \theta\mathbf{1}$, $Ba_s \leq \theta\mathbf{1}$, and $\sum_u x_u d_u \leq L$. If (1.4) holds, then*

$$\sum_{u \neq v} x_u a_v c_{uv}^2 \leq (t\theta + H)L. \tag{3.5}$$

Consequently the offdiagonal field matrix $R_s = ((B^\top B) - \mathrm{diag}(d_u))/(\tau q)$ obeys

$$\sum_{u,v} a_v x_u (R_s)_{vu}^2 \leq V := \frac{(t\theta + H)L}{(\tau q)^2}. \tag{3.6}$$

Proof. For the low-codegree offdiagonal terms use $c_{uv}^2 \leq t c_{uv}$:

$$\begin{aligned}\sum_{u \neq v : c_{uv} < t} x_u a_v c_{uv}^2 &\leq t \sum_{u,v} x_u a_v c_{uv} = t(Bx_s)^\top (Ba_s) \\ &\leq t\theta \mathbf{1}^\top B x_s = t\theta \sum_u x_u d_u \leq t\theta L.\end{aligned}$$

For the high terms, $a_v \leq 1$ and (1.4) give $\sum_u x_u \sum_{v \neq u : c_{uv} \geq t} a_v c_{uv}^2 \leq H \sum_u x_u d_u \leq HL$. Add these bounds and divide by $(\tau q)^2$. If $t > D$ no extraction is necessary: all codegrees are at most D, and the entire offdiagonal sum is at most $D\theta L$. In particular $D\theta \leq H$ when $t = H/\theta > D$. □

Proposition 3.3 (Schur correction and omitted diagonal)**.** *In target/source order let*

$$A = \begin{pmatrix} 0 & B \\ B^\top & 0 \end{pmatrix}, \qquad C = 0 \oplus B^\top B/(\tau q), \qquad T_0 = \tau q I.$$

Then $C \succeq 0$ and $A + T_0 + C \succeq 0$. For a Boolean marking $x = (x_t, x_s)$, if its marked source count is at most K, the maximum source degree is at most d, and $a_s \in [0,1]^S$, omitting the diagonal from the unknown field has comparison error

$$\left| ((R - C)x)^\top (a - x) \right| \leq dK/(\tau q), \tag{3.7}$$

where $R = 0 \oplus R_s$. If $K \leq C_K N/q$, this is at most a prescribed fixed fraction of τN provided

$$\tau^2 q^2 \geq C_0 d \tag{3.8}$$

with sufficiently large C_0.

Proof. $B^\top B$ is a Gram matrix. The Schur complement of the positive target block $\tau q I$ in $A + T_0 + C$ is exactly $\tau q I$. Equivalently its quadratic form at (y, z) is

$$\tau q \|y + Bz/(\tau q)\|_2^2 + \tau q \|z\|_2^2 \geq 0.$$

The omitted field is $-(0 \oplus \mathrm{diag}(d_u)x_s)/(\tau q)$. At marked source coordinates $|a_u - x_u| \leq 1$; at unmarked ones the omitted field is zero. Thus its absolute comparison error is at most $\sum_u d_u x_u/(\tau q) \leq dK/(\tau q)$. Substitute $K \leq C_K N/q$ to obtain (3.8). There is no free diagonal query in this argument. □

Lemma 3.4 (Frozen count and diagonal budgets)**.** *Suppose $d, L, \Sigma_2 > 0$, and each displayed real coordinate in a degree piece has frozen degree at least d/g, where $g \geq 1$ is a fixed polylogarithmic factor. Fix all dummy coordinates to zero and impose the known constraints*

$$0 \leq a \leq 1, \qquad \sum_v d_v^{\mathrm{fr}} a_v \leq L, \qquad \sum_v (d_v^{\mathrm{fr}})^2 a_v \leq \Sigma_2.$$

There is a conservative scalar $q \geq 1$ of scale

$$q = \widetilde{\Theta} \max\{1, nd/L, nd^2/\Sigma_2\} \tag{3.9}$$

such that the total feasible count is at most $C_K N/q$, and $\sum_v \tau q a_v \leq C_T \tau N$, for fixed constants and $N = \Theta(n)$. Bounds remain valid after edge deletions when the frozen constraints are kept.

Proof. The three constraints bound the count by $\min(N, gL/d, g^2\Sigma_2/d^2)$. Choose explicitly

$$q = \max\{1, nd/(gL), nd^2/(g^2\Sigma_2)\}.$$

Writing $q_0 = \max\{1, nd/L, nd^2/\Sigma_2\}$ gives $q_0/g^2 \leq q \leq q_0$, hence (3.9). Take C_K large enough that $C_K N \geq \max(N, n)$, using the fixed bounds implicit in $N = \Theta(n)$. The count bound is then $C_K N/q$. Multiplication by τq proves the diagonal mass budget. Free dummy coordinates are fixed to zero. The constraints are not redefined with smaller residual degrees, so deletion changes neither implication. □

4 Contracting the remaining graph

The algorithm maintains an edge budget M and, on a good NO history, a marked-volume bound $W \leq \alpha M$. A contraction either reports YES or retains a known edge list with both budgets halved. For dense graphs, ordinary spectral-column sampling constructs the retained set. For intermediate pieces, the extraction and variance bounds of Section 3 enable a source-supported weighted response. In both cases, the response is followed by a collision test on the edges outside the retained set. The bounded-cover-count detector supplies this test when the validated outside marking count is small. Section 5 chooses the parameters and combines these operations; the statements here specify all guarantees it needs.

The three lemmas below specify the oracle costs, output guarantees, and sign conventions used by Theorem 1.1. The dense and weighted-response proofs are given in Appendix C; the bounded-cover-count proof is given in Appendix B. They hold uniformly for the graph-known parameter choices in Section 5, at each prescribed error budget and conditional on the current classical history. All returned sets and edge lists are classical outputs, and all procedures have finite descriptions.

4.1 Dense contraction

Lemma 4.1 (Dense epoch)**.** *For $\alpha = \alpha_0 n^{-c} \leq 1$ with fixed $c \geq 0$, $m \leq M$, and the good-NO promise $W \leq \alpha M$, whenever $M \geq n^{7/4}$ there is a finite, all-input-fixed-cap, prescribed-error routine that reports YES or returns a known subset of the current edge list with*

$$m_{\mathrm{new}} \leq M/2, \qquad W_{\mathrm{new}} \leq \alpha M/2.$$

Its complete outside tests preserve every YES collision unless detected. Its fixed cap, summed over logarithmically many dense epochs, is

$$\widetilde{O}(n^{1/4+3c} + n^{2/3-c/6} + \sqrt{n}). \tag{4.1}$$

On every promised NO input, a false YES or failure of either contraction bound has probability at most the allocated error. On either input sign, an undetected disappearance of a collision has probability at most its allocated preservation error. These budgets include sampling, finite optimization, and validation.

4.2 Weighted response for intermediate pieces

Lemma 4.2 (Source-supported weighted marked response)**.** *Consider the bipartite matrices of Proposition 3.3 and a known rational compact polytope $\mathcal{F} \subseteq [0,1]^N$. Impose the frozen constraints of Lemma 3.4, fix dummy coordinates to zero, and include $Ba_s \le \theta \mathbf{1}$. The NO correctness promise is $x \in \mathcal{F}$. Take known rational $0 < K \le N$, $\tau > 0$, $0 < \beta \le 1$, and $V \ge 0$; the matrices and polytope constraints are rational as well. Assume every feasible source vector has count at most the declared $K = O(N/q)$. For every promised NO marking x and every $a \in \mathcal{F}$, assume the quadratic quantity in (3.6) is at most the declared V, and (3.8) pays the omitted diagonal. The source-supported routine has all-input cap*

$$\widetilde{O}\left(1 + \sqrt{F_0} + F_0\sqrt{K/N} + \beta^{-1/2}\right), \qquad F_0 = \frac{KV}{\tau^2 N}. \tag{4.2}$$

It reports YES or returns a known retained vertex set P with

$$e(P) \le C_E \tau N/\beta^2, \quad W(P,x) \le C_W \tau N/\beta, \quad |\{v \notin P : x_v = 1\}| \le C_O \beta N. \tag{4.3}$$

The edge cap is deterministic. Fresh fixed-budget count and incidence-volume validations certify both other caps on either sign except for allocated error. A failed validation may report YES; on promised NO inputs its probability is charged to the same error budget. No marking-dependent validation is free. On promised NO inputs the probability of reporting YES is at most the allocated error. Source-only proximal entropy is permitted; the cutoff objective still includes both *sides. The constants C_E, C_W, C_O are known, positive, and absolute. All confidence, clipping, precision, and finite-trial losses are fixed logarithms.*

Appendix C proves this statement using the total frozen diagonal mass and source-only entropy. The scalar τ in the response tolerance remains distinct from τq, and the cutoff objective includes the unsampled target coordinates.

4.3 Detecting collisions outside the retained set

Lemma 4.3 (Universal bounded-cover-count detector)**.** *Let H be any known simple edge-list graph on at most N labels, let U be a known vertex cover of H, and let $0 \le K \le N$ be a declared bound. There is a finite detector with all-input fixed cap*

$$\widetilde{O}(\sqrt{N}\, K^{1/6} + \sqrt{N}) \tag{4.4}$$

and error at most ε on every YES marking and every NO marking satisfying $\sum_{u \in U} x_u \le K$. The clock includes all required validation and error reduction. Its guarantee is uniform under copied-label restrictions.

In particular, when validated response outputs give $K = O(\alpha n)$ and $N = \Theta(n)$, this supplies the outside clock $\widetilde{O}(n^{2/3}\alpha^{1/6} + \sqrt{n})$ used in the global proof. This is a count promise on a specified cover of an arbitrary graph, not an independence-number or codegree assumption. The all-YES domain is essential; the count promise restricts only NO inputs.

5 Proof of the main theorem

5.1 Algorithm overview

The target query scale is $T = n^{13/20}$, and the volume parameter has scale $\alpha = n^{-1/10}$ up to a fixed constant. After the moment preprocessing of Theorem 2.6, the algorithm works in epochs with edge budget M and NO-volume allowance $L = \alpha M$. Each nonterminal epoch halves both allowances, so there are only logarithmically many epochs.

When $M \geq n^{7/4}$, the dense routine performs the contraction. When M is at most a fixed multiple of T^2/α, the cover-volume terminal solves the remaining graph directly. Between these regimes, cloning and the exact partition of Section 2 produce almost-biregular pieces. Unbalanced pieces are handled by smaller recursive calls. Balanced pieces are processed by edge search at low degree, or by cleanup, extraction, and weighted response at higher degree. Every set of discarded edges is tested for a collision before deletion.

The response parameters must simultaneously pay for the removed modules, the stochastic response, and the outside test, while making the retained edge and volume budgets small enough to sum across all pieces. The next subsection establishes these payments over the entire intermediate range. We then state the recursive estimate and give the complete proof, including correctness on both input signs.

