

Interactive Proof of Proximity for Bipartiteness without Mixing

Guy Weissenberg*

Abstract

Testing whether a bounded-degree n -vertex graph is bipartite or far from bipartite requires $\tilde{\Theta}(\sqrt{n})$ queries (Goldreich–Ron, *Combinatorica*’99, *Algorithmica*’02). An interactive proof of proximity (IPP) is a hybrid model of property testing and interactive proofs: the verifier faces the same task with the same query access, but it now interacts with an untrusted prover that sees the whole graph. Rothblum, Vadhan, and Wigderson (STOC’13) introduced this model and gave a one-round private-coin IPP for bipartiteness on well-mixing graphs with query and communication costs both $O(\log n)$ for fixed parameters, an exponential improvement over plain testing. The mixing assumption was used in their soundness analysis, and they asked whether it could be removed. We give a slightly modified version of the original RVW protocol and prove its soundness without the mixing promise. In the RVW protocol, soundness follows by bounding the probability of correctly guessing the parity of a hidden random walk from its start and endpoint. The best guess is wrong with probability equal to the *overlap* between the endpoint probabilities for even and odd numbers of edge traversals, and mixing gives one way to bound this overlap. We prove an overlap theorem that requires no mixing, building on the recent semidefinite programming (SDP) analysis of the Goldreich–Ron tester by Fei and Rubinfeld (2026).

1 Introduction

Property testing asks whether a large object has a given property or is far from every object with that property, while inspecting only a small part of the object [16, 9]. We consider bipartiteness in the bounded-degree graph model, where a graph of n vertices is accessed through indexed neighbor queries and distance is the number of edge insertions and deletions divided by Dn , for D the degree bound. A bipartition is easy to check once the assignment of vertices to parts is given, but bipartiteness is a global property: a graph may be far from bipartite while every neighborhood of logarithmic radius is a tree, and trees are bipartite. The fundamental work of Goldreich and Ron [10] gave a tester that runs many random walks from a sampled start vertex and rejects when two of them end at the same vertex with different parities of the number of edge traversals. For fixed D and proximity parameter it uses $\tilde{O}(\sqrt{n})$ queries, and since $\Omega(\sqrt{n})$ queries are necessary [11], the query complexity of testing bipartiteness is $\tilde{\Theta}(\sqrt{n})$.

Interactive proofs of proximity are the interactive analogue of property testing, introduced for delegating computation in sublinear time [15]. The verifier queries the graph while interacting with a powerful but untrusted prover that sees the graph in full, and the basic question is how much the prover can lower the cost of testing. Rothblum, Vadhan, and Wigderson (RVW) gave IPPs with $n^{1/2+o(1)}$ query and communication complexity for languages decided by log-space-uniform circuits of polynomial size and depth $n^{o(1)}$, including uniform NC, for a fixed proximity parameter and input length n [15, Theorem 1.1]. Separately, they gave a direct one-round private-coin protocol

*EPFL. guy.weissenberg@epfl.ch.

for bipartiteness with logarithmic query and communication costs for fixed parameters under a mixing promise [15, Section 5.1]. In this protocol, the verifier runs random walks, reveals only the start and endpoint of each, and asks the prover for the hidden traversal parity. A prover holding a bipartition reads every parity off the parts of the revealed vertices, so completeness is perfect. To analyze the cheating prover’s strategy, RVW label every endpoint by its more likely traversal parity, and because the graph is far from bipartite, many edges have both vertices labeled the same. The walk from that start traverses one of them in its last step with probability $\Omega(\varepsilon)$ when the graph is well-mixing, but without mixing no such guarantee can hold for every start: a far graph may contain large bipartite components, on which the prover’s task is easy. They asked whether this assumption could be removed [15, Section 1.1.1]: “An interesting open problem is extending this IPP to apply to all graphs (not just rapidly-mixing ones).”

Subsequent work considered noninteractive [12], honest-verifier zero-knowledge [3], public-coin [2], quantum [4], and doubly-sublinear [1] proofs of proximity for bipartiteness. In each case, the bipartiteness protocol assumes that the graph is well- or rapidly-mixing. Thus the question posed by RVW remains open for arbitrary bounded-degree graphs.

In this work we show that a small modification of the RVW protocol can be analyzed without the mixing assumption:

Theorem 1.1 (Informal main theorem). *In the bounded-degree model, bipartiteness has a one-round, private-coin interactive proof of proximity whose soundness holds for all graphs, with no mixing promise. The protocol has perfect completeness, and for constant values of the degree bound and proximity parameter, its communication and query complexities are logarithmic in the number of vertices.*

In Appendix A we prove a matching lower bound, namely that on some far graphs any parity-challenge protocol needs $\Omega(\log n)$ expected edge traversals to reject every prover with constant probability. In Appendix B, we give an alternative bipartiteness protocol and extend it to Unique Games in Appendix C, both based on a recent work by Yoshida [17].

Concurrent work. In concurrent¹ and independent work, Goldreich [8, Section 6] also obtains a one-round, private-coin IPP for bipartiteness without a mixing promise. Both works are based on RVW [15] and the SDP analysis of Fei and Rubinfeld [6]. In [8], the prover is asked to identify which endpoint of a sampled edge started a walk of randomly chosen length, given the edge, walk length, endpoint, and traversal parity. A better query bound is obtained in [8], $O(\varepsilon^{-5} \log n)$, with no dependence on D using the treatment of degree bounds in [8, Section 4], compared with our $O(\varepsilon^{-13}(1 + \log D) \log n)$.

Acknowledgements

The author thanks Tal Herman and Oded Goldreich for their comments.

Large language models were used during the preparation of this paper for editorial and consistency checks, mathematical exploration, and exposition. In particular, they helped identify that the original protocol already possessed a mild form of tolerance without requiring a bound on the number of incorrect answers, suggested clearer presentations of the completeness and soundness proofs, and motivated replacing a generic lower-bound argument over any suitable prime with a

¹We completed this work in early August, and the present text dates largely to that time. The changes in the text, compared to the early-August version, are the added comment on the digest [8] (which was posted on September 20, 2026), and the full version of the alternative protocols based on a theorem by Yoshida [17] in the appendix, which were only hinted at in the early-August version.

concrete instantiation. The author independently checked all mathematical arguments and claims about prior work and assumes responsibility for all content.

1.1 Technical overview

As in RVW, the verifier runs a random walk, reveals only its start and endpoint, and asks the prover whether the walk made an even or an odd number of edge traversals. On a bipartite graph, the prover answers correctly by checking whether the two revealed vertices lie in the same part of a bipartition. Soundness on a far graph reduces to showing that the best guess from the revealed pair errs with probability bounded away from zero. For a fixed start and endpoint, the prover answers with whichever parity is more likely, and the error at that endpoint is the smaller of the probabilities of reaching it with even parity and with odd parity. Summing these smaller probabilities over all endpoints gives the least possible error in one challenge, which we call the *overlap*. Since the start is itself uniformly random, the least error in a single challenge is the overlap *averaged* over starts.

Using the recent SDP-based techniques of Fei and Rubinfeld [6], we show that on every graph far from bipartite, a lazy walk of logarithmic length has average overlap bounded away from zero when its start is drawn from the stationary distribution. Using this bound to analyze our variant of the RVW protocol (and accounting for the fact that the verifier uniformly chooses starts whereas the overlap bound is for the stationary distribution), we obtain soundness without any mixing promise.

The parity challenge. We use $T = \Theta(\varepsilon^{-5})$ walks of length $L = \Theta(\varepsilon^{-8} \log n)$, where ε is the proximity parameter.

1. The verifier runs T lazy random walks of length L , each from a fresh uniformly random vertex.
2. It sends the prover the start and the endpoint of every walk, and the prover responds with the parity of the number of edge traversals of each hidden walk.
3. The verifier accepts if all T responses are correct.

The error of one response. We fix a start s , denote by X_L the endpoint of a walk of length L from s , and let $\sigma_L \in \{+1, -1\}$ be $+1$ for an even number of edge traversals and -1 for an odd one. For each endpoint v and $b \in \{\pm 1\}$, we define

$$p_{L,s}^b(v) := \mathbb{P}[X_L = v, \sigma_L = b \mid X_0 = s].$$

The prover sees only the pair (s, X_L) , so its best response is the more likely parity. We define

$$\text{ov}_L(s) := \sum_{v \in [n]} \min\{p_{L,s}^{+1}(v), p_{L,s}^{-1}(v)\}, \quad \alpha_L := \mathbb{E}_{s \sim \text{Unif}([n])}[\text{ov}_L(s)].$$

By Lemma 2.6, $\text{ov}_L(s)$ is the minimum guessing error for a walk from s , namely the overlap of the endpoint probabilities from s for even and odd numbers of edge traversals, and α_L is its average over a uniform start. Zero overlap means that the endpoint determines the parity. The prover may choose the more likely parity separately for each revealed pair (s, v) , and since the T walks are independent, this remains optimal even if it chooses all responses jointly from the full message. Its largest acceptance probability is therefore²

$$\left(\frac{1}{n} \sum_{s \in [n]} \sum_{v \in [n]} \max\{p_{L,s}^{+1}(v), p_{L,s}^{-1}(v)\} \right)^T = \left(1 - \frac{1}{n} \sum_{s \in [n]} \text{ov}_L(s) \right)^T = (1 - \alpha_L)^T \leq e^{-T\alpha_L}. \quad (1.1)$$

²For the first equality, we use $\max\{a, b\} = a + b - \min\{a, b\}$ and, for every start s , $\sum_{v \in [n]} (p_{L,s}^{+1}(v) + p_{L,s}^{-1}(v)) = 1$.