5.2 Parameter choices and query costs

All scalar data supplied to the response routine are chosen rational. The displayed power formulas specify ideal scales: choose rational values within fixed factors for T, α, round the moment allowances upward, use integer polylogarithmic degree factors, and compute q and β from these actual values. Define τ by the exact allocation formula (5.5). Round the load threshold upward within a factor two, choose a rational H within a fixed factor of (5.6), and set $t = H/\theta$ exactly. These choices retain feasibility of the true marking and all cost bounds by the perturbation argument below. They give a rational response polytope and rational field matrices. All absolute output constants are fixed positive integer upper bounds, so the allocation constant κ is rational as well. For readability we use the ideal scale notation in exponent formulas. Set

$$p = 13/20, \qquad c = 1/10 = 4 - 6p, \qquad T = n^p, \qquad \alpha = \alpha_0 n^{-c}, \qquad 0 < \alpha_0 \leq \beta_0 < 1. \tag{5.1}$$

The scalar α is an epoch volume parameter, not the independence number $\alpha(G)$ appearing in the introduction. For a balanced intermediate piece, let $M = n^\mu$ and $d = n^\delta$. At the level of polynomial exponents, suppress fixed and logarithmic factors and write $L = \alpha M$ with exponent $\ell = \mu - c$. The terminal and dense boundaries give

$$2p + c \leq \mu \leq 7/4, \quad 0 \leq \delta \leq \mu - 1, \qquad \ell = \mu - c. \tag{5.2}$$

Moment preprocessing supplies the known allowance $\Sigma_2 = C_p n^{3-2p} \log(2n)$, so (3.9) has exponent

$$Q = \max\{0, L_q, Q_q\}, \qquad L_q = 1 + \delta - \mu + c, \qquad Q_q = 2\delta + 2p - 2. \tag{5.3}$$

Choose

$$\beta = \max\{\alpha, \min(\beta_0, T^2/(nd))\}, \qquad b = \min\{c, \max(0, \delta + 1 - 2p)\}, \tag{5.4}$$

so $\beta = n^{-b}$ at exponent level. Let $\mathcal{J} \geq 1$ be a graph-known polylogarithmic upper bound on the number of pieces in one epoch, and put

$$\tau = \frac{\kappa \alpha \beta M}{\mathcal{J} N}, \qquad \kappa = \frac{1}{4 \max\{1, C_E, C_W\}}. \tag{5.5}$$

Here $N = \Theta(n)$ is the padded piece size. This explicit allocation has scale $\widetilde{\Theta}(\alpha\beta M/n)$ and will be summed in Section 5. The raw load $\theta_0 = N/(T\sqrt{q})$ has exponent $\vartheta = 1 - p - Q/2$. From (5.5), τ has exponent $\ell - b - 1$. Define

$$H = T\tau^4 q^{7/2}/L, \qquad R = p + 3\ell - 4b - 4 + 7Q/2. \tag{5.6}$$

For the ideal exponent calculation in the nonterminal load branch take $\theta = \theta_0$ and $t = H/\theta$, so $t\theta + H = 2H$. Real thresholds are permitted in Theorem 1.2; dyadic rounding down by a factor less than two also preserves $t\theta \le H$. Since the actual piece has $m = O(nd)$ and $D \le d$, the three summands in (1.3) have exponents

$$P_1 = \tfrac{1}{2} + \tfrac{2}{3}\delta + \tfrac{1}{6}\vartheta - \tfrac{1}{3}R, \tag{5.7}$$
$$P_2 = \tfrac{1}{2} + \tfrac{1}{2}\delta + \tfrac{1}{2}\vartheta - \tfrac{1}{2}R, \tag{5.8}$$
$$P_3 = \tfrac{1}{2} + \delta - \tfrac{1}{2}R. \tag{5.9}$$

These are exact affine expressions for the ideal polynomial exponents within each max/min cell. Constant and logarithmic factors are not being identified with one. Each such factor shifts a logarithm to base n by $O(\log\log(2n)/\log n)$. The maps defining Q, b, ϑ, R, P_i use finitely many affine operations and maxima or minima, and are therefore Lipschitz with absolute constants. Consequently all upper payments proved below lift to $\widetilde{O}(n^p)$ for the actual parameters, and the strictly positive threshold and diagonal margins survive beyond a fixed base range. The next argument controls the whole domain, without relying on a grid. This perturbation argument pays logarithmic losses. The implementations are proved separately in the appendices.

Lemma 5.1 (Entire-domain payment)**.** *For every point in* (5.2) *with* $\delta \ge u := 2p - 1 = 3/10$,

$$P_1, P_3 \le 13/20, \qquad P_2 \le 3/5, \qquad R - \vartheta \ge 1/10, \qquad \delta - 2b \ge 1/5. \tag{5.10}$$

If $\delta \le u$, *edge search costs* $\widetilde{O}(n^p)$ *without field processing.*

Proof. Since $Q \ge L_q, Q_q$, the following weighted inequalities hold even when a floor is negative:

$$\tfrac{5}{4}Q \ge L_q + \tfrac{1}{4}Q_q, \quad 2Q \ge \tfrac{3}{2}L_q + \tfrac{1}{2}Q_q, \quad \tfrac{7}{4}Q \ge \tfrac{3}{2}L_q + \tfrac{1}{4}Q_q. \tag{5.11}$$

Expansion of (5.7)–(5.9) gives

$$\begin{aligned} P_1 &= 2 - p/2 + 2\delta/3 - \mu + c + 4b/3 - 5Q/4, \\ P_2 &= 3 - p + \delta/2 - 3\mu/2 + 3c/2 + 2b - 2Q, \\ P_3 &= 5/2 - p/2 + \delta - 3\mu/2 + 3c/2 + 2b - 7Q/4. \end{aligned}$$

Apply the respective inequalities in (5.11) to obtain

$$P_1 \le 3/2 - p - 5\delta/6 + 4b/3, \tag{5.12}$$
$$P_2 \le 5/2 - 2p - 2\delta + 2b, \tag{5.13}$$
$$P_3 \le 3/2 - p - \delta + 2b. \tag{5.14}$$

For $u \le \delta \le u + c = 2/5$ we have $b = \delta - u$. The three right sides have slopes $1/2, 0, 1$, respectively. Their values at $\delta = 2/5$ are $13/20, 3/5, 13/20$. For $\delta \ge 2/5$ we have $b = c$; all three expressions decrease. This proves the cost inequalities over the full domain.

Similarly $4Q \ge 3L_q + Q_q$ implies

$$\begin{aligned} R - \vartheta &= 2p + 3\mu - 3c - 4b - 5 + 4Q \\ &\ge 4p - 4 + 5\delta - 4b. \end{aligned}$$

On $[3/10, 2/5]$ the last expression is $\delta - 1/5 \ge 1/10$; above $2/5$ it is $5\delta - 9/5 \ge 1/5$. On the first interval $\delta - 2b = 2u - \delta \ge 1/5$; above it $\delta - 2b = \delta - 2c \ge 1/5$. Finally if $\delta \le u$, then $m = \widetilde{O}(nd) \le \widetilde{O}(n^{2p})$. A search over the known edge list uses $O(\sqrt{m})$ oracle calls to the two-endpoint predicate and is complete on every marking. □

Corollary 5.2 (Three separate payments)**.** *In the high-degree, nonterminal load branch, the module detector costs $\widetilde{O}(T)$, the threshold t is at least $n^{1/10}$ up to logarithms, and the omitted diagonal satisfies (3.8) beyond a fixed base range. Under Lemma 4.2, the response main and secondary clocks cost $\widetilde{O}(T)$.*

Proof. The first two conclusions follow from Lemma 5.1, with the preceding logarithmic perturbation convention. For the diagonal, $Q \ge L_q$ gives

$$\log_n(\tau q) = \ell - b - 1 + Q \ge \delta - b,$$

at exponent level. Hence $\tau^2 q^2/d$ has exponent at least $\delta - 2b \ge 1/5$. This strict polynomial margin absorbs the fixed and logarithmic constants in (3.8).

Equations (3.6) and (5.6), together with $K = O(N/q)$ and $t\theta = O(H)$, imply

$$F_0 \le \widetilde{O}(HL/(\tau^4 q^3)), \qquad F_0\sqrt{K/N} \le \widetilde{O}(HL/(\tau^4 q^{7/2})) = \widetilde{O}(T).$$

For the secondary clock, $\theta_0 = N/(T\sqrt{q}) > 1$ implies $q < (N/T)^2$. Since (5.6) gives $F_0 = \widetilde{O}(T\sqrt{q})$,

$$\sqrt{F_0} \le \widetilde{O}(\sqrt{T}\, q^{1/4}) \le \widetilde{O}(\sqrt{N}) = \widetilde{O}(T).$$

Also $\beta^{-1/2} \le \alpha^{-1/2} = O(n^{1/20})$. All confidence and composition overheads are fixed powers of logarithms, as proved by the implementation lemmas. □

The module-cost inequalities are tight at $(\mu, \delta) = (7/5, 2/5)$, where $P_1 = P_3 = 13/20$ and $P_2 = 3/5$.

5.3 Smaller-instance recursion

Lemma 5.3 (Aspect-ratio induction)**.** *Fix $p > 1/2$. Suppose a parent execution uses at most $C_J \log^k(2n)$ piece calls in total. Every piece has side sizes $1 \le a \le b \le C_N n$. Choose a fixed B with $B(p - 1/2) > k + 3$ and call a piece unbalanced if $a/b \le \log^{-B}(2n)$. Under the inductive cap $As^p \log^K(2s)$ for every $s < n$, all unbalanced calls together cost at most $(A/4)n^p \log^K(2n)$ for sufficiently large n, with fixed total error. Balanced pieces must be handled nonrecursively.*

The proof is given in Appendix A.

5.4 Completion of the proof

Proof of Theorem 1.1. We combine the proved routines, keeping their promises explicit. Set the parameters in (5.1); for any fixed finite initial range read all input bits. All calls are assigned conditional error budgets from a fixed polylogarithmic call bound.

1. Initial invariant and epoch boundaries. Apply Theorem 2.6 at exponent p. On a good NO history, the residual satisfies $W = O(n^{2-p})$ and $S_2 = \widetilde{O}(n^{3-2p})$. Edge deletion and splitting preserve these bounds by Lemma 2.7. Initialize the edge budget $M = \lceil n^2 \rceil$ and volume allowance $L = \alpha M$. For sufficiently large n, the first-moment bound fits this allowance, since $p > c$. Each successful nonterminal epoch replaces M by $M/2$, retaining an exact edge list and good-NO volume at most $\alpha M/2$. There are $O(\log(2n))$ epochs.

2. Dense and final epochs. While $M \geq M_1 := n^{7/4}$, use Lemma 4.1. Its field exponent is $1/4 + 3c = 11/20 < p$, and its outside exponent is $2/3 - c/6 = 13/20 = p$. It supplies the required halving invariant. When $M \leq M_* := C_* T^2/\alpha$, the good-NO volume is at most $C_* T^2$. Use Lemma 2.4 with $U = V$ and that declared NO bound. Its cap is $\widetilde{O}(\sqrt{n} + T)$; every YES marking is in its domain, with no volume assumption. Constants C_* only shift boundaries by $O(1/\log n)$ at exponent level and are absorbed in the prescribed slack.

3. Exact intermediate partition. For $M_* < M < M_1$, clone at $D = \lceil M/n \rceil$, orient every edge once, and bucket as in Section 2. Apply Lemma 2.8 in each bucket. There are $O(\log(2n))$ orientation buckets, and each contributes $O(\log^3(2n))$ pieces. Choose a predetermined $\mathcal{J} = C_{\mathcal{J}} \log^4(2n)$, rounded up, that bounds their number in every epoch. Over all epochs there are $O(\log^5(2n))$ parent piece calls. Each piece has actual almost-biregular degrees; no density discarded during trimming is omitted from the exact edge partition. Use Lemma 5.3 for unbalanced pieces.

For a balanced piece, pad by known zeros to $N = \Theta(n)$ labels. Write a, b for its unpadded side sizes and $g_0 = 8h^2$ for the degree-loss factor of Lemma 2.8. From $ad_L/g_0 \leq e \leq ad_L$ and $bd_R/g_0 \leq e \leq bd_R$, balancedness gives $d_L/d_R, d_R/d_L \leq g_0 \log^B(2n)$. Thus $d = \max(d_L, d_R)$ is a frozen upper bound, and every real coordinate has frozen degree at least d/g with $g = g_0^2 \log^B(2n)$. The piece has $O(nd)$ edges. Dyadic caps are at most twice the clone bound, so $d \leq 2\lceil M/n \rceil$. The global good-NO allowances $L = \alpha M$ and $\Sigma_2 = C_p n^{3-2p} \log(2n)$ therefore give feasible frozen constraints for the true marking. Preserve these constraints during all subsequent deletions. Lemma 3.4 gives (3.9). Use the explicit count cap $K = \min\{N, C_K N/q\}$ on either side. If $K < 1$, a promised NO input has no marked source, and Lemma 2.4 with declared volume zero settles the piece in $\widetilde{O}(\sqrt{N})$, with every YES input still in its domain. Otherwise $1 \leq K \leq N$, as required by the cleanup call. The exponent domain is (5.2), up to the already allocated logarithmic slack.