Soundness therefore reduces to one question: at which walk length L does every graph far from bipartite have average overlap α_L bounded away from zero?

We note that the Goldreich–Ron tester succeeds by comparing many walks, but from its success probability alone we cannot conclude that a prover has a constant chance of answering a single challenge incorrectly.

Bounding the overlap with mixing. RVW use mixing to prove a stronger, per-start bound on the overlap. A graph is *well-mixing* when the L -step walk from every start ends at each vertex with probability at least $1/(2n)$. For a fixed start s , every vertex is labeled by its more likely parity, with ties broken arbitrarily, and since the graph is far, at least εDn edges violate this 2-partition. Comparing the parity probabilities at lengths $L - 1$ and L using holding steps, RVW bound the overlap below by $\Omega(1/(nD))$ per violating edge, through walks that traverse it in the final step. Hence $\text{ov}_L(s) \geq \varepsilon Dn \cdot \Omega(1/(nD)) = \Omega(\varepsilon)$ from every start, and $\Theta(1/\varepsilon)$ challenges suffice. Without mixing, the walk from s may rarely reach the violating edges: a start in a bipartite connected component, for example, has zero overlap for every walk length.

Bounding the overlap without mixing. RVW lower-bound the overlap from every start. For soundness, we can allow perfect answers from some starts, provided the average guessing error is bounded away from zero. Let

$$\mu(v) := \deg_G(v)/(2m)$$

be the degree-weighted distribution, where m is the number of edges. The main technical claim, stated formally in [Theorem 2.7](#), is that there are absolute constants $c, C > 0$ such that every graph that is ε -far from bipartite and every integer $L \geq C\varepsilon^{-8} \log_2 n$ satisfy

$$\mathbb{E}_{s \sim \mu}[\text{ov}_L(s)] \geq c\varepsilon^4.$$

The lazy walk and RVW’s walk have different stationary distributions: in the lazy walk the transition probability along $v \rightarrow u$ is $1/(2 \deg_G(v))$ and $\mu(v)/(2 \deg_G(v)) = 1/(4m)$, so μ is stationary, whereas in RVW’s walk both directions of every edge have transition probability $1/(2D)$, so the uniform distribution is stationary.

We use the SDP techniques of Fei and Rubinfeld [6]. Suppose that G is ε -far from bipartite but that the stationary average is small, so that for most starts the endpoint of a walk almost determines its traversal parity. The more likely parity may depend on the start, so these guesses do not directly define a single 2-partition. Nevertheless, using their SDP analysis, we can construct a near-optimal Max-Cut SDP solution and round it to a 2-partition with fewer than εDn violating edges, contradicting farness.

The walk and starting distribution. We use the lazy walk because [Theorem 2.7](#) is proved for it, and its stationary distribution is the degree-weighted μ . The verifier’s query access samples a uniform vertex for free, whereas sampling from μ would cost extra queries, so we keep uniform starts and use the direct comparison below, losing a factor of ε in the stated overlap bound. RVW’s uniform stationary distribution already matches the verifier’s starting distribution, so they incur no such loss.

- **Completeness.** A prover holding a bipartition answers perfectly by reporting whether the two revealed vertices lie in the same part.
- **Soundness.** By the overlap theorem, for $L = \Theta(\varepsilon^{-8} \log n)$, $\mathbb{E}_{s \sim \mu}[\text{ov}_L(s)] \geq c\varepsilon^4$. This average is over the stationary distribution μ , whereas the protocol starts its walks uniformly. In the

direct comparison below, we lose a factor of ε in the stated lower bound. Since every vertex has degree at most D , the uniform weight $1/n$ is at least $2m/(Dn)$ times the stationary weight $\mu(s)$ for every start,

$$\frac{2m}{Dn} \mu(s) = \frac{\deg_G(s)}{Dn} \leq \frac{1}{n},$$

and a far graph has $m \geq \varepsilon Dn$ edges, so $2m/(Dn) \geq 2\varepsilon$. Multiplying $\text{ov}_L(s) \geq 0$ by these weights and summing over s gives

$$\alpha_L = \frac{1}{n} \sum_{s \in [n]} \text{ov}_L(s) \geq \frac{2m}{Dn} \mathbb{E}_{s \sim \mu}[\text{ov}_L(s)] \geq 2\varepsilon \cdot c\varepsilon^4.$$

By Equation (1.1), every prover is accepted with probability at most $e^{-T\alpha_L}$, so we obtain constant soundness by choosing $T = \Theta(\varepsilon^{-5})$ with a sufficiently large constant. For constant D and ε , both query and communication costs are logarithmic in n .

2 Preliminaries

Definition 2.1 (Bounded-degree graph and oracle). A bounded-degree graph is a simple undirected graph $G = ([n], E)$ of maximum degree at most D , where $n \geq 2$ and $1 \leq D \leq n - 1$. We write $u \sim v$ when $uv \in E$. The neighbor oracle maps $(v, i) \in [n] \times [D]$ to the i -th neighbor of v in a fixed ordering when $i \leq \deg_G(v)$, and to $\perp$ otherwise. For two such graphs $G = ([n], E)$ and $H = ([n], E')$, we define

$$\Delta(G, H) := \frac{|E \setminus E'| + |E' \setminus E|}{Dn} = \frac{|E \Delta E'|}{Dn}.$$

Definition 2.2 (2-partitions and proximity). A 2-partition is a pair $\mathcal{P} = (V_0, V_1)$ such that $V_0 \cap V_1 = \emptyset$ and $V_0 \cup V_1 = [n]$. An edge is *violating* if its two vertices lie in the same part. The set of violating edges is

$$F_{\mathcal{P}} := \left(E \cap \binom{V_0}{2} \right) \cup \left(E \cap \binom{V_1}{2} \right).$$

A bipartition is a 2-partition with $F_{\mathcal{P}} = \emptyset$.

For fixed n and D , the class in Definition 2.1 is finite, and for any nonempty property Π of graphs in this class we define

$$\Delta(G, \Pi) := \min_{H \in \Pi} \Delta(G, H).$$

The minimum exists by finiteness. For $\eta \geq 0$, the graph G is η -close to Π if $\Delta(G, \Pi) \leq \eta$, and it is η -far from Π if $\Delta(G, \Pi) \geq \eta$.

We now define the model, following [15].

Definition 2.3 (One-round tolerant interactive proof of proximity). For fixed n, D , a nonempty property Π of bounded-degree graphs, and parameters $0 \leq \xi < \varepsilon$, a one-round (ξ, ε) -tolerant interactive proof of proximity for Π is an interactive proof system (P, V) in which one verifier message is followed by one prover response. The prover P receives the graph G in full, while the verifier V has access to its neighbor oracle. Both parties know n, D, Π , the tolerance ξ , and the proximity parameter ε . All probabilities are over the verifier's coins. The output $\mathsf{V}^{\mathsf{P}, G} \in \{0, 1\}$ satisfies the following conditions.

- *Completeness.* If $\Delta(G, \Pi) \leq \xi$, then $\mathbb{P}[\mathsf{V}^{\mathsf{P}, G} = 1] \geq 2/3$.

• *Soundness.* If $\Delta(G, \Pi) \geq \varepsilon$, then $\mathbb{P}[\mathbf{V}^{\mathbf{P}^*, G} = 1] \leq 1/3$ for every prover strategy $\mathbf{P}^*$. Completeness is perfect on Π when $\mathbb{P}[\mathbf{V}^{\mathbf{P}, G} = 1] = 1$ for every $G \in \Pi$.

A more general formulation allowing multiple rounds appears in [15], but we only need the one-round version here. The case $\xi = 0$ is an ordinary interactive proof of proximity.

We use Π for the property of bipartite graphs. Deleting the violating edges of a 2-partition makes the graph bipartite, and conversely, if a graph $H \in \Pi$ has bipartition $\mathcal{P}$, then obtaining H from G requires deleting every edge of $F_{\mathcal{P}}$. Hence

$$\Delta(G, \Pi) = \frac{1}{Dn} \min_{\mathcal{P}} |F_{\mathcal{P}}|. \quad (2.1)$$

Equivalently, G is ξ -close to bipartite when $\min_{\mathcal{P}} |F_{\mathcal{P}}| \leq \xi Dn$, and it is ε -far when $\min_{\mathcal{P}} |F_{\mathcal{P}}| \geq \varepsilon Dn$.

Definition 2.4 (Simple and lazy random walks). From a non-isolated vertex v , a simple-random-walk step moves to a uniformly random neighbor of v . The lazy random walk $(X_r)_{r \geq 0}$ on state space $[n]$ stays where it is with probability $1/2$ at each step from a non-isolated vertex and otherwise makes one simple-random-walk step, while isolated vertices are absorbing. Equivalently, its transition probabilities are

$$\mathbb{P}[X_{r+1} = u \mid X_r = v] := \begin{cases} 1, & u = v, \deg_G(v) = 0, \\ \frac{1}{2}, & u = v, \deg_G(v) > 0, \\ \frac{1}{2 \deg_G(v)}, & uv \in E, \\ 0, & u \neq v, uv \notin E. \end{cases}$$

A step with $X_{r+1} = X_r$ is a *holding step*, and a step with $X_{r+1} \neq X_r$ is an *edge traversal*.

A walk from a non-isolated start never reaches an isolated vertex, so each of its steps is a holding step with probability $1/2$, independently. The number of traversals among r lazy steps is therefore $\text{Bin}(r, 1/2)$, and given j traversals, the endpoint distribution is that of a j -step simple random walk.

For $m := |E| > 0$, we define the distribution

$$\mu(v) := \frac{\deg_G(v)}{2m}, \quad v \in [n].$$

The distribution μ is stationary for the lazy random walk and is the averaging measure in the overlap theorem below.

Definition 2.5 (Overlap). For the lazy random walk $(X_r)_{r \geq 0}$, let $\sigma_r \in \{+1, -1\}$ denote the parity of the number of edge traversals during the first r steps: $\sigma_r = +1$ for even parity and $\sigma_r = -1$ for odd parity. Equivalently, $\sigma_0 = +1$; after an edge traversal the sign is reversed, and after a holding step it is unchanged. For an integer $r \geq 0$, vertices $s, v \in [n]$, and $b \in \{\pm 1\}$, we define the joint probabilities

$$p_{r,s}^b(v) := \mathbb{P}[X_r = v, \sigma_r = b \mid X_0 = s] \quad (2.2)$$

and the overlap

$$\text{ov}_r(s) := \sum_{v \in [n]} \min\{p_{r,s}^{+1}(v), p_{r,s}^{-1}(v)\}. \quad (2.3)$$

Lemma 2.6 (Minimum guessing error). *Let ν be any distribution on $[n]$, and let $(X_r)_{r \geq 0}$ be the lazy random walk with $X_0 \sim \nu$. Then*

$$\mathbb{E}_{X_0 \sim \nu}[\text{ov}_r(X_0)] = \min_{a: [n] \times [n] \rightarrow \{\pm 1\}} \mathbb{P}[a(X_0, X_r) \neq \sigma_r].$$

Proof. We fix a function $a: [n] \times [n] \rightarrow \{\pm 1\}$. Conditioned on $X_0 = s$, the contribution of an endpoint v to its error is $p_{r,s}^{-1}(v)$ if $a(s, v) = +1$, and $p_{r,s}^{+1}(v)$ if $a(s, v) = -1$. Since a is a function of the pair (s, v) , its value may be chosen separately for every pair, so the minimum contribution for that pair is $\min\{p_{r,s}^{+1}(v), p_{r,s}^{-1}(v)\}$. We conclude by summing over v and averaging over $X_0 \sim \nu$. $\square$

We now present the overlap theorem, which is the main technical tool. We derive it in [Section 4](#) as a corollary of the analysis of the Goldreich–Ron tester by Fei and Rubinfeld [6].

Theorem 2.7 (Overlap theorem). *There are absolute constants $c, C > 0$ such that the following holds. Let $0 < \varepsilon < \frac{1}{2}$ and let $L \geq C\varepsilon^{-8} \log_2 n$ be an integer. If G is ε -far from bipartite, then*

$$\mathbb{E}_{s \sim \mu}[\text{ov}_L(s)] \geq c\varepsilon^4.$$

3 The protocol

Theorem 3.1 (Main theorem). *For every $n \geq 2$, $1 \leq D \leq n - 1$, and $0 < \varepsilon \leq 1$, bipartiteness has a one-round private-coin interactive proof of proximity with perfect completeness and soundness error at most $1/3$ on every graph that is ε -far from bipartite. Its query complexity is $O(\varepsilon^{-13}(1 + \log D) \log n)$, and its communication complexity is $O(\varepsilon^{-5} \log n)$. All queries are made before the verifier sends its message. The verifier runs in time polynomial in $1/\varepsilon$ and $\log n$, and on bipartite graphs the honest prover runs in time polynomial in n and $1/\varepsilon$.*

A variant of the RVW parity-challenge protocol (with the lazy walk of [Definition 2.4](#)).

Parameters. Let c and C be the constants of [Theorem 2.7](#), and set

$$\varepsilon_0 := \min\{\varepsilon, 1/4\}, \quad L := \lceil C\varepsilon_0^{-8} \log_2 n \rceil, \quad T := \left\lceil \frac{2}{c\varepsilon_0^5} \right\rceil.$$

The T walks are independent.

1. **Walks.** For each $k \in [T]$, the verifier samples $S_k \sim \text{Unif}([n])$ and runs a lazy random walk from S_k for exactly L steps. We denote the endpoint by V_k and write $b_k \in \{0, 1\}$ for the parity of the number of edge traversals, with $b_k = 0$ for even parity.
2. **Interaction.** The verifier sends $((S_k, V_k))_{k \in [T]}$. The prover returns bits $(a_k)_{k \in [T]}$.
3. **Decision.** The verifier accepts if and only if $a_k = b_k$ for every $k \in [T]$.

Replacing ε by $\varepsilon_0 = \min\{\varepsilon, 1/4\}$ loses nothing: a uniformly random 2-partition violates each edge with probability $1/2$, so every graph has a 2-partition with at most $m/2 \leq Dn/4$ violating edges and therefore has distance at most $1/4$ from bipartiteness. Hence no graph is ε -far for $\varepsilon > 1/4$.

Completeness. On a bipartite graph, the honest prover fixes a bipartition $\mathcal{P}$ and responds $a_k = 1$ when S_k and V_k lie in different parts, and $a_k = 0$ otherwise. At every edge traversal, the walk moves to the other part, so $a_k = b_k$ for every k and the verifier accepts with probability 1.

Soundness. The average in [Theorem 2.7](#) is over the stationary distribution, whereas the protocol's walks start uniformly, so we first pass from the average over μ to the uniform average using $\deg_G \leq D$, and then bound an arbitrary prover by the best guess at each challenge.

Lemma 3.2. *If G is ε -far from bipartite, then for every prover strategy the verifier accepts with probability at most $1/3$.*

Proof. Since $\varepsilon_0 \leq \varepsilon$, the graph is ε_0 -far from bipartite, so $m \geq \min_{\mathcal{P}} |F_{\mathcal{P}}| \geq \varepsilon_0 Dn > 0$, and using $\deg_G \leq D$ we obtain

$$\frac{1}{n} \sum_{s \in [n]} \text{ov}_L(s) \geq \frac{1}{nD} \sum_{s \in [n]} \deg_G(s) \text{ov}_L(s) = \frac{2m}{Dn} \mathbb{E}_{s \sim \mu}[\text{ov}_L(s)] \geq 2c\varepsilon_0^5.$$

For the final inequality, we used $2m/(Dn) \geq 2\varepsilon_0$ and the bound $\mathbb{E}_{s \sim \mu}[\text{ov}_L(s)] \geq c\varepsilon_0^4$ of [Theorem 2.7](#).

We fix an arbitrary prover strategy. For a possible value

$$M := ((s_k, v_k))_{k \in [T]}$$

of the verifier's message $((S_k, V_k))_{k \in [T]}$, we write $(a_k(M))_{k \in [T]}$ for the prover's responses to M .

The k -th walk has start s_k , endpoint v_k , and even parity with probability $\frac{1}{n} p_{L, s_k}^{+1}(v_k)$, and it has start s_k , endpoint v_k , and odd parity with probability $\frac{1}{n} p_{L, s_k}^{-1}(v_k)$. For either value of $a_k(M)$, the probability that the walk has this start and endpoint with $b_k = a_k(M)$ is therefore at most $\frac{1}{n} \max\{p_{L, s_k}^{+1}(v_k), p_{L, s_k}^{-1}(v_k)\}$. Since the walks are independent, the probability that the message is M and $a_k(M) = b_k$ for every k factors into the product of the corresponding one-walk probabilities. Therefore

$$\begin{aligned} \mathbb{P}[a_k = b_k \ \forall k \in [T]] &= \sum_{M \in ([n] \times [n])^T} \mathbb{P}[(S_k, V_k)_{k \in [T]} = M \text{ and } a_k(M) = b_k \ \forall k \in [T]] \\ &\leq \sum_{M \in ([n] \times [n])^T} \prod_{k=1}^T \frac{1}{n} \max\{p_{L, s_k}^{+1}(v_k), p_{L, s_k}^{-1}(v_k)\} \\ &= \left(\frac{1}{n} \sum_{s \in [n]} \sum_{v \in [n]} \max\{p_{L, s}^{+1}(v), p_{L, s}^{-1}(v)\} \right)^T \\ &= \left(1 - \frac{1}{n} \sum_{s \in [n]} \text{ov}_L(s) \right)^T \leq \exp(-2c\varepsilon_0^5 T) \leq e^{-4} < \frac{1}{3}. \end{aligned}$$

For the last equality, we used $\max\{x, y\} = x + y - \min\{x, y\}$ and $\sum_{v \in [n]} (p_{L, s}^{+1}(v) + p_{L, s}^{-1}(v)) = 1$. If the prover answers each challenge with the more likely parity, the verifier accepts with probability $(1 - \frac{1}{n} \sum_{s \in [n]} \text{ov}_L(s))^T$. The prover cannot do better by choosing all its answers together. $\square$

Complexities. The verifier makes queries only to run its walks, and the communication consists of two vertices from the verifier and one bit from the prover per challenge.