4. Low degree and raw-load endpoint. If $d \leq T^2/n$, solve the piece by complete edge search, using the final assertion of Lemma 5.1. Otherwise set β and τ by (5.4) and (5.5), and compute the raw load $\theta_0 = N/(T\sqrt{q})$. The next two branches each make one cleanup call, using the indicated threshold and count cap K. If $\theta_0 \leq 1$, apply Lemma 2.5 with $\theta = 1$. Then $\sqrt{NK} \leq O(N/\sqrt{q}) \leq O(T)$, so cleanup costs $\widetilde{O}(T)$. On a good NO history the remaining source-cover volume is $\mathbf{1}^\top B x_s \leq N$. The entire source side is a known vertex cover, and Lemma 2.4 solves the piece in $\widetilde{O}(\sqrt{N})$. No target volume bound is needed. Cleanup searches the deleted neighbor union before deleting it, so its preservation is all-sign.

5. High-degree, nonterminal-load module and field step. In the remaining branch $\theta_0 > 1$, use a rational $\theta \in [\theta_0, 2\theta_0]$, where $\theta_0 = N/(T\sqrt{q})$. Apply Lemma 2.5 with this threshold. Its cleanup cap is $\widetilde{O}(\sqrt{N} + N/(\sqrt{q}\theta)) = \widetilde{O}(T)$. The cleanup proves $Bx_s \leq \theta \mathbf{1}$ on good NO histories. Keep the source marked-volume cap $L = \alpha M$ and impose the known response constraint $Ba_s \leq \theta \mathbf{1}$. The true NO marking remains feasible.

Choose H as in (5.6) and $t = H/\theta$. The positive threshold margin in Lemma 5.1 ensures $t \geq 1$ beyond the fixed base range. Apply extraction with the piece's frozen source upper bound $D_{\text{piece}} = d$. If $t \leq d$, use Theorem 1.2 and completely test all removed modules. Their cap is $\widetilde{O}(T)$ by Corollary 5.2. If $t > d$, remove nothing and use the last branch of Proposition 3.2. Edge deletion only decreases loads and marked volumes; the original frozen constraints remain feasible.

On the residual, Propositions 3.2 and 3.3 prove the offdiagonal variance and omitted-diagonal bias. Corollary 5.2 pays both the main and secondary clocks of Lemma 4.2. Invoke that lemma, with its source-only entropy but all-coordinate cutoff objective, and its charged fresh edge/count/volume validations.

6. Outside edges and halving. For a response set P, let H_{out} contain exactly the current edges with at least one endpoint in $U = V_{\text{piece}} \setminus P$. This known U is a vertex cover of H_{out}. Test H_{out} before deleting its edges. On a successfully validated output,

$$\sum_{u \in U} x_u \leq C_O \beta N, \qquad \sum_{u \in U} x_u \deg_{H_{\text{out}}}(u) \leq d \sum_{u \in U} x_u \leq C_O \beta N d.$$

The latter is an algebraic consequence of the frozen upper degree bound, not a separate volume assumption. If $\beta = \alpha$, use Lemma 4.3 with declared count cap $\min\{N, C_O \alpha N\}$, costing $\widetilde{O}(n^{2/3}\alpha^{1/6} + \sqrt{n}) = \widetilde{O}(T)$. If $\beta > \alpha$, (5.4) gives $\beta n d \leq T^2$: either $\beta = T^2/(nd)$ or $\beta = \beta_0 \leq T^2/(nd)$. Lemma 2.4, with the declared volume $C_O \beta N d$, then costs $\widetilde{O}(T)$. Both detectors have an all-YES domain.

Finally take $\mathcal{J}$ to be a predetermined upper bound on the number of pieces in this epoch, and use (5.5) on each balanced piece i. Its deterministic retained edge cap and validated NO volume cap satisfy

$$e(P_i) \leq \frac{C_E \kappa}{\mathcal{J}} \frac{\alpha}{\beta_i} M \leq \frac{M}{4\mathcal{J}}, \qquad W(P_i, x) \leq \frac{C_W \kappa}{\mathcal{J}} \alpha M \leq \frac{\alpha M}{4\mathcal{J}}.$$

Solved low-degree, terminal, or unbalanced pieces retain no edges. Sum the at most $\mathcal{J}$ retained edge lists, which are edge-disjoint. By Lemma 2.7 their total budgets are at most $M/4$ and $\alpha M/4$, hence in particular $M/2$ and $\alpha M/2$. This gives the required simultaneous allocation. Restore every retained edge through its original edge identifier and merge these exact lists on the original physical labels; no induced closure is taken. Marked volume is additive over these disjoint edge lists. Although identifying clones can increase their second moment, the restored graph is an edge subset of the fixed original graph returned by preprocessing. Its second moment is therefore still at most the original allowance Σ_2. The next epoch starts on at most n original labels and constructs fresh clones and frozen constraints, so the number of labels does not accumulate over epochs.

7. Correctness on both input signs. On a good NO history, cleanup promises, frozen feasible sets, the variance bounds, and response validations all apply, and all retained invariants follow. Conditional first-failure accounting limits false YES or an invalid invariant to the assigned total error. On a YES input, track an original marked edge. It is in a completely tested low-degree or recursive piece, in an all-YES cover terminal, in a searched cleanup neighbor union, in a completely tested removed module, in a guarded outside call, or in the exact retained edge list. All-sign preservation and all-YES domain clauses prevent an untested disappearance on a good history. The final cover call detects any retained collision. Off-promise NO children on a global YES input need not have NO correctness: a YES answer is already globally correct, and an actual YES child has completeness.

All sample caps, optimizer attempts, and abort conventions are those specified by the implementation proofs. A cap exhaustion may report YES and is charged only on NO inputs. Every branch has its fixed query cap. Choose the sum of all NO and preservation failure budgets below $1/3$, using a conservative fraction for each class of calls.

8. Closing strong induction. The explicit balanced work, dense and terminal steps, and global preprocessing have total cap $C_B n^p \log^{K_0}(2n)$. The unbalanced recursive calls are strictly smaller and their total cap is at most $(A/4)n^p \log^K(2n)$ by Lemma 5.3. Choose the aspect cutoff first, then $K \geq K_0$, then A large enough to absorb C_B and the finite initial range. The contraction uses $p - 1/2 = 3/20 > 0$. This proves the unconditional cap (1.2) with bounded error and all-input fixed clocks. □

A Auxiliary proofs

A.1 Preprocessing proofs

These proofs use only the published primitives stated in Section 2, independently of Lemmas 4.1–4.3.

A.1.1 Source-neighborhood cleanup

Proof of Lemma 2.5. If $K < 1$, return the input graph: promised NO inputs have no marked source, and this deletion-free operation is safe on all inputs. Otherwise use amplitude estimation [12, Theorem 12] with an even integer precision $M_{\mathrm{AE}} \geq C\sqrt{N}$ of order $\sqrt{N}$, and take the median of a predetermined $O(\log(2/\varepsilon))$ number of repetitions. The predicate is the known source mask AND the marking bit. For its masked count $k = \sum_{v \in S} x_v$, the amplitude-estimation guarantee gives

$$|N\tilde{a} - k| \leq \frac{2\pi\sqrt{k(N-k)}}{M_{\mathrm{AE}}} + \frac{\pi^2 N}{M_{\mathrm{AE}}^2} \leq 0.1\sqrt{k} + 0.01$$

after choosing C large. Confidence amplification makes the probability of violating this inequality at most $\varepsilon/8$. When $k = 0$ every repetition returns zero exactly. Nearest-integer rounding gives $\widehat{k} = 0$ when $k = 0$, and $k/3 \leq \widehat{k} \leq 2k$ when $k \geq 1$, on a good transcript. Clip to $[0, \min(N, 2\lceil K \rceil)]$; on promised transcripts this keeps the comparison, and always $\widehat{k} \leq 4K$. If $\widehat{k} = 0$, return the graph.

Otherwise request

$$s = \left\lceil (12\widehat{k}/\theta) \log(8N/\varepsilon) \right\rceil$$

verified uniform marked-source samples, using fixed-budget search trials [13] and the estimate $\widehat{k}$. For completeness, each trial chooses with equal probability between querying a uniform label and applying j Grover iterations, with j uniform in $\{0, \ldots, m_{\mathrm{tr}} - 1\}$ and $m_{\mathrm{tr}} = \left\lceil 2\sqrt{N/\widehat{k}} \right\rceil$; its measured candidate is always verified. Each trial has cap $O(1 + \sqrt{N/\widehat{k}})$. For $1 \leq k \leq N/2$, put $\phi = \arcsin\sqrt{k/N}$. The Grover branch has average success probability

$$\frac{1}{m_{\mathrm{tr}}} \sum_{j=0}^{m_{\mathrm{tr}}-1} \sin^2((2j+1)\phi) = \frac{1}{2} - \frac{\sin(4m_{\mathrm{tr}}\phi)}{4m_{\mathrm{tr}}\sin(2\phi)} \geq \frac{3}{8},$$

since $\widehat{k} \leq 2k$ implies $m_{\mathrm{tr}} \sin(2\phi) \geq 2$. For $k > N/2$ the uniform-query branch alone gives total success probability greater than $1/4$. Thus a good estimate guarantees success probability at least $3/16$. Allow only $C'(s + \log(8/\varepsilon))$ trials; if too few succeed, return YES. A Chernoff bound pays this failure on promised NO inputs. A fresh trial is symmetric among marked labels, and $\Pr(\text{success and output } v) = p_{\mathrm{tr}}/k$ for each marked source label v. Equivalently, the success indicator and a uniform marked label can be coupled independently. Hence, conditional on enough successes before the cap, their labels are independent uniform samples, conditional on the fixed count estimate. Future search trials use fresh randomness. Take the trial-cap constant large enough that its

exhaustion has probability at most $\varepsilon/8$ on promised NO inputs. This is a fixed clock, not an expected-time sampler.

Let Z be the sampled marked vertices and form the known union $U = N_G(Z)$. Search U for a mark with cap $O(\sqrt{N}\log(8/\varepsilon))$, verifying any answer. A found label and a sampled neighbor certify a collision. Otherwise delete U. On a NO marking U is unmarked. On any marking, deleting a marked vertex requires a failed union search, whose conditional probability is charged regardless of whether the count estimate was good.

For a vertex having $k_u > \theta$ marked neighbors in S, its chance of surviving a good capped sample history is at most

$$(1 - k_u/k)^s \le e^{-sk_u/k} \le (\varepsilon/(8N))^4.$$

Union bound over all vertices. If $k \le \theta$, the conclusion holds without sampling. The trial cap costs

$$O\left(\frac{\sqrt{N\widehat{k}}}{\theta}\log(8N/\varepsilon) + \sqrt{N}\log(8/\varepsilon)\right),$$

which gives (2.4) on every transcript, not only promised ones. The count and trial-cap failures each cost at most $\varepsilon/8$. Give the verified union search failure budget $\varepsilon/8$. The displayed survival bound, summed over at most N vertices, is less than $\varepsilon/8$. Thus the overall promised-NO failure is below ε, and all-sign preservation needs only the union-search budget. □

A.1.2 Global degree moments

Proof of Theorem 2.6. Apply Lemma 2.5 with all labels as source, $K = n$, and $\theta_0 = n^{1-p}$. Every vertex then has at most θ_0 marked neighbors on a good NO history. Double counting gives $W \le B := n\theta_0 = n^{2-p}$.

Fix a dyadic degree threshold d. At every subsequent invocation recompute $S_d = \{v : \deg_{\text{current}}(v) \ge d\}$. If its declared marked-volume cap is L_j, its marked count is at most $K_j = \min(n, L_j/d)$. If $K_j < 1$, there is no marked source on the good promised history. Otherwise clean up this source at $\theta_j = \sqrt{nK_j}/n^p$. Each call costs $\widetilde{O}(n^p)$. Its new source volume is at most $n\theta_j$. Vertices still of degree at least d belonged to the old source, since degrees only decrease. Thus the new declared tail cap may be

$$L_{j+1} \le B\sqrt{L_j/(nd)}. \tag{A.1}$$

The fixed point is $L_* = B^2/(nd) = n^{3-2p}/d$. For $L_j \ge L_*$ the ratio satisfies $L_{j+1}/L_* \le \sqrt{L_j/L_*}$. The initial declared allowance is B, and $d \le n$, so its ratio to L_* is at most n^p. For example, $r_n = \lceil \log_2 \max\{1, p\log_2 n\}\rceil + 1$ rounds suffice: $n^{p/2^{r_n}} \le \sqrt{2} < 2$. If the declared allowance is already below $2L_*$, skip the rounds. If a declared $K_j < 1$, retain the current graph and record zero tail allowance; this is valid on good NO histories and deletes nothing on either sign. In particular small thresholds $d \le \theta_0$ are already paid by $W \le B$.

Do this for the $O(\log(2n))$ dyadic thresholds up to n. Subsequent deletions cannot spoil an earlier tail cap. A dyadic threshold immediately below a real d proves (2.6) with a factor two. Summing shells gives

$$S_2 \le \sum_{d \text{ dyadic}} 2d \sum_{v : d \le \deg_H(v) < 2d} x_v \deg_H(v) \le C_p n^{3-2p}\log(2n).$$

The total number of calls is fixed and polylogarithmic. Use the adaptive error convention in Section 2 for their promises and error budgets. The all-sign preservation statement follows from the all-sign search safety in Lemma 2.5. □

A.1.3 Moment inheritance

Proof of Lemma 2.7. If degrees split as $d = \sum_j d_j$, then $\sum_j d_j^2 \le (\sum_j d_j)^2 = d^2$ and $\sum_j x_v d_j = x_v d$. Sum at each original label. These same identities apply to degrees in an exact edge partition. Monotonicity under deletion is immediate for the nonnegative degrees and bits. □

A.1.4 Exact almost-biregular partition

Proof of Lemma 2.8. In a current residual with m edges, bucket each side by the dyadic upper cap of its actual positive degree. Let $q \le h^2$ be the product of the numbers of occupied buckets. A largest bucket pair (A, B) has at least m/q edges. Its caps satisfy $|A| \le 2m/d_L$ and $|B| \le 2m/d_R$. On a separate copy of this pair repeatedly remove vertices of candidate degree below $d_L/(8q)$ or $d_R/(8q)$. Charge every removed edge to its first deleted endpoint. At most

$$|A|d_L/(8q) + |B|d_R/(8q) \le m/(2q)$$

edges are lost. The surviving piece has at least $m/(2q) > 0$ edges, and its final actual degrees meet the asserted bounds. Rational thresholds below one simply remove isolates.

Record only this surviving edge list and delete only those edges from the global residual. Edges trimmed from the candidate remain available for later extractions. Recompute all degrees and repeat. After r extractions the residual has at most $m_0 e^{-r/(2h^2)}$ edges. Once this is below one the integer residual is empty. Edge disjointness and the degree-sum identity prove all remaining claims. □

A.2 Composition and terminal proofs

A.2.1 OR of unequal clocks

Proof of Lemma 2.3. Group the caps into dyadic bins. A bin containing r_k routines of cap at most 2^{k+1} costs $\widetilde{O}(2^k\sqrt{r_k})$ by bounded-error quantum search [14]. One may instead amplify complete children to inverse-polynomial error and use a compute–phase–uncompute simulation with a hybrid error bound. Select the algorithm index coherently; each query is routed to the original physical address. Disjoint variable sets are not required. Evaluate the logarithmically many bins sequentially, allocating their errors in advance. Cauchy–Schwarz bounds the sum of bin costs by a logarithmic factor times the square root in (2.2). Padding clocks is allowed; query histories are not reused as a free data structure.

Intermediate measurements and classical feed-forward in a child can be deferred and implemented reversibly, retaining their work registers. For fixed known instance data the constructions have finitely many measurement transcripts, and every known-data computation terminates on each transcript. They therefore admit a finite coherent circuit, with shorter query branches padded by identity queries. This can be computationally expensive but does not increase the query cap. In particular the recursive composition does not assume free access to the outcome of a measured child execution. □

A.2.2 Known-cover volume terminal

Proof of Lemma 2.4. If H has no edges, return NO without queries. Otherwise U is nonempty. Use the program in [4, Section 3.2], on the space with orthonormal basis $\{e_0\} \cup \{e_u : u \in U\}$ and target e_0. For each $u \in U$, include a column $e_0 - e_u$, available when $x_u = 1$. For each ordered incidence (u, v) with $u \in U$ and $uv \in E(H)$, include a separate column e_u, available when $x_v = 1$.

A marked edge has an endpoint $u \in U$; its two available columns sum to e_0, giving positive witness size at most 2. On a NO marking, put $w = e_0 + \sum_{u \in U} x_u e_u$. Then $e_0^\top w = 1$, every available $e_0 - e_u$ annihilates w, and every available incidence column e_u also annihilates w, because $x_v = 1$ implies $x_u = 0$. Its negative witness size, the sum of squared inner products with *all* columns, is exactly

$$|U| - \sum_{u \in U} x_u + \sum_{u \in U} x_u \deg_H(u) \le |U| + L.$$

Apply span-program conversion on the partial Boolean domain

$$\{x : \mathrm{GC}_H(x) = 1\} \ \cup \ \{x : \mathrm{GC}_H(x) = 0, \ \sum_{u \in U} x_u \deg_H(u) \le L\}.$$

The positive and negative witness bounds give query complexity $O(\sqrt{2(|U| + L)})$ on this domain [3, Theorem 4][15, 16]. Balance the two witness bounds by rescaling all columns by $\lambda = (2/(|U|+L))^{1/4}$: the positive bound becomes $2/\lambda^2$ and the negative bound becomes $\lambda^2(|U| + L)$, both $\sqrt{2(|U| + L)}$. The conversion uses a prescribed clock on every oracle input; restricting its correctness domain does not give an expected-time or promise-dependent stopping rule. Amplification gives (2.3). All repeated incidence columns have been counted, and no query discovers the promised L. □

A.2.3 Aspect-ratio induction

Proof of Lemma 5.3. Split the larger side into $r = \lceil b/a \rceil$ blocks of size at most a. Each block with the smaller side is a complete child instance on at most $2a$ labels, and $2a \le 2C_N n \log^{-B}(2n) < n$ beyond a fixed base range. Amplify each child to error at most $\left(\varepsilon/(C_{\mathrm{amp}} \mathcal{R} r^2 \log^2(2\mathcal{R}/\varepsilon))\right)^2$, where $\mathcal{R} \ge 1$ is the known maximum number of parent piece calls. Choose the absolute constant C_{amp} large enough for the compute–phase–uncompute hybrid bound. This uses only $O(\log(2n/\varepsilon))$ repetitions. Indeed, a compute–phase–uncompute call with child error δ has clean-input norm error $O(\sqrt{\delta})$. A hybrid over $O(\sqrt{r}\log(2\mathcal{R}/\varepsilon))$ evaluations therefore contributes at most $\varepsilon/(10\mathcal{R})$ for the displayed choice, after increasing C_{amp}. Allocate the ideal OR's amplified error the same amount. These children's edge lists partition the piece. Quantum OR of their complete inductive algorithms costs, after amplification,

$$CA\sqrt{b/a}\,(2a)^p \log^{K+2}(2n) \le C' A b^p (a/b)^{p-1/2} \log^{K+2}(2n).$$

Shared smaller-side bits are handled by the known address map, not disjoint-variable composition. The two extra logarithmic factors pay child error reduction and the OR's confidence amplification; the parent-call error budgets have already been included. Summing over the parent piece calls gives total cap

$$C'' A n^p \log^{K+k+2-B(p-1/2)}(2n),$$

which is at most the claimed quarter for sufficiently large n. Choose B first. Once the explicit balanced construction has a finite logarithmic overhead $K_0(B)$, choose $K \ge K_0(B)$ and then A to absorb that construction and the finite base range. Every child is strictly smaller, so this closes strong induction without applying the desired theorem to a same-size copied graph. □

B Bounded-cover-count detection

This section supplies the promised-query algorithm used for edges outside a selected vertex set. The construction is a bipartite specialization of the value-dependent learning-graph method of Belovs [2, Section 4]. We give its matrices explicitly, so that the marked-count promise, setup normalization, and absence of a logarithmic conversion loss are part of the proof.

B.1 The imported adversary characterization

Let $f : \mathcal{D} \to \{0,1\}$, where $\mathcal{D} \subseteq \{0,1\}^m$ is a finite promised domain. The general adversary bound has the positive-semidefinite dual formulation

$$\mathrm{Adv}^{\pm}(f) = \min_{X_1,\dots,X_m \succeq 0} \max_{x \in \mathcal{D}} \sum_{j=1}^{m} X_j[x,x], \qquad \sum_{j : x_j \neq y_j} X_j[x,y] = 1 \quad \text{whenever } f(x) \neq f(y). \tag{B.1}$$

All matrices are indexed by $\mathcal{D}$. We import this formulation and the constant-error query characterization

$$Q_\varepsilon(f) = \Theta_\varepsilon(\mathrm{Adv}^{\pm}(f)), \qquad 0 < \varepsilon < 1/2 \text{ fixed}, \tag{B.2}$$

from the general adversary theory; see [15, 16] and the formulation in [2, Definition 1 and Theorem 2]. In particular, a feasible family in (B.1) with objective at most T gives a fixed-cap constant-error algorithm with $O(T)$ queries, without an additional $\log m$ factor. Only this upper-bound direction is used below.

B.2 The precise bipartite promise

Let $H = (L, R, E)$ be a known bipartite graph on disjoint label sets L, R, with $|L| = a$ and $|R| = b$. For an integer $0 \le K \le a$, write

$$\mathcal{D}_K = \left\{ z \in \{0,1\}^{L \cup R} : \sum_{u \in L} z_u \le K \right\}.$$

There is no restriction on the right marking bits. The promise on the left applies to *both* output signs.

Proposition B.1 (Bounded-left-mark primitive)**.** *Suppose $a \ge 2$, $b \ge 1$, and $1 \le K \le a$. For every integer $1 \le r \le \lfloor a/2 \rfloor$,*

$$Q_{1/3}(\mathrm{GC}_H|_{\mathcal{D}_K}) = O\left(r\sqrt{K/a} + \sqrt{a} + \sqrt{\frac{ab}{r+1}} \right). \tag{B.3}$$

The implicit constant is independent of H, a, b, K, r.

Proof. If the function is constant, zero queries suffice. Otherwise choose, for each positive $x \in \mathcal{D}_K$, one marked edge $(u_x, v_x) \in E$, with $u_x \in L$ and $v_x \in R$. These choices index a single family of matrices over the entire known promised domain; they are not operations performed on the unknown input by the query algorithm.

Put

$$B = \binom{a-1}{r}, \qquad p = B^{-1}.$$

For each r-subset $S \subseteq L$, fix an order $s_1(S), \dots, s_r(S)$ independent of the input. Define three disjoint families of formal query arcs:

1. A setup arc (S, ℓ) for every $|S| = r$ and $1 \le \ell \le r$. It queries $s_\ell(S)$ and has preceding loaded set $\{s_1(S), \dots, s_{\ell-1}(S)\}$. The full tag S is retained, even when two setup prefixes happen to agree.

2. A left-witness arc (S, j) for every $|S| = r$ and $j \in L \setminus S$. It queries j after the loaded set S.

3. A right-witness arc (T, j) for every $T \subseteq L$ with $|T| = r + 1$ and $j \in R$. It queries j after the loaded set T.

The family tag distinguishes arcs belonging to different stages. For an arc e, write $j(e)$ for its queried coordinate and P_e for its preceding loaded set.

For a positive input x, use the following B paths, each with weight p: choose $S \subseteq L \setminus \{u_x\}$, $|S| = r$, load S along its tagged setup path, then query u_x, and finally query v_x. Let $\chi_e(x)$ indicate that x uses arc e. No arc is repeated on two different paths for a fixed positive x. This is immediate for setup and left-witness arcs. For a final arc, its preceding set $T = S \cup \{u_x\}$ determines $S = T \setminus \{u_x\}$ uniquely. Thus $\chi_e(x)$ is a 0–1 indicator, without a hidden path-multiplicity factor.