Claim 3.3. *The query complexity is $O(\varepsilon^{-13}(1 + \log D) \log n)$, and the communication complexity is $O(\varepsilon^{-5} \log n)$. The verifier runs in time polynomial in $1/\varepsilon$ and $\log n$, and on bipartite graphs the honest prover runs in time polynomial in n and $1/\varepsilon$.*

Proof. The query (v, i) returns a vertex for $i \leq \deg_G(v)$ and $\perp$ for larger i , so the verifier finds $\deg_G(v)$ by binary search over $[D]$ with $O(1 + \log D)$ queries at each lazy-walk step, and it makes one further query to sample a uniformly random neighbor when the walk traverses an edge. The verifier therefore makes

$$O(TL(1 + \log D)) = O(\varepsilon_0^{-13}(1 + \log D) \log n) = O(\varepsilon^{-13}(1 + \log D) \log n)$$

queries. The verifier sends $2T$ vertices and the prover sends T bits, so the communication complexity is $O(T \log n) = O(\varepsilon^{-5} \log n)$. The verifier performs TL walk steps, each in time polynomial in $\log n$, and compares T bits, so its running time is polynomial in $1/\varepsilon$ and $\log n$. The honest prover answers the T challenges according to a bipartition, in total time polynomial in n and $1/\varepsilon$. $\square$

Proof of Theorem 3.1. By Lemma 3.2 and Claim 3.3, the protocol has the stated soundness and costs. The protocol has one round, and all queries are made when the walks are run, before the verifier's message is sent. $\square$

Discussion. We retain RVW's parity challenge, but modify the walk and adjust its length and the number of repetitions. Our dependence on ε is worse than theirs: on a well-mixing graph RVW use $T = \Theta(1/\varepsilon)$ walks of length $\Theta(\log n)$, whereas our protocol uses $T = \Theta(\varepsilon^{-5})$ walks of length $\Theta(\varepsilon^{-8} \log n)$. We did not aim to optimize the dependence on ε .

Weak tolerance. Parnas, Ron, and Rubinfeld [14] observed that testers with uniformly distributed queries have a weak tolerance guarantee. For IPPs, a sufficient condition requires uniformity also when conditioned on prover messages. Here all queries precede interaction, and the honest prover uses a bipartition of a closest bipartite graph.

Lemma 3.4 (Weak tolerance). *Suppose that the honest prover answers according to a 2-partition minimizing $|F_{\mathcal{P}}|$. If G is bipartite, the prover is accepted with probability 1. If G is ξ -close to bipartite for*

$$\xi := \frac{1}{3TLD} = \Theta\left(\frac{\varepsilon^{13}}{D \log n}\right),$$

then the verifier accepts with probability at least $2/3$. Hence the same protocol is a (ξ, ε) -tolerant IPP.

Proof. Let $\mathcal{P}$ be the prover's 2-partition, so that the prover responds $a_k = 1$ when S_k and V_k lie in different parts of $\mathcal{P}$, and $a_k = 0$ otherwise. If G is bipartite, then $F_{\mathcal{P}} = \emptyset$, and if G is ξ -close to bipartite, then $|F_{\mathcal{P}}| \leq \xi Dn$ by Equation (2.1). When $F_{\mathcal{P}} = \emptyset$, the walk moves to the other part at every edge traversal, so $a_k = b_k$ for every k and the verifier accepts with probability 1. We may therefore assume that $F_{\mathcal{P}} \neq \emptyset$, and in particular that G has a non-isolated vertex.

For a single walk, we first bound the probability of traversing an edge of $F_{\mathcal{P}}$ at a given step. We write X_r for its vertex after r steps and define

$$h_r(v) := \frac{\mathbb{P}[X_r = v]}{\deg_G(v)}$$

for every non-isolated v . Since X_0 is uniform,

$$h_0(v) \leq \frac{1}{n}.$$

For every non-isolated v ,

$$h_{r+1}(v) = \frac{1}{2}h_r(v) + \frac{1}{2\deg_G(v)} \sum_{u \sim v} h_r(u) \leq \max_{u: \deg_G(u) > 0} h_r(u).$$

Hence $h_r(v) \leq 1/n$ for every $r \geq 0$, and

$$\mathbb{P}[\{X_r, X_{r+1}\} \in F_{\mathcal{P}}] = \frac{1}{2} \sum_{\{u,v\} \in F_{\mathcal{P}}} (h_r(u) + h_r(v)) \leq \frac{|F_{\mathcal{P}}|}{n}.$$

For each $k \in [T]$, we write $X_{k,r}$ for the vertex of the k -th walk after r steps. If the k -th walk traverses no edge of $F_{\mathcal{P}}$, it moves to the other part at every traversal, and $a_k = b_k$, so

$$\{a_k \neq b_k\} \subseteq \bigcup_{r=0}^{L-1} \{\{X_{k,r}, X_{k,r+1}\} \in F_{\mathcal{P}}\}.$$

Thus

$$\mathbb{P}[\exists k \in [T] : a_k \neq b_k] \leq \sum_{k=1}^T \sum_{r=0}^{L-1} \mathbb{P}[\{X_{k,r}, X_{k,r+1}\} \in F_{\mathcal{P}}] \leq \frac{TL|F_{\mathcal{P}}|}{n} \leq TLD\xi = \frac{1}{3}.$$

The verifier therefore accepts with probability at least $2/3$. $\square$

4 Proof of the overlap theorem

Fei and Rubinfeld [6] prove that the Goldreich–Ron tester rejects far graphs using $\Theta(\sqrt{n})$ walks of logarithmic length, and their analysis has two stages. Informally, the first stage is a birthday-paradox argument: among the tester’s $\Theta(\sqrt{n})$ walks, a collision between two walks of opposite parity is likely unless the even and odd endpoint distributions are far apart in total variation distance. In the second stage they show, in our language, that small overlap at logarithmic walk length is impossible on a far graph: the square roots of the even and odd endpoint distributions combine into a near-optimal solution of the basic semidefinite relaxation of Max-Cut, whose Goemans–Williamson rounding [7] gives a 2-partition with fewer violating edges than a far graph allows. Our theorem is about the overlap, so only the second stage is needed.

Notation. We use the notation of [Section 2](#) and introduce the following proof-specific notation.

- Let A be the adjacency matrix of G .
- Let $L^2(\mu)$ be the weighted space with $\langle f, g \rangle := \mathbb{E}_{v \sim \mu}[f(v)g(v)]$ and $\|f\|_1 := \mathbb{E}_{v \sim \mu}[|f(v)|]$.
- For a non-isolated vertex s , let $\delta_s(v) := \mathbf{1}[v = s]/\mu(s)$.
- Let I be the identity operator on $L^2(\mu)$, and let P be the simple-random-walk operator,

$$(Pf)(v) := \frac{1}{\deg_G(v)} \sum_{u \sim v} f(u)$$

for every non-isolated vertex v . The value of P on isolated vertices is immaterial, since $\mu(v) = 0$ there.

- Following [6], we set

$$W := \frac{I + P}{2}, \quad S := \frac{I - P}{2},$$

so W is the operator of the lazy random walk of [Definition 2.4](#). For each integer $r \geq 0$, we set

$$Q_r := W^r + S^r, \quad R_r := W^r - S^r.$$

- For a non-isolated vertex s and an integer $r \geq 0$, let

$$f_s^{(r)} := \sqrt{\frac{1}{2}Q_r\delta_s} - \sqrt{\frac{1}{2}R_r\delta_s},$$

where the square roots are taken coordinatewise.

- With $F_{\mathcal{P}}$ as in [Definition 2.2](#), the normalized Max-Cut value of G is

$$\text{MaxCut}(G) := \max_{\mathcal{P}} \frac{m - |F_{\mathcal{P}}|}{m} = 1 - \frac{Dn}{m} \Delta(G, \Pi), \quad (4.1)$$

where the equality follows from [Equation \(2.1\)](#).

- We denote by BASICSDP_A the semidefinite program

$$\max \left\{ \frac{1}{2} - \frac{\text{tr}(AX)}{4m} : X \succeq 0, X_{ii} \leq 1 \text{ for every } i \in [n] \right\}.$$

By taking square roots, we can combine the parity information from all starts into SDP vectors of norm at most one. Our aim is an average inner product close to -1 across edges, as for opposite signs in a bipartition.

Claim 4.1. *For every non-isolated vertex s , integer $r \geq 0$, and vertex $v \in [n]$,*

$$p_{r,s}^{+1}(v) = \frac{1}{2}(Q_r \delta_s)(v) \mu(v), \quad p_{r,s}^{-1}(v) = \frac{1}{2}(R_r \delta_s)(v) \mu(v). \quad (4.2)$$

In particular, $Q_r \delta_s$ and $R_r \delta_s$ are nonnegative at every non-isolated vertex.