We next specify the Gram matrices. For a setup arc choose positive weights $w_{e,0}, w_{e,1}$. For each assignment $\sigma \in \{0,1\}^{P_e}$, define real vectors indexed by $z \in \mathcal{D}_K$ as follows:

$$\psi_{e,\sigma}(z) = \begin{cases} w_{e,1}^{-1/2}, & f(z) = 1,\ \chi_e(z) = 1,\ z|_{P_e} = \sigma,\ z_{j(e)} = 1, \\ w_{e,1}^{1/2}, & f(z) = 0,\ z|_{P_e} = \sigma,\ z_{j(e)} = 0, \\ 0, & \text{otherwise}, \end{cases}$$

$$\phi_{e,\sigma}(z) = \begin{cases} w_{e,0}^{-1/2}, & f(z) = 1,\ \chi_e(z) = 1,\ z|_{P_e} = \sigma,\ z_{j(e)} = 0, \\ w_{e,0}^{1/2}, & f(z) = 0,\ z|_{P_e} = \sigma,\ z_{j(e)} = 1, \\ 0, & \text{otherwise}. \end{cases}$$

Here $f = \mathrm{GC}_H|_{\mathcal{D}_K}$. On either witness stage every used positive arc queries a 1. For those arcs retain $\psi_{e,\sigma}$ with a positive weight $w_{e,1}$ and set $\phi_{e,\sigma}$ identically to zero. In particular no zero weight is inserted into a denominator. For each queried label j, set

$$X_j = p \sum_{e:j(e)=j} \sum_{\sigma \in \{0,1\}^{P_e}} \left(\psi_{e,\sigma} \psi_{e,\sigma}^T + \phi_{e,\sigma} \phi_{e,\sigma}^T \right). \tag{B.4}$$

Every X_j is positive semidefinite. The overall factor is p, not p^2.

Cross-output constraints. Fix a positive x and a negative y in $\mathcal{D}_K$. For an arc used by x, its contribution to $X_{j(e)}[x, y]$ when $x_{j(e)} \neq y_{j(e)}$ is p if $x|_{P_e} = y|_{P_e}$, and zero otherwise. This follows directly from the reciprocal weights in the Gram vectors. On each of the B paths for x, at least one queried bit differs between x and y: otherwise $y_{u_x} = y_{v_x} = 1$, contradicting that y is negative. The first differing bit contributes exactly p to the filtered sum in (B.1). Earlier bits are not in that sum, and after this first disagreement the preceding assignments no longer agree, so all later contributions are zero. Because no used arc belongs to two paths for this fixed x, summing gives

$$\sum_{j:x_j \neq y_j} X_j[x, y] = Bp = 1.$$

Symmetry gives the same constraint with the signs reversed. There is no required equality for a pair with equal outputs. Hence (B.4) is feasible for the dual.

Setup diagonal budget. Use the same setup weights on every setup arc:

$$w_{e,0} = \sqrt{a/K}, \qquad w_{e,1} = \sqrt{K/a}. \tag{B.5}$$

For clarity, the four setup charges from one matching assignment are

	$z_{j(e)}=0$	$z_{j(e)}=1$
$f(z)=1$ and $\chi_e(z)=1$	$p/w_{e,0}$	$p/w_{e,1}$
$f(z)=0$	$pw_{e,1}$	$pw_{e,0}$

At any fixed input and arc there is exactly one matching preceding assignment σ; summing over σ creates no additional diagonal factor.

If a positive x has $k_x \le K$ marked left labels, its witness u_x is one of them. A uniform setup set $S \subseteq L \setminus \{u_x\}$ has expected marked count $r(k_x-1)/(a-1)$. The positive setup diagonal is consequently

$$\left(r - \frac{r(k_x-1)}{a-1}\right)\frac{1}{w_{e,0}} + \frac{r(k_x-1)}{a-1}\frac{1}{w_{e,1}} \le r\sqrt{K/a} + \frac{rK}{a}\sqrt{a/K} = 2r\sqrt{K/a}.$$

We used $(k_x-1)/(a-1) \le K/a$, which follows from $1 \le k_x \le K \le a$.

For a negative y with $k_y \le K$ marked left labels, all $\binom{a}{r}$ setup sets occur in the negative diagonal, regardless of the chosen witnesses for positive inputs. Each left coordinate occurs in $\binom{a-1}{r-1}$ such sets. Thus the negative setup diagonal is exactly

$$\frac{\binom{a}{r}}{\binom{a-1}{r}} r\left[(1-k_y/a)w_{e,1} + (k_y/a)w_{e,0}\right] \le \frac{a}{a-r} r\left(\sqrt{K/a} + \frac{K}{a}\sqrt{a/K}\right) \le 4r\sqrt{K/a},$$

because $r \le a/2$. The factor $a/(a-r)$ is retained explicitly; no assumption $r = o(a)$ is needed.

Left-witness diagonal budget. Every positive input uses B left-witness arcs, so its total diagonal at this stage is $1/w_{e,1}$. There are $\binom{a}{r}(a-r)$ arcs in this stage. A negative input contributes at most $pw_{e,1}$ per arc, giving total at most

$$p\binom{a}{r}(a-r)w_{e,1} = aw_{e,1}.$$

Choose $w_{e,1} = a^{-1/2}$. Both signs then have stage budget at most $\sqrt{a}$.

Right-witness diagonal budget. Again a positive input uses B arcs, with diagonal $1/w_{e,1}$. There are $\binom{a}{r+1}b$ possible arcs, so the negative diagonal is at most

$$p\binom{a}{r+1}b\, w_{e,1} = \frac{ab}{r+1}\, w_{e,1}.$$

Taking $w_{e,1} = \sqrt{(r+1)/(ab)}$ bounds either sign by $\sqrt{ab/(r+1)}$.

Combining the three stages, the objective of (B.4) is at most

$$4r\sqrt{K/a} + \sqrt{a} + \sqrt{ab/(r+1)}.$$

The imported conversion (B.2) proves (B.3). No intermediate stage is an error-prone algorithm, so there is no stage-by-stage amplification or associated logarithmic cost. □

B.3 Padding, endpoint cases, copies, and confidence

Corollary B.2 (Uniform bounded-mark interface)**.** *Let H be a known bipartite graph with at most N labels per side, and let $K \geq 0$ be a known marking cap. On the domain where at most K left labels are marked, on both signs, graph collision has a fixed-cap algorithm of query complexity*

$$O\Big(\sqrt{N}\,(K+1)^{1/6}\Big) \tag{B.6}$$

for constant error. The bound remains valid for known repeated labels and known-zero dummy labels. Error at most δ is obtained with an additional factor $O(\log(2/\delta))$.

Proof. An empty side or an empty edge set gives the constant-zero function. Replace K by $\min\{\lfloor K \rfloor, |L|\}$, which leaves its Boolean promise unchanged and can only decrease the claimed bound. If this cap is zero, the promised function is identically zero and needs no queries. Known-zero labels pad both sides to size N without increasing the number of marked labels. Cases $N \leq 1$ are direct constant-size computations. For $N \geq 2$ and $1 \leq K \leq N$, take

$$r = \max\left\{1, \left\lfloor \frac{N}{2(K+1)^{1/3}} \right\rfloor\right\}.$$

Then $1 \leq r \leq \lfloor N/2 \rfloor$, $r \leq 1 + N/(2(K+1)^{1/3})$, and $r + 1 \geq N/(2(K+1)^{1/3})$. Consequently

$$r\sqrt{K/N} \leq 1 + \tfrac{1}{2}\sqrt{N}\,(K+1)^{1/6},$$
$$\sqrt{N^2/(r+1)} \leq \sqrt{2N}\,(K+1)^{1/6}.$$

Proposition B.1, applied to the padded graph, proves (B.6).

For repeated labels, first construct the algorithm with independent label bits and then restrict to assignments in which copies of an original variable agree. A copied-label query is simulated by computing its known original address, querying that bit, and uncomputing the address. In a coherent superposition containing known-zero dummy labels, a controlled swap can route the bit-query target to an auxiliary $|+\rangle$ register on dummy branches. The oracle acts trivially there because $X|+\rangle = |+\rangle$. Thus each padded query slot uses at most one original bit query; no controlled unknown unitary is assumed for free.

Finally, constant-error decision amplification by a fixed number of repetitions gives the stated $O(\log(2/\delta))$ factor. The constructed algorithm has a fixed query cap even when physically run outside its promised domain; only its correctness guarantee is then absent. Conditional on a prior classical transcript fixing the known graph and the cap K, the same pointwise guarantee applies whenever the actual input satisfies the promise. No union bound over all possible transcripts is required for this later call. □

B.4 A count guard and the universal cover theorem

The preceding primitive assumes a left-mark cap on both input signs. A gap-counting guard converts it to the all-YES domain required in the main algorithm.

Theorem B.3 (Universal bounded-cover-count detector)**.** *Let H be a known simple edge-list graph on at most N labels, let U be a known vertex cover, and let $0 \leq K \leq N$ be a declared real bound. There is a finite quantum detector with all-input cap*

$$O\Big(\sqrt{N}\,(K+1)^{1/6}\log(2/\varepsilon)\Big)$$

and error at most ε on every YES marking and every NO marking satisfying $\sum_{u\in U} x_u \leq K$. The guarantee is uniform under known copied-label restrictions.

Proof. Empty graphs return NO without queries, so assume $N \geq 1$. Write $t = \sum_{u \in U} x_u$ and $k = \lfloor K \rfloor$. First represent H as a bipartite graph B with left side a separate copy of U and right side a separate copy of $V(H)$. Include $(u_{\mathrm{L}}, v_{\mathrm{R}})$ for each $u \in U$ and each original edge $\{u, v\}$. Copy the original marking to both sides. Every edge has an endpoint in U, so collision in B is equivalent to collision in H. An edge internal to U may give two bipartite edges, which does not affect this equivalence. Each side has at most N labels, and the number of marked left labels is exactly t.

If $k = 0$, search U for any marked label, using a fixed-cap bounded-error OR algorithm with error ε. This costs $O(\sqrt{N}\log(2/\varepsilon))$ queries. Every YES marking has a marked endpoint in U, whereas every promised NO marking has $t = 0$. Thus the test has the required partial-domain correctness even though it need not decide collision on the excluded NO markings.

Assume now $k \geq 1$. If $k \geq |U|$, apply Corollary B.2 to B with cap $|U|$. Otherwise count marked labels in U, padding the counting universe to N by known-zero labels. The amplitude-estimation theorem of [12, Theorem 12], specialized to counting, gives an estimate $\widehat{t}$ with

$$|\widehat{t} - t| \leq \frac{2\pi\sqrt{t(N-t)}}{M} + \frac{\pi^2 N}{M^2}$$

with probability at least $8/\pi^2$ using $O(M)$ queries. Choose in advance

$$M = \left\lceil 32\pi\sqrt{N/k} \right\rceil.$$

The displayed error is at most

$$\frac{\sqrt{tk}}{16} + \frac{k}{1024} \leq \frac{t+k}{32} + \frac{k}{1024} \leq \frac{t+k}{16}. \tag{B.7}$$

A median of a predetermined odd number $O(\log(2/\varepsilon))$ of independent estimates satisfies (B.7) except with probability $\varepsilon/2$. Its query cap is fixed on every input.

If $\widehat{t} > 2k$, output YES. Otherwise invoke Corollary B.2 on B with cap $k' = \min\{|U|, 4k\}$ and error $\varepsilon/2$. On the good estimation event, two implications justify this rule. First, if $t \leq k$ then

$$\widehat{t} \leq t + \frac{t+k}{16} \leq \frac{9}{8}k < 2k,$$

so the guard does not accept a promised NO input. Second, if the guard does not accept, then

$$\frac{15t - k}{16} \leq \widehat{t} \leq 2k, \qquad \text{hence} \qquad t \leq \frac{33}{15}k < 4k.$$

Together with $t \leq |U|$, this proves the cap $t \leq k'$ on *both* input signs before the primitive is invoked.

For a promised NO marking the good counting event therefore leads to an on-promise primitive call, which returns NO except with probability $\varepsilon/2$. For an arbitrary YES marking, acceptance by the guard is correct; otherwise the same good event places the primitive call on its promise domain, where it returns YES except with probability $\varepsilon/2$. In either case the total error is at most ε. No correctness assertion is needed for NO markings with $t > K$.

The count guard costs $O(\sqrt{N/k}\log(2/\varepsilon))$, and the primitive costs $O(\sqrt{N}\,(k'+1)^{1/6}\log(2/\varepsilon))$. Since $k \geq 1$, $k' \leq 4k$, and $k \leq K$, their sum is the claimed bound. Both constituent algorithms have fixed all-input caps; off-promise behavior cannot lengthen an execution. The copied-label argument in the preceding corollary applies to both the counting predicate and the bipartite primitive. □

Theorem B.3 establishes the universal bounded-cover-count interface without an additional hypothesis. Appendix C proves the dense-epoch and weighted-response routines. This proof requires no computational certificate.

C Contraction implementations

C.1 Dense contraction

The dense part of the construction uses ordinary sampling of marking bits. In particular it does not require the source-supported stochastic response used for intermediate pieces. We give its proof, including the simultaneous edge and marked-volume contractions.

Lemma C.1 (Cutoff regularizer)**.** *For $0 < b \le 1/8$, define continuously on $[0,1]$*

$$h_b(u) = -u\log u - (1-u)\log(1-u) + \frac{u}{u+b} + \frac{1-u}{1-u+b} - \frac{1}{1+b}.$$

Then $h_b(0) = h_b(1) = 0$, $0 \le h_b < 2$, and, on $(0,1)$, $h_b'' \le -4$. Also $h_b'(0+) = +\infty$, $h_b'(1-) = -\infty$, and $h_b'(3b) \ge 1/(32b)$. Consequently the endpoint tangent loss

$$\Lambda_b(s,u) = h_b(u) + (s-u)h_b'(u), \qquad s \in \{0,1\},$$

is nonnegative and satisfies $\Lambda_b(1,u) \ge 5/(256b)$ when $0 < u \le 3b$.

Proof. Direct differentiation gives

$$h_b'(u) = \log\frac{1-u}{u} + \frac{b}{(u+b)^2} - \frac{b}{(1-u+b)^2},$$
$$h_b''(u) = -\frac{1}{u} - \frac{1}{1-u} - \frac{2b}{(u+b)^3} - \frac{2b}{(1-u+b)^3} \le -4.$$

The endpoint values and slopes follow from these formulas. Concavity and symmetry imply $0 \le h_b(u) \le h_b(1/2) < \log 2 + 1 < 2$. For $b \le 1/8$, the logarithmic term in $h_b'(3b)$ is nonnegative, and $32b^2 \le (1-2b)^2$, so

$$h_b'(3b) \ge \frac{1}{16b} - \frac{b}{(1-2b)^2} \ge \frac{1}{32b}.$$

The tangent-loss conclusions follow from concavity, the zero endpoint values, and $1 - 3b \ge 5/8$. □

Lemma C.2 (Container from a volume allowance)**.** *Let H be a known simple graph on $N \ge 1$ vertices, with adjacency matrix A. Fix known rational $L \ge 0$, $\tau > 0$, $0 < b \le 2^{-40}$, and $0 < \varepsilon < 1/2$. There is a finite routine of all-input query cap*

$$O\Big((1 + L/\tau^2 + b^{-1/2})\log(2/\varepsilon)\Big)$$

that reports YES or returns a known set P. Every returned set has $e_H(P) \le 3\tau N/b^2$. Except with probability ε, on either sign every returned set has at most $2^{16}bN$ marked labels outside it. On every NO input with $W_H(x) \le L$, except with probability ε, the routine returns a set and

$$W_{H[P]}(x) \le 3\tau N/b.$$

Proof. Set $C = (-A - \tau I)_+$ and $B = A + \tau I + C \succeq 0$. Both matrices are determined by the known graph. A uniform label $i \in [N]$ and one bit query produce the entire vector $Y = NCe_ix_i$. Its mean is Cx. Functional calculus gives $C^2 \preceq A^2$, and hence, on the promised NO domain,

$$\mathbb{E}\|Y - \mathbb{E}Y\|_2^2 \le N\sum_{i:x_i=1}(C^2)_{ii} \le N\sum_{i:x_i=1}\deg_H(i) \le NL.$$

Each coordinate of Ce_i is known, so evaluating this vector does not require a separate marking query for each coordinate.

For $L > 0$, average $s = \lceil 1600L/\tau^2 \rceil$ independent samples. The expected squared error is at most NL/s. With $r = \tau\sqrt{N}/10$, Markov's inequality shows that the average lies within distance r of Cx with probability at least $15/16$. Form $b_0 = \lceil 32 \log(4/\varepsilon) \rceil$ independent averages. Hoeffding's inequality implies, with probability at least $1 - \varepsilon/4$, that more than $3b_0/4$ of them are within r of Cx. Choose an average whose closed $2r$ ball contains more than $b_0/2$ averages, or choose any average if none qualifies. On the specified good event such a choice exists, and its $2r$ ball contains a good average. The selected value $\widehat{z}$ therefore satisfies

$$\|\widehat{z} - Cx\|_2 \le 3r < \tau\sqrt{N}. \tag{C.1}$$

For $L = 0$, take $\widehat{z} = 0$: on the promised domain every marked column of C has squared norm zero by the same diagonal inequality. The chosen sample cap is fixed even off the NO promise.

Minimize on the cube the known function

$$\Phi(z) = z^T B z - 2\widehat{z}^T z - \tau \sum_v h_b(z_v).$$

Its Hessian on the interior is at least $4\tau I$, and the infinite inward slopes of h_b exclude a boundary minimizer. Thus it has a unique interior minimizer z, satisfying $2Bz - 2\widehat{z} = \tau h_b'(z)$. Assume now the input is NO and (C.1) holds. Put $d = z - x$ and $e = \widehat{z} - Cx$. Since $x^T A x = 0$, comparison of $\Phi(z)$ with $\Phi(x)$ yields

$$z^T A z + d^T C d \le \tau \left(|x| - \|z\|_2^2 + \sum_v h_b(z_v) \right) + 2e^T d \le 5\tau N. \tag{C.2}$$

Here $\|d\|_2 \le \sqrt{N}$ and $0 \le h_b < 2$; both terms on the left are nonnegative. Stationarity also gives the exact identity

$$\sum_v \Lambda_b(x_v, z_v) = \sum_v h_b(z_v) - \frac{2}{\tau} z^T A x + 2(|x| - z^T x) - \frac{2}{\tau} d^T B d + \frac{2}{\tau} d^T e. \tag{C.3}$$

Its left side is nonnegative, $B \succeq 0$, and $z^T A x \ge 0$. Consequently

$$\sum_v \Lambda_b(x_v, z_v) \le 6N, \qquad x^T A z \le 3\tau N.$$

By Lemma C.1, fewer than $384bN$ marked coordinates satisfy $z_v < 3b$.

Compute $\widetilde{z}$ with $\|\widetilde{z} - z\|_\infty \le b/8$, and let $P = \{v : \widetilde{z}_v \ge 2b\}$. Then

$$\{v : z_v \ge 3b\} \subseteq P \subseteq \{v : z_v \ge b\}.$$

By (C.2), $e_H(P) \le 5\tau N/(2b^2)$. Furthermore,

$$bW_{H[P]}(x) \le \sum_{v \in P} x_v \sum_{u \in P : uv \in E(H)} z_u \le x^T A z \le 3\tau N.$$

These estimates also show that at most $384bN$ marked labels lie outside P on the NO good-field event.

Check the known edge bound $e_H(P) \le 3\tau N/b^2$; report YES if it fails. Next apply fixed-budget quantum approximate counting to $x_v \mathbf{1}_{v \notin P}$. Using a constant multiplicative gap, it costs

$O(b^{-1/2}\log(2/\varepsilon))$ to return LOW or HIGH such that, except with probability $\varepsilon/4$, LOW implies a count at most $12288bN$, whereas HIGH implies a count at least $6144bN$. This follows from amplitude estimation with additive error $O(\sqrt{p}/T+1/T^2)$ at $T=\Theta(b^{-1/2})$, amplified by independent repetitions. Both outcomes may be valid in the overlap. Return YES on HIGH and P on LOW. A NO good-field input forces LOW when counting succeeds. Conversely, on either sign LOW validates the asserted outside-count bound independently of field accuracy. All claims follow by the union bound.

For finite computation, take rational parameters; fixed-factor conservative rationalization preserves the stated bounds. Matrix spectral data and sample means are algebraic and can be computed to certified precision. The objective is effectively continuous on the compact cube, including its entropy endpoints. Finite rational grid search with certified objective gap $2\tau(b/8)^2$ suffices for optimizer error at most $b/8$, by strong convexity. Known rounding error can be absorbed by the strict slack in (C.1) and the cutoff inclusions. This may require large computation, but no further marking queries. Thus every branch has a finite description and the displayed fixed cap. □

Proof of Lemma 4.1. The current graph is always represented as an edge subset on the original n labels. Set $b=2^{-40}\alpha$, $L=\alpha M$, and $\tau=b^2M/(6n)$. Apply Lemma C.2. Every returned set has at most $3\tau n/b^2=M/2$ induced edges. On the promised NO good event its induced marked volume is at most $3\tau n/b=bM/2\le\alpha M/2$. Its query cost is

$$O\Big((1+n^2/(M\alpha^3)+\alpha^{-1/2})\log(2/\varepsilon)\Big).$$

Test the original edges having at least one endpoint outside P. Their outside endpoints form a known vertex cover. The fresh validation in the container construction bounds its marked count by $2^{16}bn$ on either sign, except for allocated error. Theorem B.3 therefore tests this edge list in $\widetilde{O}(\sqrt{n}+n^{2/3}\alpha^{1/6})$ queries. If the detector reports YES, report YES; otherwise retain precisely the known original edges of $H[P]$.

Every original collision is either in the outside test or in this retained edge list. The first case is detected except for its charged error, and the second case survives. On NO inputs satisfying the volume promise, both simultaneous contractions hold and all false YES probabilities are within the allocated budgets. Conditional on any previous transcript, the current graph and declared caps are fixed, so fresh errors can be summed by first-failure accounting.

Finally $M\ge n^{7/4}$ implies $n^2/(M\alpha^3)=O(n^{1/4+3c})$. The validation term $\alpha^{-1/2}=O(n^{c/2})$ is absorbed in this power. Halving M gives $O(\log n)$ epochs; allocate the prescribed total error across them and absorb the resulting fixed logarithmic factors. □

C.2 Source-supported weighted response

C.2.1 Statement and known-data preparation

Let A be the adjacency matrix of a known simple graph on $N\ge 2$ labels, and let I be a known subset of source labels. For a vertex set P, write

$$e(P)=|E(G[P])|,\qquad W(P,x)=\sum_{v\in P}x_v\deg_{G[P]}(v).$$

An independent Boolean marking is a vector $x\in\{0,1\}^N$ with $x^TAx=0$.

Theorem C.3 (Source-supported weighted response)**.** *Let $\mathcal{F}\subseteq[0,1]^N$ be a known rational compact polytope, and let $\mathcal{S}$ be its independent Boolean points. Let $0<K\le N$, $\tau>0$, and $0<\beta\le 1$ be*

known rational parameters. Let T be a known nonnegative rational diagonal matrix and let C, R be known rational matrices satisfying

$$C \succeq 0, \qquad M = A + T + C \succeq 0, \qquad R_{ji} = 0 \text{ unless } i, j \in I.$$

Assume, for fixed absolute constants C_T, B, that

$$\sum_{i \in I} a_i \le K, \qquad \sum_i T_{ii} a_i \le C_T \tau N \qquad (a \in \mathcal{F}), \tag{C.4}$$

$$|((R - C)x)^T (a - x)| \le B\tau N \qquad (a \in \mathcal{F},\ x \in \mathcal{S}), \tag{C.5}$$

$$\sum_{j,i} a_j x_i R_{ji}^2 \le V \qquad (a \in \mathcal{F},\ x \in \mathcal{S}), \tag{C.6}$$

where $V \ge 0$ is a known rational bound. Put $F_0 = KV/(\tau^2 N)$. For every prescribed $0 < \varepsilon < 1/2$, there is a finite procedure with all-input query cap

$$\widetilde{O}(1 + \sqrt{F_0} + F_0\sqrt{K/N} + \beta^{-1/2}),$$

where the suppressed factors are fixed powers of $\log(2N)$ and $\log(2/\varepsilon)$. It reports YES or returns a classical vertex set P. Every returned set obeys the deterministic edge bound

$$e(P) \le C_E \tau N/\beta^2.$$

Except with probability ε, on either input sign a returned set also obeys

$$W(P, x) \le C_W \tau N/\beta, \qquad \sum_{v \notin P} x_v \le C_O \beta N.$$

On every input $x \in \mathcal{S}$, the probability of reporting YES is at most ε. The constants C_E, C_W, C_O depend only on C_T and B and may be chosen explicitly in advance.