Proof. The number J of traversals among the first r steps is $\text{Bin}(r, 1/2)$, and conditioned on $J = j$, the endpoint has the distribution of a j -step simple random walk from s ([Definition 2.4](#)), which assigns to v the probability $(P^j \delta_s)(v) \mu(v)$ [[6](#), Proposition 2.10], since $(P^j \delta_s)(v)$ is the j -step transition probability from v to s divided by $\mu(s)$, and $\mu(v) P^j(v, s) = \mu(s) P^j(s, v)$ for the simple random walk. By [Equation \(2.2\)](#),

$$p_{r,s}^{+1}(v) = \sum_{\substack{0 \leq j \leq r \\ j \text{ even}}} \mathbb{P}[J = j] \mathbb{P}[X_r = v \mid J = j] = \sum_{\substack{0 \leq j \leq r \\ j \text{ even}}} 2^{-r} \binom{r}{j} (P^j \delta_s)(v) \mu(v),$$

and since I and P commute, by the binomial theorem,

$$Q_r = 2^{-r} ((I + P)^r + (I - P)^r) = 2^{1-r} \sum_{\substack{0 \leq j \leq r \\ j \text{ even}}} \binom{r}{j} P^j.$$

From these expressions we obtain the identity for $p_{r,s}^{+1}$; for $p_{r,s}^{-1}$, we sum over odd j instead. Non-negativity follows, since $p_{r,s}^b(v) \geq 0$ for every $b \in \{\pm 1\}$ and $\mu(v) > 0$ at non-isolated v . $\square$

Claim 4.2. *For every non-isolated vertex s and integer $r \geq 0$,*

$$\text{ov}_r(s) = \frac{1}{4} (2 - \|Q_r \delta_s - R_r \delta_s\|_1). \quad (4.3)$$

Proof. Since $\sum_{v \in [n]} (p_{r,s}^{+1}(v) + p_{r,s}^{-1}(v)) = 1$, we obtain $\mathbb{E}_{v \sim \mu}[Q_r \delta_s + R_r \delta_s] = 2$ by summing [Equation \(4.2\)](#) over v . By [Equations \(2.3\)](#) and [\(4.2\)](#) and $2 \min\{a, b\} = a + b - |a - b|$,

$$\begin{aligned} \text{ov}_r(s) &= \sum_{v \in [n]} \min\{p_{r,s}^{+1}(v), p_{r,s}^{-1}(v)\} \\ &= \frac{1}{4} \mathbb{E}_{v \sim \mu}[Q_r \delta_s + R_r \delta_s - |Q_r \delta_s - R_r \delta_s|] \\ &= \frac{1}{4} (2 - \|Q_r \delta_s - R_r \delta_s\|_1). \end{aligned} \quad \square$$

We will apply the bounds of Fei and Rubinfeld at a length divisible by four, and to pass from it to the chosen length L , we first prove that the overlap is monotone in the length.

Lemma 4.3 (Monotonicity of the overlap). *For every $s \in [n]$ and all positive integers $r \leq r'$,*

$$\text{ov}_r(s) \leq \text{ov}_{r'}(s).$$

Proof. When s is isolated, both overlaps are zero, and otherwise $Q_r - R_r = 2S^r$, so by Equation (4.3), $\text{ov}_r(s) = \frac{1}{2} (1 - \|S^r \delta_s\|_1)$. For the simple-random-walk operator,

$$\begin{aligned} \|Pf\|_1 &= \sum_{\substack{v \in [n] \\ \deg_G(v) > 0}} \frac{\deg_G(v)}{2m} \left| \frac{1}{\deg_G(v)} \sum_{u \sim v} f(u) \right| \\ &\leq \frac{1}{2m} \sum_{\substack{v \in [n] \\ \deg_G(v) > 0}} \sum_{u \sim v} |f(u)| = \sum_{u \in [n]} \frac{\deg_G(u)}{2m} |f(u)| = \|f\|_1, \end{aligned}$$

where the terms for isolated vertices are zero. Consequently,

$$\|Sf\|_1 = \frac{1}{2} \|f - Pf\|_1 \leq \frac{1}{2} (\|f\|_1 + \|Pf\|_1) \leq \|f\|_1.$$

Thus $\|S^{r+1} \delta_s\|_1 \leq \|S^r \delta_s\|_1$ for every r , and $\text{ov}_{r+1}(s) \geq \text{ov}_r(s)$. $\square$

Statements from [6]. We restate the following three statements from [6].

Corollary 4.4 ([6, Corollary 3.6]). *The optimal value of BASICSDP_A is at least*

$$1 - \frac{1}{2} \cdot \inf_{k \in \mathbb{Z}, k > 0} \left\{ \mathbb{E}_{v \sim \mu} \left[\left\langle f_v^{(k)}, Pf_v^{(k)} \right\rangle + 1 \right] \right\}.$$

By [6, Proposition 3.5], the matrix $X_{ij} = \mathbb{E}_{v \sim \mu} [f_v^{(k)}(i) f_v^{(k)}(j)]$, with zero rows and columns at isolated vertices, satisfies the SDP constraints for every positive integer k . The optimal value of BASICSDP_A is therefore at least

$$\frac{1}{2} - \frac{\text{tr}(AX)}{4m} = 1 - \frac{1}{2} \mathbb{E}_{v \sim \mu} \left[\left\langle f_v^{(k)}, Pf_v^{(k)} \right\rangle + 1 \right].$$

Corollary 4.5 ([6, Corollary 3.10]). *Let t be a positive integer, let $v \in [n]$ be non-isolated, and let $\gamma := 2 - \|Q_{4t} \delta_v - R_{4t} \delta_v\|_1$. Then*

$$\frac{1}{t} \sum_{r=t}^{2t-1} \left(\left\langle f_v^{(2r)}, Pf_v^{(2r)} \right\rangle + 1 \right) \leq 2\sqrt{\gamma} + 2 \left(\frac{2}{t} \right)^{1/4} + \frac{1}{t} \log_2 \frac{1}{\mu(v)}. \quad (4.4)$$

Theorem 4.6 ([6, Theorem 2.2]). *For any $\beta \in [0, 1]$, if the optimal value of BASICSDP_A is at least $1 - \beta$, then the Max-Cut value of G is at least $1 - \sqrt{\beta}$.*

Proof of Theorem 2.7. We take $c := 1/288$ and $C := 2^{19}$, and assume that G is ε -far from bipartite. Then $m > 0$, and by Equation (4.1),

$$1 - \text{MaxCut}(G) = \frac{Dn}{m} \Delta(G, \Pi) \geq \frac{\varepsilon Dn}{m} \geq 2\varepsilon, \quad (4.5)$$

since $m \leq Dn/2$. Thus $\text{MaxCut}(G) \leq 1 - 2\varepsilon$, and the rest of the proof uses the assumption that G is far only through this bound. We may also assume that G has no isolated vertices, since deleting them changes none of m , $\text{MaxCut}(G)$, the optimal value of BASICSDP_A , and $\mathbb{E}_{s \sim \mu}[\text{ov}_r(s)]$ for every r , and after relabeling the remaining vertices as $[n]$, the assumed lower bound on L remains valid because n only decreases.

We set

$$\theta := \frac{\varepsilon^4}{144}, \quad t := \left\lfloor \frac{L}{4} \right\rfloor.$$

By the definitions of c and θ , $c\varepsilon^4 = \frac{\theta}{2}$ and $3\sqrt{\theta} = \frac{\varepsilon^2}{4}$. Since $L \geq C\varepsilon^{-8} \log_2 n \geq 8$, we have $t \geq L/8$, and since $\frac{C}{8} = 2^{16} \geq 2 \cdot 144^2$,

$$t \geq \frac{L}{8} \geq 2 \cdot 144^2 \varepsilon^{-8} \log_2 n = 2\theta^{-2} \log_2 n. \quad (4.6)$$

Suppose, toward a contradiction, that

$$\mathbb{E}_{s \sim \mu}[\text{ov}_L(s)] < \frac{\varepsilon^4}{288} = \frac{\theta}{2}.$$

For each $s \in [n]$, let

$$\gamma_s := 2 - \|Q_{4t}\delta_s - R_{4t}\delta_s\|_1.$$

Since $4t \leq L$, by [Lemma 4.3](#) and [Equation \(4.3\)](#),

$$\mathbb{E}_{s \sim \mu}[\gamma_s] = 4\mathbb{E}_{s \sim \mu}[\text{ov}_{4t}(s)] < 2\theta. \quad (4.7)$$

We apply [Corollary 4.5](#) with $v = s$ and take the expectation over $s \sim \mu$. By Cauchy–Schwarz and [Equation \(4.7\)](#), $\mathbb{E}_{s \sim \mu}[2\sqrt{\gamma_s}] \leq 2\sqrt{\mathbb{E}_{s \sim \mu}[\gamma_s]} \leq 2\sqrt{2\theta}$, and by Jensen’s inequality,

$$\mathbb{E}_{s \sim \mu} \left[\log_2 \frac{1}{\mu(s)} \right] \leq \log_2 \mathbb{E}_{s \sim \mu} \left[\frac{1}{\mu(s)} \right] = \log_2 n,$$

where the equality holds because no vertex is isolated. Hence

$$\frac{1}{t} \sum_{r=t}^{2t-1} \mathbb{E}_{s \sim \mu} \left[\left\langle f_s^{(2r)}, Pf_s^{(2r)} \right\rangle + 1 \right] \leq 2\sqrt{2\theta} + 2 \left(\frac{2}{t} \right)^{1/4} + \frac{\log_2 n}{t} \leq 2\sqrt{2\theta} + 2\sqrt{\theta} + \frac{\theta^2}{2} \leq 6\sqrt{\theta}.$$

The second inequality uses [Equation \(4.6\)](#) and the third uses $\theta \leq 1$. Hence some $r \in \{t, \dots, 2t-1\}$ satisfies

$$\mathbb{E}_{s \sim \mu} \left[\left\langle f_s^{(2r)}, Pf_s^{(2r)} \right\rangle + 1 \right] \leq 6\sqrt{\theta}.$$

Since $2r$ is a positive integer, we may take $k = 2r$ in [Corollary 4.4](#), so the optimal value of BASICSDP_A is at least

$$1 - \frac{1}{2} \inf_{k \in \mathbb{Z}, k > 0} \mathbb{E}_{s \sim \mu} \left[\left\langle f_s^{(k)}, Pf_s^{(k)} \right\rangle + 1 \right] \geq 1 - \frac{1}{2} \mathbb{E}_{s \sim \mu} \left[\left\langle f_s^{(2r)}, Pf_s^{(2r)} \right\rangle + 1 \right] \geq 1 - 3\sqrt{\theta} = 1 - \frac{\varepsilon^2}{4}.$$

By [Theorem 4.6](#) with $\beta = \varepsilon^2/4$, $\text{MaxCut}(G) \geq 1 - \frac{\varepsilon}{2}$, contradicting [Equation \(4.5\)](#). Therefore $\mathbb{E}_{s \sim \mu}[\text{ov}_L(s)] \geq \frac{\varepsilon^4}{288}$, as required. $\square$

References

- [1] N. Amir, O. Goldreich, and G. N. Rothblum. Doubly sub-linear interactive proofs of proximity. *Electronic Colloquium on Computational Complexity*, TR24-143, 2024. Full version, revision 1, 2025.
- [2] G. Arnon and G. N. Rothblum. On prover-efficient public-coin emulation of interactive proofs. In *2nd Conference on Information-Theoretic Cryptography*, volume 199 of *LIPICs*, pages 3:1–3:15, 2021.
- [3] I. Berman, R. D. Rothblum, and V. Vaikuntanathan. Zero-knowledge proofs of proximity. In *9th Innovations in Theoretical Computer Science Conference*, volume 94 of *LIPICs*, pages 19:1–19:20, 2018.
- [4] M. Dall’Agnol, T. Gur, S. Roy Moulik, and J. Thaler. Quantum proofs of proximity. *Quantum*, 6:834, 2022.
- [5] G. Davidoff, P. Sarnak, and A. Valette. *Elementary Number Theory, Group Theory, and Ramanujan Graphs*, volume 55 of *London Mathematical Society Student Texts*. Cambridge University Press, 2003.
- [6] Y. Fei and R. Rubinfeld. Testing bipartiteness in logarithmic rounds. arXiv:2606.13583, 2026. Theorem numbers refer to version 1.
- [7] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995.
- [8] O. Goldreich. Testing Bipartite in the Bounded-Degree Graph Model, Revisited (A digest of the paper of Fei and Rubinfeld (2026)). *Electronic Colloquium on Computational Complexity*, TR26-197, 2026. <https://eccc.weizmann.ac.il/report/2026/197/>.
- [9] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45(4):653–750, 1998.
- [10] O. Goldreich and D. Ron. A sublinear bipartiteness tester for bounded degree graphs. *Combinatorica*, 19(3):335–373, 1999.
- [11] O. Goldreich and D. Ron. Property testing in bounded degree graphs. *Algorithmica*, 32(2):302–343, 2002.
- [12] T. Gur and R. D. Rothblum. Non-interactive proofs of proximity. *Computational Complexity*, 27(1):99–207, 2018.
- [13] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.
- [14] M. Parnas, D. Ron, and R. Rubinfeld. Tolerant property testing and distance approximation. *Journal of Computer and System Sciences*, 72(6):1012–1042, 2006.
- [15] G. N. Rothblum, S. P. Vadhan, and A. Wigderson. Interactive proofs of proximity: delegating computation in sublinear time. In *Proceedings of the 45th Annual ACM Symposium on*

Theory of Computing, pages 793–802, 2013. Theorem and lemma numbers refer to the full version, <https://guyrothblum.wordpress.com/wp-content/uploads/2014/11/rvw13.pdf>, accessed 2026.

- [16] R. Rubinfeld and M. Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.
- [17] Y. Yoshida. Tolerant testing for Unique Games. arXiv:2605.17760, 2026. Theorem numbers refer to version 1.

A Optimality for parity-challenge protocols

In a *parity-challenge protocol*, the verifier samples a finite, arbitrarily correlated tuple of walks $\Gamma_1, \dots, \Gamma_T$ in G , where T may depend on its coins, and sends one message revealing their starts and endpoints. The prover returns one bit per walk, chosen jointly after seeing the whole message, and the verifier rejects only if some bit differs from the traversal parity of its walk. We write N_i for the number of edge traversals of Γ_i and define the protocol’s cost as $\mathbb{E}[\sum_{i=1}^T N_i]$. The protocol of Section 3 belongs to this class.

Theorem A.1 (Logarithmic lower bound for parity-challenge protocols). *There are absolute constants $\delta, \gamma > 0$ and infinitely many n with the following property. There is a simple connected 6-regular n -vertex graph G , in the bounded-degree model with degree bound $D = 6$, that is δ -far from bipartite and on which every parity-challenge protocol has a prover that is rejected with probability at most $\mathbb{E}[\sum_{i=1}^T N_i]/(\gamma \log n)$. Consequently, any family of parity-challenge protocols that rejects every prover with probability at least some fixed $\kappa > 0$ on every δ -far graph has expected cost $\Omega(\log n)$ on these graphs.*

Proof. We fix a sufficiently large prime $\ell \equiv 1 \pmod{20}$ and take $G = X^{5,\ell}$, the Lubotzky–Phillips–Sarnak graph [13]. The parameter 5 makes the graph 6-regular, and $\ell \equiv 1 \pmod{20}$ gives $\ell \equiv 1 \pmod{4}$ and makes 5 a quadratic residue modulo ℓ , the case in which $X^{5,\ell}$ is simple, connected, and non-bipartite, with $n = \ell(\ell^2 - 1)/2$ vertices and girth $g \geq 2 \log_5 \ell$ [5, Theorem 4.2.2 and Corollary 4.3.6]. By Dirichlet’s theorem, there are infinitely many such primes.

For a 2-partition $\mathcal{P} = (V_0, V_1)$ with violating edges $F_{\mathcal{P}}$ as in Definition 2.2, we define its sign vector $\chi \in \{-1, 1\}^n$ by $\chi_v = 1$ for $v \in V_0$ and $\chi_v = -1$ for $v \in V_1$. We write A for the adjacency matrix of G . Since G is a Ramanujan graph, every eigenvalue of A other than ± 6 has absolute value at most $2\sqrt{5}$ [13], and -6 is not an eigenvalue because G is connected and non-bipartite, so $\lambda_{\min}(A) \geq -2\sqrt{5}$. In $\chi^T A \chi$ each edge appears in both orientations, with value 2 for a violating edge and -2 otherwise, so since G has $3n$ edges, the Rayleigh bound gives

$$-2\sqrt{5}n \leq \chi^T A \chi = 4|F_{\mathcal{P}}| - 6n, \quad |F_{\mathcal{P}}| \geq \frac{6 - 2\sqrt{5}}{4}n.$$

Thus G is δ -far from bipartite for $\delta = (6 - 2\sqrt{5})/24$, using degree bound $D = 6$.

We set $\gamma := 1/(4 \log 5)$, chosen so that $2\gamma \log n < g$: since $n < \ell^3$,

$$2\gamma \log n < \frac{3}{2} \log_5 \ell < 2 \log_5 \ell \leq g.$$

If two walks with the same start and endpoint each make fewer than $\gamma \log n$ edge traversals but have different parity, then deleting their holding steps and concatenating one with the reverse of

the other gives an odd closed walk shorter than g , which contains an odd cycle no longer than itself, contradicting the definition of girth. Thus any two such walks have the same parity.

The prover returns the parity of a shortest path for each revealed pair. If $N_i < \gamma \log n$, deleting the holding steps from Γ_i gives an N_i -edge walk, while a shortest path with the same start and endpoint has length at most N_i . Both walks then make fewer than $\gamma \log n$ traversals, so they have the same parity, and the response for Γ_i is correct. Rejection therefore implies $\sum_i N_i \geq \gamma \log n$, and by Markov's inequality we bound the rejection probability by $\mathbb{E}[\sum_{i=1}^T N_i]/(\gamma \log n)$. $\square$

B Yoshida's multiscale alternative

Yoshida [17] proves a related multiscale theorem for the overlap of a walk stopped at an independent geometric time, which we restate in our notation in order to compare the two overlaps.

B.1 Geometric overlap

For $m > 0$, we define $\tau(G) := m^{-1} \min_{\mathcal{P}} |F_{\mathcal{P}}|$, with $\tau(G) = 0$ when $m = 0$. By Equation (4.1) and $2m \leq Dn$, every graph that is ε -far from bipartite has $\tau(G) \geq 2\varepsilon$. For $L > 0$, we use a geometric time T_L with mean L , independent of the walk:

$$\mathbb{P}[T_L = t] = \frac{1}{L+1} \left(\frac{L}{L+1} \right)^t, \quad t = 0, 1, 2, \dots$$

The overlap at the stopping time T_L is

$$\begin{aligned} \widetilde{\text{ov}}_L(s) &:= \sum_{v \in [n]} \min_{\sigma = \pm 1} \mathbb{P}[X_{T_L} = v, \sigma_{T_L} = \sigma \mid X_0 = s] \\ &= \sum_{v \in [n]} \min \left\{ \sum_{t \geq 0} \mathbb{P}[T_L = t] p_{t,s}^{+1}(v), \sum_{t \geq 0} \mathbb{P}[T_L = t] p_{t,s}^{-1}(v) \right\}. \end{aligned}$$

As in Lemma 2.6, this is the minimum error in guessing the parity from the start and endpoint when the stopping time is not revealed.