All required NO inputs lie in $\mathcal{S}$. The theorem does not assert NO correctness for independent markings outside $\mathcal{F}$. Rationality is an effective-description condition, not a numerical conditioning assumption. The diagonal $T = \tau q I$ used in the main construction satisfies (C.4) whenever the frozen constraints give $\sum_i a_i \le C_T N/q$.

Enumerate $\mathcal{S}$ using known graph and polytope data. If it is empty, the algorithm can report YES without queries. Otherwise replace $\mathcal{F}$ by $\mathcal{P} = \mathrm{conv}(\mathcal{S})$. This preserves every promised NO input and all the displayed assumptions. It is a finite, query-free computation, although it can be exponential in N. For each coordinate not fixed on $\mathcal{S}$, choose two points of $\mathcal{S}$ attaining zero and one. Their average, denoted $a^{(1)}$, uses at most $2N$ points and has every nonfixed coordinate in $[1/(2N), 1 - 1/(2N)]$. If no coordinate is free, use the unique point. Fixed coordinates are handled on their known face throughout.

For $K < 1$, every promised source bit is zero and the field Rx is zero. The zero-field construction below applies immediately. Hence in the sampling argument assume $K \ge 1$. Append the source slack

$$a_0 = 2K - \sum_{i \in I} a_i \in [K, 2K]$$

and use generalized relative entropy on $I \cup \{0\}$:

$$D(p, a) = \sum_{i \in I \cup \{0\}} [p_i \log(p_i/a_i) - p_i + a_i].$$

Identically zero coordinates are omitted. Both augmented vectors have mass $2K$, so the linear terms cancel. For any $x \in \mathcal{S}$,

$$D(x, a^{(1)}) \le K \log(2N) + 2K \log 2 =: \mathcal{R}. \tag{C.7}$$

Indeed each marked free source contributes at most $\log(2N)$, fixed marked coordinates contribute zero, and $x_0/a_0^{(1)} \le 2$. This gives the required radius without assuming a well-conditioned description of $\mathcal{F}$.

C.2.2 Cutoff losses and weighted identities

It suffices to prove the result for $0 < \beta \le 1/64$. For larger $\beta \le 1$, run the procedure with $1/64$; multiplying the edge and volume constants by at most 64^2 gives the stated bounds and only changes absolute query constants.

Use the cutoff function of Lemma C.1, with parameter β:

$$h_\beta(u) = -u \log u - (1-u)\log(1-u) + \frac{u}{u+\beta} + \frac{1-u}{1-u+\beta} - \frac{1}{1+\beta}, \qquad 0 \le u \le 1,$$

with the continuous endpoint convention. Its properties are $h_\beta(0) = h_\beta(1) = 0$, $0 \le h_\beta < 2$, and $h''_\beta \le -4$. Concavity gives nonnegative endpoint tangent losses

$$L_h(b, u) = h_\beta(u) + (b-u)h'_\beta(u) \ge 0, \qquad b \in \{0, 1\}, \quad 0 < u < 1.$$

For completeness, when $u \le 3\beta$,

$$L_h(1, u) \ge \frac{(1-u)\beta}{(u+\beta)^2} - 1 \ge \frac{1-3\beta}{16\beta} - 1 \ge \frac{1}{32\beta}.$$

The first inequality discards the nonnegative entropy loss $-\log u$ and the nonnegative tangent contribution $(1-u)^2/(1-u+\beta)^2$ of the second rational term. The last inequality uses $\beta \le 1/64$. The upper bound $h_\beta < 2$ follows from $\log 2 + 2 - 1/(1+\beta) < 2$, and differentiating gives $h''_\beta \le -4$. At a fixed Boolean coordinate both a_i and x_i equal that endpoint; its loss is defined as zero, with no singular derivative evaluated.

Write $H(a) = \sum_i h_\beta(a_i)$. For a fixed field z define

$$G(a) = a^T M a - \tau H(a), \qquad F_z(a) = G(a) - 2z^T a.$$

The known part G is strongly convex on the free coordinates. For $x \in \mathcal{S}$, $d = a - x$, $e = z - Cx$, and $g = \nabla F_z(a)^T d$, direct expansion and convexity give

$$J := a^T A a + d^T C d \le x^T T x - a^T T a + \tau H(a) + 2e^T d + g, \tag{C.8}$$

$$L_h(x, a) + \frac{2}{\tau} a^T A x + \frac{2}{\tau} d^T M d = H(a) + \frac{2}{\tau} x^T T(x - a) + \frac{2}{\tau} e^T d + \frac{g}{\tau}. \tag{C.9}$$

Here $L_h(x, a)$ is the sum of coordinate losses. To verify the first identity before applying convexity, expand

$$F_z(a) - F_z(x) = J + a^T T a - x^T T x - \tau H(a) - 2e^T d$$

and use $F_z(a) - F_z(x) \le g$. For the second, expand

$$g = 2d^T M d + 2a^T A x + 2x^T T(a - x) - 2e^T d - \tau \nabla H(a)^T d.$$

This uses $x^T Ax = 0$. By (C.4), $x^T Tx \le C_T \tau N$ and $x^T T(x-a) \le C_T \tau N$. Thus if

$$|e^T(a-x)| \le B'\tau N \quad (a \in \mathcal{P}) \tag{C.10}$$

for an absolute B', and a random feasible a satisfies $\mathbb{E}g \le \tau N/64$, then

$$\mathbb{E}J \le (C_T + 2 + 2B' + 1/64)\tau N, \tag{C.11}$$

$$\mathbb{E}\left[L_h(x,a) + 2a^T Ax/\tau + 2d^T Md/\tau\right] \le (2 + 2C_T + 2B' + 1/64)N. \tag{C.12}$$

Every quantity on these left sides is nonnegative. These estimates depend on the total frozen diagonal mass, even though the stochastic field and entropy below are supported only on I.

C.2.3 Clipping and a fixed-budget source count

Suppose $V > 0$. Clip every entry of R to $[-b, b]$, where $b = 16V/(\tau N)$, and denote the result by $\bar{R}$. The elementary bound $|R_{ji} - \bar{R}_{ji}| \le R_{ji}^2/b$ gives, because both a and x are feasible,

$$|(a-x)^T(R-\bar{R})x| \le b^{-1}\sum_{ji}(a_j + x_j)x_i R_{ji}^2 \le 2V/b = \tau N/8. \tag{C.13}$$

Clipping preserves (C.6) and source support. If $V = 0$, weighted Cauchy–Schwarz below shows $a^T Rx = x^T Rx = 0$ for all feasible a and promised x; the field can be omitted on the promised domain and no count or sampling is needed.

Let $k = \sum_{i \in I} x_i$ and take $C_0 = 2^{20}$,

$$\Delta = \frac{N}{C_0 F_0}.$$

Fixed-budget amplitude estimation followed by median amplification produces a rational estimate $\widehat{k}$, clamped to $[0, \min(2K, N)]$, with

$$|\widehat{k} - k| \le .1\sqrt{k\Delta} + .01\Delta \tag{C.14}$$

except with prescribed error. Its cap is $O((1+\sqrt{F_0})\log(2/\varepsilon))$. Indeed the usual amplitude-estimation count error is $2\pi\sqrt{k(N-k)}/J + \pi^2 N/J^2$ with a fixed positive success probability above $1/2$. An even $J \ge C\sqrt{N/\Delta}$, a sufficiently large absolute C, and arbitrarily fine rational rounding give (C.14) with slack. Clamping cannot increase error when $k \le K$. The count uses the masked source marking predicate, implementable with $O(1)$ physical marking queries per invocation.

Weighted Cauchy–Schwarz yields, uniformly over $a \in \mathcal{P}$,

$$|a^T \bar{R}x| \le \sqrt{KkV}, \qquad |x^T \bar{R}x| \le \sqrt{KkV}. \tag{C.15}$$

If $\widehat{k} \le 2\Delta$, the good-count event implies $k \le 3\Delta$. Use $z = 0$; the additional comparison bias is at most $2\sqrt{3\Delta KV} = 2\sqrt{3/C_0}\,\tau N$. Thus (C.10) holds with $B' = B + 1$. Minimize the known objective G to a certified first-order gap $\tau N/1024$, and apply the validations below.

If $\widehat{k} > 2\Delta$, the good-count event implies $k > \Delta$ and $.89k \le \widehat{k} \le 1.11k$. Define the effective mean field $z = (\widehat{k}/k)\bar{R}x$. The further deterministic comparison error is at most

$$2\sqrt{KV}\left(.1\sqrt{\Delta} + .01\Delta/\sqrt{k}\right) \le .22\sqrt{\Delta KV} = .22\tau N/\sqrt{C_0}.$$

Together with (C.5) and (C.13), this proves (C.10) with $B' = B + 1$ also in the high branch. The algorithm never computes k or z exactly; they are fixed quantities used only to analyze the noisy updates.

C.2.4 Source-only implicit updates and a paid sampler

Condition on a good count transcript in the high branch. For a fresh uniform marked source label i, use the known classical column

$$Z = \widehat{k}\,\bar{R}e_i, \qquad \mathbb{E}Z = z.$$

If a is determined by the previous successful labels and $u = 2(Z - z)$, then

$$\mathbb{E}\sum_{j\in I} a_j u_j^2 \le 4\widehat{k}^2 V/k \le 8\widehat{k}V, \qquad |u_j| \le 4\widehat{k}b. \tag{C.16}$$

Choose

$$\eta = \frac{\tau N}{4096\widehat{k}V}, \qquad L = \max\left\{1, \left\lceil \frac{1024\mathcal{R}}{\eta\tau N} \right\rceil \right\}.$$

Starting at $a^{(1)}$, perform L known-data implicit updates

$$a^{(t+1)} = \arg\min_{a\in\mathcal{P}} \left\{ \eta[G(a) - 2Z_t^T a] + D(a, a^{(t)}) \right\}. \tag{C.17}$$

The entropy divergence involves sources and their slack only. The objective G and cutoff term still involve every free coordinate. Hence the target coordinates have a known convex optimization step without an uncharged target-dependent stochastic field.

For exact updates the three-point Bregman identity gives, for fixed $x \in \mathcal{S}$ and $g_t = \nabla F_z(a^{(t+1)})^T(a^{(t+1)} - x)$,

$$\begin{aligned}\eta g_t \le\ & D(x, a^{(t)}) - D(x, a^{(t+1)}) - D(a^{(t+1)}, a^{(t)}) \\ & + \eta u_t^T(a^{(t+1)} - x).\end{aligned}$$

Split the final term at $a^{(t)}$. Only $u_t^T(a^{(t)} - x)$ has zero conditional mean. For the dependent part, coordinatewise entropy conjugacy gives

$$\begin{aligned}&\eta u_t^T(a^{(t+1)} - a^{(t)}) - D(a^{(t+1)}, a^{(t)}) \\ &\qquad \le \sum_{j\in I} a_j^{(t)}(e^{\eta u_{t,j}} - 1 - \eta u_{t,j}) \le \eta^2 \sum_{j\in I} a_j^{(t)} u_{t,j}^2.\end{aligned}$$

The first inequality maximizes each coordinate over the nonnegative orthant; the slack has zero noise. The second uses $|\eta u_{t,j}| \le 64/4096 < 1$ and $e^v - 1 - v \le v^2$ for $|v| \le 1$. Summing, using (C.7) and (C.16), and dropping a nonnegative final divergence proves

$$\mathbb{E}\frac{1}{L}\sum_{t=1}^{L} g_t \le \frac{\mathcal{R}}{\eta L} + 8\eta\widehat{k}V \le \frac{\tau N}{1024} + \frac{\tau N}{512}.$$

Choosing a uniformly random update therefore supplies the required expected residual. This calculation does not assert independence of the new iterate and its sample.