Theorem B.1 (Yoshida [17, Theorem 4.9]). *There are absolute constants $0 < \alpha', c', c_\lambda \leq 1$ and $C' \geq 1$ such that the following holds. For $0 < \rho < 1/10$, we define*

$$\lambda := c_\lambda \frac{\rho}{1 + \log(1/\rho)}, \quad \mathcal{Q}_\lambda := \{2^j \lambda : j \in \mathbb{Z}_{\geq 0}, 2^j \lambda \leq 1\} \cup \{1\}.$$

If $\tau(G) \geq \rho$, then there is a scale $q \in \mathcal{Q}_\lambda$ such that, for $L_q := C' q \lambda^{-2} \log n$,

$$\mathbb{P}_{s \sim \mu}[\widetilde{\text{ov}}_{L_q}(s) \geq \alpha'] \geq c' q.$$

It is useful to compare the two notions of overlap. For fixed t , $\text{ov}_t(s)$ is the minimum error in guessing σ_t from (s, X_t) , whereas $\widetilde{\text{ov}}_L(s)$ is the minimum error in guessing σ_{T_L} from (s, X_{T_L}) when T_L is random and not revealed. By Jensen's inequality, for every s ,

$$\widetilde{\text{ov}}_L(s) \geq \sum_{t \geq 0} \mathbb{P}[T_L = t] \text{ov}_t(s).$$

The inequality can be strict: for a given endpoint, the most likely parity may depend on the stopping time. Therefore, we cannot in general deduce a fixed-time lower bound from Theorem B.1

using this comparison. Additionally, in [Theorem 2.7](#), we bound $\mathbb{E}_{s \sim \mu}[\text{ov}_L(s)]$ from below for *every* integer $L \geq C\varepsilon^{-8} \log_2 n$. In [Theorem B.1](#), the lower bound instead holds for *some* $q \in \mathcal{Q}_\lambda$, which may depend on G . These changes between the two notions force us to change the protocol, as we do next.

B.2 The bipartiteness protocol

Theorem B.2. *For every $n \geq 2$, $1 \leq D \leq n - 1$, and $0 < \varepsilon \leq 1$, bipartiteness has a one-round private-coin interactive proof of proximity with perfect completeness and soundness error at most $1/3$ on every graph that is ε -far from bipartite. Its query complexity is $O(\varepsilon^{-3} \log^3(1/\varepsilon)(1 + \log D) \log n)$, and its communication complexity is $O(\varepsilon^{-2} \log(1/\varepsilon) \log n)$. All graph queries are made before the verifier sends its message.*

Bipartiteness protocol with geometric stopping.

Parameters. Let λ , $\mathcal{Q}_\lambda$, and L_q be as in [Theorem B.1](#), with

$$\rho := \min\{\varepsilon, 1/20\}.$$

For each $q \in \mathcal{Q}_\lambda$, we set

$$M_q := \left\lceil \frac{2}{\alpha' c' \rho q} \right\rceil, \quad W := \sum_{q \in \mathcal{Q}_\lambda} M_q L_q, \quad B := 12W.$$

The walks are independent.

1. **Walks.** For each $q \in \mathcal{Q}_\lambda$ and $k \in [M_q]$, the verifier samples $S_{q,k} \sim \text{Unif}([n])$ and runs a lazy random walk from $S_{q,k}$ until an independent geometric time of mean L_q . We denote the endpoint by $V_{q,k}$ and write $b_{q,k} \in \{0, 1\}$ for the parity of the number of edge traversals, with $b_{q,k} = 0$ for even parity. If the verifier cannot complete all the walks within B lazy steps, it accepts.
 2. **Interaction.** The verifier sends $((S_{q,k}, V_{q,k}))_{q \in \mathcal{Q}_\lambda, k \in [M_q]}$. The prover returns bits $(a_{q,k})_{q,k}$.
 3. **Decision.** The verifier accepts if and only if $a_{q,k} = b_{q,k}$ for every $q \in \mathcal{Q}_\lambda$ and $k \in [M_q]$.
-

Proof of [Theorem B.2](#). Completeness. If G is bipartite, the honest prover answers according to a bipartition. The traversal parity is even exactly when the start and endpoint lie in the same part, so every response is correct. The verifier also accepts if the cutoff is exceeded, so completeness is perfect.

Soundness. If G is ε -far from bipartite, then $m \geq \min_{\mathcal{P}} |F_{\mathcal{P}}| \geq \rho D n$ and $\tau(G) \geq \rho D n / m \geq 2\rho$. By [Theorem B.1](#), some $q \in \mathcal{Q}_\lambda$ satisfies $\mathbb{E}_{s \sim \mu}[\widetilde{\text{ov}}_{L_q}(s)] \geq \alpha' c' q$. To pass to the uniform average, we use $\mu(s) \leq D/(2m)$:

$$\frac{1}{n} \sum_s \widetilde{\text{ov}}_{L_q}(s) \geq \frac{2m}{Dn} \mathbb{E}_{s \sim \mu}[\widetilde{\text{ov}}_{L_q}(s)] \geq 2\alpha' c' \rho q.$$

We first consider the protocol without the cutoff. For each start and endpoint, the prover's best response is the more likely parity, as in [Lemma 3.2](#). Since the challenges are independent, for every prover,

$$\mathbb{P}[\text{accept without the cutoff}] \leq \prod_{q' \in \mathcal{Q}_\lambda} \left(1 - \frac{1}{n} \sum_s \widetilde{\text{ov}}_{L_{q'}}(s) \right)^{M_{q'}} \leq (1 - 2\alpha' c' \rho q)^{M_q} \leq e^{-4}.$$

The total number N of lazy steps without the cutoff satisfies $\mathbb{E}[N] = W$, so

$$\mathbb{P}[N > B] \leq \frac{W}{B} \leq \frac{1}{12}.$$

If $N \leq B$, the verifier accepts only if all the prover's answers are correct. Thus, for every prover,

$$\mathbb{P}[\text{accept}] \leq \mathbb{P}[\text{accept without the cutoff}] + \mathbb{P}[N > B] \leq e^{-4} + \frac{1}{12} < \frac{1}{3}. \quad (\text{B.1})$$

Complexities. The verifier takes at most $B = O(W)$ lazy steps, with $O(1 + \log D)$ queries per step by [Claim 3.3](#). Without the cutoff, the expected number of steps for each scale is $M_q L_q = \Theta(\log n / (\rho \lambda^2))$, uniformly over q . Since $\rho = \Theta(\varepsilon)$, $\lambda = \Theta(\varepsilon / \log(1/\varepsilon))$, and $|\mathcal{Q}_\lambda| = \Theta(1 + \log(1/\lambda)) = \Theta(\log(1/\varepsilon))$,

$$W = \sum_{q \in \mathcal{Q}_\lambda} M_q L_q = \Theta\left(\frac{|\mathcal{Q}_\lambda| \log n}{\rho \lambda^2}\right) = \Theta(\varepsilon^{-3} \log^3(1/\varepsilon) \log n).$$

The query complexity is therefore

$$O(W(1 + \log D)) = O(\varepsilon^{-3} \log^3(1/\varepsilon)(1 + \log D) \log n).$$

For each challenge, the verifier sends two vertices and the prover returns one bit. Since $\sum_{q \in \mathcal{Q}_\lambda} 1/q = \Theta(1/\lambda)$ and $\sum_{q \in \mathcal{Q}_\lambda} M_q = \Theta(1/(\rho \lambda))$, the communication complexity is

$$O\left(\left(\sum_{q \in \mathcal{Q}_\lambda} M_q\right) \log n\right) = O\left(\frac{\log n}{\rho \lambda}\right) = O(\varepsilon^{-2} \log(1/\varepsilon) \log n).$$

□

C An extension to Unique Games

Yoshida [17] uses the same general approach for bipartiteness and Unique Games. We extend the preceding protocol to Unique Games by asking the prover for one starting label and its images along several independent walks from the same vertex.

C.1 Constraints and geometric ambiguity

Definition C.1 (Unique Games instance). We fix an alphabet $[Q]$, where $Q \geq 2$. A Unique Games instance is $U = (G, \{\pi_{uv}\}_{uv \in E})$, where G is as in [Definition 2.1](#) and each oriented edge uv has an associated permutation π_{uv} of $[Q]$, with $\pi_{vu} = \pi_{uv}^{-1}$. For a query (v, i) , the verifier receives the i -th neighbor u together with π_{vu} , or $\perp$ if no such neighbor exists.

Definition C.2 (Labelings and proximity). For a labeling $x: [n] \rightarrow [Q]$, the set of violated constraints is $F_x := \{uv \in E : x(v) \neq \pi_{uv}(x(u))\}$. Writing Π_{UG} for satisfiability, we use the distance

$$\Delta(U, \Pi_{\text{UG}}) := \frac{1}{Dn} \min_x |F_x|.$$

This is the minimum number of constraint deletions needed for satisfiability, normalized by Dn . As for bipartiteness, the distance is the same if insertions are allowed. For $m > 0$, we write $\tau_{\text{UG}}(U) := m^{-1} \min_x |F_x|$, with $\tau_{\text{UG}}(U) = 0$ when $m = 0$. An instance is ε -far from satisfiable when $\Delta(U, \Pi_{\text{UG}}) \geq \varepsilon$.

We require completeness on satisfiable instances and soundness on instances that are ε -far from satisfiable, as in [Definition 2.3](#).

Definition C.3 (Geometric ambiguity). For a lazy walk $P = (u_0, \dots, u_t)$, the induced permutation is

$$\Pi_P := \pi_{u_{t-1}u_t} \circ \dots \circ \pi_{u_0u_1},$$

where the permutation is the identity for a holding step or the empty path. For $P = (X_0, \dots, X_{T_L})$, we define

$$p_{L,s}^{a,b}(v) := \mathbb{P}[X_{T_L} = v, \Pi_P(a) = b \mid X_0 = s].$$

The endpoint marginal $\sum_b p_{L,s}^{a,b}(v)$ is independent of a . Yoshida's geometric ambiguity, denoted here by $\text{amb}_L(s)$, is

$$\text{amb}_L(s) := \min_{a \in [Q]} \sum_{v \in [n]} \left(\sum_{b \in [Q]} p_{L,s}^{a,b}(v) - \max_{b \in [Q]} p_{L,s}^{a,b}(v) \right) = 1 - \max_{a \in [Q]} \sum_{v \in [n]} \max_{b \in [Q]} p_{L,s}^{a,b}(v). \quad (\text{C.1})$$

For a fixed starting label a , the prover's best guess at endpoint v is a label b maximizing $p_{L,s}^{a,b}(v)$. The stationary distribution is $\mu(v) = \deg_G(v)/(2m)$.

When $Q = 2$ and adjacent vertices must have different labels (the bipartiteness case), $\text{amb}_L(s) = \widetilde{\text{ov}}_L(s)$.

Lemma C.4 (Yoshida [17, Lemma 3.12]). *For each fixed $Q \geq 2$, there are constants $0 < c_Q \leq 1$ and $C_Q \geq 1$ with the following property. For $0 < \rho < 1/10$, we define*

$$\mathcal{Q}_\rho := \{2^j \rho^2 : j \in \mathbb{Z}_{\geq 0}, j < \lceil \log_2(\rho^{-2}) \rceil\} \cup \{1\}.$$

For each $r \in \mathcal{Q}_\rho$, let

$$\alpha_r := c_Q \rho r^{-1/2}, \quad L_r := C_Q \rho^{-3} r^{1/2} \log n.$$

If $\tau_{\text{UG}}(U) \geq \rho$, then some $r \in \mathcal{Q}_\rho$ satisfies

$$\mathbb{P}_{s \sim \mu}[\text{amb}_{L_r}(s) \geq \alpha_r] \geq c_Q r.$$

C.2 The protocol

Theorem C.5. *For each fixed $Q \geq 2$, every $n \geq 2$, $1 \leq D \leq n - 1$, and $0 < \varepsilon \leq 1$, Unique Games satisfiability has a one-round private-coin interactive proof of proximity with perfect completeness and soundness error at most $1/3$ on every instance that is ε -far from satisfiable. Its query complexity is $O_Q(\varepsilon^{-5} \log(1/\varepsilon)(1 + \log D) \log n)$, and its communication complexity is $O_Q(\varepsilon^{-3} \log n)$. All graph and constraint queries are made before the verifier sends its message.*

Unique Games protocol with geometric stopping.

Parameters. Let $\mathcal{Q}_\rho$, α_r , and L_r be as in [Lemma C.4](#), with

$$\rho := \min\{\varepsilon, 1/20\}.$$

For each $r \in \mathcal{Q}_\rho$, we set

$$M_r := \left\lceil \frac{4}{c_Q \rho r} \right\rceil, \quad K_r := \left\lceil \frac{\log(2Q)}{\alpha_r} \right\rceil.$$

The cutoff is B , where

$$W := \sum_{r \in \mathcal{Q}_\rho} M_r K_r L_r, \quad B := 12W.$$

The starts are independent, and the walks are independent conditional on their starts.

1. **Walks.** For each $r \in \mathcal{Q}_\rho$ and $k \in [M_r]$, the verifier samples $S_{r,k} \sim \text{Unif}([n])$ and runs K_r lazy random walks from $S_{r,k}$, each stopped at an independent geometric time of mean L_r . We denote the endpoint of the j -th walk by $V_{r,k,j}$ and write $\Pi_{P_{r,k,j}}$ for its induced permutation. If the verifier cannot complete all the walks within B lazy steps, it accepts.
2. **Interaction.** The verifier sends

$$((S_{r,k}, (V_{r,k,j})_{j \in [K_r]}))_{r \in \mathcal{Q}_\rho, k \in [M_r]}.$$

The prover returns one starting label $a_{r,k} \in [Q]$ and endpoint labels $(b_{r,k,j})_{j \in [K_r]} \in [Q]^{K_r}$ for each (r, k) .

3. **Decision.** The verifier accepts if and only if

$$b_{r,k,j} = \Pi_{P_{r,k,j}}(a_{r,k}) \quad \text{for all } r, k, j.$$

Proof of Theorem C.5. Completeness. On a satisfiable instance, the prover fixes a satisfying labeling x and responds with $a_{r,k} = x(S_{r,k})$ and $b_{r,k,j} = x(V_{r,k,j})$. Every constraint is satisfied, so $\Pi_{P_{r,k,j}}(x(S_{r,k})) = x(V_{r,k,j})$ for every walk. All responses are therefore correct. The verifier also accepts if the cutoff is exceeded, so completeness is perfect.

Soundness. For a fixed scale r and start s , we consider one batch in the protocol without the cutoff and write $A_r(s)$ for its largest acceptance probability. For each possible list of endpoints $v_1, \dots, v_{K_r}$, the prover chooses one common starting label a and the most likely image of a at each endpoint. Since the walks are independent conditional on s ,

$$A_r(s) = \sum_{v_1, \dots, v_{K_r} \in [n]} \max_{a \in [Q]} \prod_{j=1}^{K_r} \max_{b \in [Q]} p_{L_r, s}^{a,b}(v_j) \leq \sum_{a \in [Q]} \left(\sum_{v \in [n]} \max_{b \in [Q]} p_{L_r, s}^{a,b}(v) \right)^{K_r} \leq Q(1 - \text{amb}_{L_r}(s))^{K_r}.$$

For the first inequality, we bound the maximum over a by the sum over the Q possible starting labels. With $K_r \geq \log(2Q)/\alpha_r$,

$$\text{amb}_{L_r}(s) \geq \alpha_r \implies A_r(s) \leq Qe^{-\alpha_r K_r} \leq \frac{1}{2}. \quad (\text{C.2})$$

If U is ε -far, then $m \geq \min_x |F_x| \geq \rho Dn$ and $\tau_{\text{UG}}(U) \geq 2\rho$. By Lemma C.4, there is a scale $r \in \mathcal{Q}_\rho$ for which the set

$$H_r := \{s \in [n] : \text{amb}_{L_r}(s) \geq \alpha_r\}$$

has $\mu(H_r) \geq c_Q \rho r$. Since $\deg_G \leq D$,

$$\frac{|H_r|}{n} \geq \frac{2m}{Dn} \mu(H_r) \geq 2c_Q \rho r, \quad \frac{1}{n} \sum_s A_r(s) \leq 1 - \frac{|H_r|}{2n} \leq 1 - c_Q \rho r.$$

Since the hidden permutations of different batches are independent conditional on the full message, the prover maximizes its probability of answering all batches correctly by choosing the response to each batch separately. Thus, for every prover,

$$\mathbb{P}[\text{accept without the cutoff}] \leq \prod_{r' \in \mathcal{Q}_\rho} \left(\frac{1}{n} \sum_s A_{r'}(s) \right)^{M_{r'}} \leq (1 - c_Q \rho r)^{M_r} \leq e^{-4}.$$

Without the cutoff, the expected total number of lazy steps is W . As in Equation (B.1), the probability of exceeding B is at most $1/12$, so

$$\mathbb{P}[\text{accept}] \leq e^{-4} + \frac{1}{12} < \frac{1}{3}.$$

Complexities. The verifier takes at most $B = 12W$ lazy steps, with $O(1 + \log D)$ queries per step. For fixed Q ,

$$M_r = \Theta_Q(\rho^{-1}r^{-1}), \quad K_r = \Theta_Q(\rho^{-1}r^{1/2}), \quad L_r = \Theta_Q(\rho^{-3}r^{1/2} \log n).$$

For every scale r , the expected number of steps without the cutoff is $M_r K_r L_r = \Theta_Q(\rho^{-5} \log n)$. There are $|\mathcal{Q}_\rho| = \Theta(\log(1/\rho))$ scales, so

$$W = \sum_{r \in \mathcal{Q}_\rho} M_r K_r L_r = \Theta_Q(\rho^{-5} |\mathcal{Q}_\rho| \log n) = \Theta_Q(\varepsilon^{-5} \log(1/\varepsilon) \log n).$$

The query complexity is therefore

$$O(W(1 + \log D)) = O_Q(\varepsilon^{-5} \log(1/\varepsilon)(1 + \log D) \log n).$$

For each batch, the verifier sends the starting vertex and K_r endpoints, and the prover returns $K_r + 1$ labels. Since $\sum_{r \in \mathcal{Q}_\rho} r^{-1/2} = \Theta(\rho^{-1})$ and $\sum_{r \in \mathcal{Q}_\rho} M_r K_r = \Theta_Q(\rho^{-3})$, the communication complexity in bits is

$$O\left(\sum_{r \in \mathcal{Q}_\rho} M_r (K_r + 1)(\log n + \log Q)\right) = O_Q(\varepsilon^{-3} \log n).$$

□