Finite optimization. The exact minimizer in (C.17) has every nonfixed cube coordinate in $(0, 1)$. Otherwise moving toward the positive feasible point $a^{(1)}$ decreases the negative entropy contribution by order $s\log(1/s)$ while smooth terms change by $O(s)$. This contradicts minimality. At each step it suffices to find a rational feasible point whose every nonfixed coordinate lies strictly between zero and one and whose certified variational-inequality gap is at most $\zeta = \eta\tau N/1024$ for the known convex objective in (C.17). The maximum of its gradient pairing against $b \in \mathcal{P}$ is attained at a point

of the enumerated finite set $\mathcal{S}$, so it is a known-data finite calculation. Interval evaluation certifies the logarithms. Interior rational feasible points with arbitrarily small gap exist by continuity near the exact minimizer; enumerating and refining such points terminates. Explicitly, for the known objective Ψ_t in (C.17), certify

$$\max_{b\in\mathcal{S}} \nabla\Psi_t(a)^T(a-b) \le \zeta.$$

Search for a point with gap strictly below ζ so that interval certification terminates; existence follows by approaching the exact minimizer. This criterion also implements the zero-field minimization, using $\Psi = G$ and tolerance $\tau N/1024$. The Bregman bound gains at most $\zeta/\eta = \tau N/1024$ per averaged step, leaving $\mathbb{E}g < \tau N/64$. This finite procedure incurs no further marking queries.

A fixed-cap verified uniform sampler. Each independent trial mixes, with probability $1/2$ each, a direct uniform-label query and a Grover experiment from the uniform N-label state. In the latter choose the number of Grover iterations uniformly from $0,\dots,m-1$, where $m = \lceil 2\sqrt{N/\widehat{k}}\rceil$. Measure a label and verify its source marking. Every trial costs $O(1+\sqrt{N/\widehat{k}})$ queries on every input. On the good-count event its success probability is bounded below by a positive absolute constant. For $k \le N/2$, put $\phi = \arcsin\sqrt{k/N}$; the Grover branch has mean success probability

$$\frac{1}{2} - \frac{\sin(4m\phi)}{4m\sin(2\phi)}.$$

Since $m \ge 2\sqrt{N/(1.11k)}$ and $\sin(2\phi) \ge \sqrt{2k/N}$, this is at least $3/8$. For $k > N/2$ the direct branch succeeds with probability at least $1/2$. Thus the mixed trial succeeds with probability at least $3/16$ in all cases. Conditional on success its label is exactly uniform over the marked source set, by the symmetry of uniform initialization and the masked marking oracle.

For L successful labels allow only $C(L+\log(2/\delta))$ trials, with a sufficiently large absolute C. The probability of failing to obtain L labels on a good count transcript is at most δ by a binomial tail bound. For fixed input and count transcript each trial factors into a success indicator and an independent uniform successful label. Different trials are independent. The cap event depends only on the indicators. Consequently the successful-label sequence remains iid uniform conditional on obtaining L successes before the cap, exactly as required above.

On every transcript entering this branch, $\widehat{k} > 2\Delta$ and $\widehat{k} \le \min(2K, N)$, so the capped trajectory has, including fixed logarithms, query cost

$$\begin{aligned}\widetilde{O}\left(1+\sqrt{N/\widehat{k}} + \frac{\mathcal{R}\widehat{k}V}{\tau^2N^2}(1+\sqrt{N/\widehat{k}})\right) &= \widetilde{O}\left(1+\sqrt{F_0} + \frac{\mathcal{R}V\sqrt{K}}{\tau^2N^{3/2}}\right)\\ &= \widetilde{O}\left(1+\sqrt{F_0}+F_0\sqrt{K/N}\right).\end{aligned}$$

The final equality uses (C.7). This is an all-input cap; an incorrect count or an off-promise input can cause a capped trial failure but cannot extend the runtime. Reporting YES upon such failure is charged to the NO failure budget.

C.2.5 Conversion and all-sign validation

Take a random update as above and set explicitly

$$U = C_T + 2 + 2(B+1) + 1/64, \qquad Z_0 = 2 + 2C_T + 2(B+1) + 1/64,$$

$$C_1 = 12U, \qquad C_2 = 6Z_0, \qquad C_3 = 12Z_0.$$

Equations (C.11) and (C.12) and Markov's inequality give failure probability at most $1/12$ for each of the following three bounds. Thus all hold simultaneously with probability at least $3/4$:

$$a^T Aa \le C_1 \tau N, \qquad x^T Aa \le C_2 \tau N, \qquad L_h(x, a) \le C_3 N,$$

For $P_a = \{i : a_i \ge 2\beta\}$, nonnegativity and the cutoff losses imply

$$e(P_a) \le \frac{C_1 \tau N}{8\beta^2}, \qquad W(P_a, x) \le \frac{C_2 \tau N}{2\beta}, \qquad \sum_{i \notin P_a} x_i \le 32 C_3 \beta N.$$

The same implications hold deterministically for the approximate zero-field minimizer on its good count event. All iterates are rational, so thresholding is an exact finite computation.

Write $A_W = C_2 \tau N/(2\beta)$ and $A_O = 32 C_3 \beta N$. Choose

$$C_E = C_1/2, \qquad C_W = 2C_2, \qquad C_O = 128 C_3$$

for $\beta \le 1/64$. Reject a proposed set if its known edge count exceeds $C_E \tau N/\beta^2$. For any remaining set perform two fresh count tests, each with fixed query budget and a constant-factor gap between its accept and reject thresholds. The first counts marks in $[N] \setminus P_a$, using the ambient N labels and thresholds $2A_O, 4A_O$. The second counts marked oriented incidences of $G[P_a]$, using a known universe of size $2e(P_a)$ and thresholds $2A_W, 4A_W$. An incidence query is one marking query to its known tail. The deterministic edge filter makes the ratio of the second universe size to its threshold $O(1/\beta)$; the first has the same ratio. Amplitude estimation therefore implements each overlapping LOW/HIGH test in $O(\beta^{-1/2} \log(2/\delta))$ queries. Here is an explicit gap test, including thresholds below one. For a known universe of size m and parameter $A > 0$, return LOW without queries when $m \le 4A$. Otherwise use amplitude estimation with $J \ge 100\pi\sqrt{m/A}$, rounded up to an even integer, and median amplification. Its estimate $\widehat{k}$ of the marked count k obeys

$$|\widehat{k} - k| \le 0.02\sqrt{kA} + 0.0001A$$

except for the assigned error. Return LOW if $\widehat{k} \le 3A$ and HIGH otherwise. If $k \le 2A$, the upper error bound gives $\widehat{k} < 3A$; if $k > 4A$, the lower bound gives $\widehat{k} > 3A$. Thus the desired guarantees hold for every $A > 0$, without an integrality assumption. The cap is $O((1 + \sqrt{m/A}) \log(2/\delta))$. Since $m/A = O(1/\beta)$ in both applications, it is within the asserted validation budget.

The count tests have the following overlapping guarantees outside their allocated bad events: LOW implies the true count is at most $4A$, and HIGH implies it is at least $2A$, with $A = A_W$ or A_O as appropriate. Consequently a candidate whose true count is at most A passes, and every returned LOW certifies the declared cap on every marking. This is why all-sign count and volume guarantees follow even though the stochastic comparison proof is only for promised NO inputs.

Perform the initial count with error at most $\varepsilon/8$ and then $r = \lceil \log_2(8/\varepsilon) \rceil$ independent trajectories, each with fresh sampler coins and validations. Give each sampler cap error at most $\varepsilon/(8r)$, and each of the two validations in an attempt error at most $\varepsilon/(16r)$. Return the first validated candidate; if none passes, report YES. Conditional on a good promised-NO count transcript, couple every trajectory to its ideal iid successful-label stream. Each ideal candidate has failure probability at most $1/4$, independently across attempts, so all ideal candidates fail with probability at most $(1/4)^r \le \varepsilon/8$. On the complement of all sampler and validation errors any good candidate is returned or an earlier valid candidate has already been returned. The error sum is at most $\varepsilon/8 + \varepsilon/8 + \varepsilon/8 + \varepsilon/8 < \varepsilon$. On arbitrary input signs only the $2r$ validation error events are needed to guarantee the caps of a returned set; their sum is at most $\varepsilon/8$. All branches and trial counts are capped in advance. This completes the proof of Theorem C.3.

C.2.6 Application to the manuscript's Schur field

For target-by-source incidence B, use

$$A = \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix}, \quad T = \tau q I, \quad C = 0 \oplus \frac{B^T B}{\tau q}, \quad R = 0 \oplus \frac{B^T B - \operatorname{diag}(d_u)}{\tau q}.$$

The Schur identity gives $C \succeq 0$ and $A + T + C \succeq 0$. The frozen total feasible count gives (C.4), and the stated source cap gives its source-count part. For a source marking count at most K and maximum source degree d,

$$|((R - C)x)^T (a - x)| \leq dK/(\tau q).$$

Thus the manuscript's condition $\tau^2 q^2 \geq C_0 d$, together with $K \leq C_K N/q$, supplies (C.5) with an absolute constant. The residual codegree estimate supplies (C.6) uniformly over feasible responses and promised NO inputs. All known polytopes and scalar choices can be taken rational. The theorem therefore proves Lemma 4.2 under its stated frozen mass and variance premises. Section 5 verifies these premises for every invocation and combines this routine with the other proved algorithms.

Use of AI: The author prompted Claude and ChatGPT interactively and verified the proofs manually.

References

[1] F. Magniez, M. Santha, and M. Szegedy, *Quantum algorithms for the triangle problem*, SIAM Journal on Computing 37(2) (2007), 413–424. arXiv:quant-ph/0310134.

[2] A. Belovs, *Learning-Graph-Based Quantum Algorithm for k-distinctness*, 2012, Section 4. arXiv:1205.1534.

[3] D. Gavinsky and T. Ito, *A quantum query algorithm for the graph collision problem*, 2012. arXiv:1204.1527.

[4] A. Ambainis, K. Balodis, J. Iraids, R. Ozols, and J. Smotrovs, *Parameterized Quantum Query Complexity of Graph Collision*, 2013. arXiv:1305.1021.

[5] S. Jeffery, R. Kothari, and F. Magniez, *Nested Quantum Walks with Quantum Data Structures*, Proceedings of SODA 2013, 1474–1485; Appendix A of the arXiv version. arXiv:1210.1199.

[6] S. Jeffery, R. Kothari, and F. Magniez, *Improving Quantum Query Complexity of Boolean Matrix Multiplication Using Graph Collision*, ICALP 2012, 522–532. arXiv:1112.5855.

[7] S. Jeffery, R. Kothari, F. Le Gall, and F. Magniez, *Improving Quantum Query Complexity of Boolean Matrix Multiplication Using Graph Collision*, Algorithmica 76(1) (2016), 1–16. doi:10.1007/s00453-015-9985-x. Author-hosted journal article.

[8] K. Balodis and J. Iraids, *Quantum Lower Bound for Graph Collision Implies Lower Bound for Triangle Detection*, 2015. arXiv:1507.03885.

[9] A. Cornelissen, A. S. Gilani, and S. Patro, *Quantum algorithms for path and cycle containment problems*, 2026, version 2. arXiv:2605.09017v2.

[10] H. Aaronson, T. Gur, and J. Li, *Pseudo-deterministic Quantum Algorithms*, 2026. arXiv:2602.17647.

[11] A. M. Childs, S. Kimmel, and R. Kothari, *The quantum query complexity of read-many formulas*, 2011. arXiv:1112.0548.

[12] G. Brassard, P. Høyer, M. Mosca, and A. Tapp, *Quantum Amplitude Amplification and Estimation*, Contemporary Mathematics 305 (2002), 53–74. arXiv:quant-ph/0005055.

[13] M. Boyer, G. Brassard, P. Høyer, and A. Tapp, *Tight bounds on quantum searching*, Fortschritte der Physik 46 (1998), 493–506. arXiv:quant-ph/9605034.

[14] P. Høyer, M. Mosca, and R. de Wolf, *Quantum search on bounded-error inputs*, ICALP 2003, 291–299. arXiv:quant-ph/0304052.

[15] B. W. Reichardt, *Reflections for quantum query algorithms*, SODA 2011, 560–569. arXiv:1005.1601.

[16] T. Lee, R. Mittal, B. W. Reichardt, R. Špalek, and M. Szegedy, *Quantum query complexity of state conversion*, FOCS 2011, 344–353. arXiv:1011.3020.

[17] A. Juditsky, A. Nemirovski, and C. Tauvel, *Solving variational inequalities with Stochastic Mirror-Prox algorithm*, Stochastic Systems 1(1) (2011), 17–58. arXiv:0809.0815.