Electronic Colloquium on Computational Complexity, Report No. 129 (2025)

Expansion without Connectivity: A Property Testing Perspective

Irit Dinur Oded Goldreich

September 7, 2025

Abstract

We consider the query complexity of testing whether a bounded-degree graph is expanding, regardless of whether or not it is connected.

Whereas prior work studied testing the property of being an expander (equiv., testing the set of expander graphs), here we study testing the set of graphs that consist of connected components that are each an expander. Within the context of simplicial complexes, such graphs are called coboundary expanders. Loosely speaking, we show that testing expansion of n-vertex graphs requires $\Omega(n^{1/2})$ queries and can be done using $O(n^{0.667})$ queries. Recall that testing whether a n-vertex graph is an expander has query complexity $\widetilde{\Theta}(n^{1/2})$.

Our upper bound combines a distribution tester for generalized uniformity (i.e., uniformity over an unspecified subset of the domain) with a tester for expander graphs.

We also consider the problem of testing the girth of a graph. We prove that, as long as the girth of the (not necessarily connected) graph is at most logarithmic in its size, the query complexity must be exponential in the girth. This matches the upper bound established by the straightforward tester.

Keywords: Property Testing, Graph Properties, Bounded-Degree Graphs, Expander Graphs, Coboundary Expansion, Girth, Distribution Testing

ISSN 1433-8092

Contents

1 Introduction

The study of testing graph properties is one of the main subareas within property testing (see textbook [4]). Within this sub-area, the bounded-degree graph model (introduced by Goldreich and Ron [6] and reviewed in [4, Chap. 9]) is one of the two main models (the other being the dense graph model, introduced by Goldreich, Goldwasser and Ron [5] and reviewed in [4, Chap. 8]).

Loosely speaking, in the bounded-degree graph model, for a fixed degree bound $d \in \mathbb{N}$, we represent n-vertex graphs by an incidence function of the form $g : [n] \times [d] \to [n] \cup \{\perp\}$ such that $g(v,i)$ is the i^{th} neighbor of v (and $g(v,i) = \perp$ if v has less than i neighbors). Distance between n-vertex graphs is defined as the fraction of entries on which their (mutually closest) incidence functions differ. The tester (for a fixed graph property) is given oracle access to such an incidence function and has to distinguish between the case that the graph has the property and the case that it is ϵ-far from the property, where $\epsilon > 0$ is called the proximity parameter.

A detailed definition of the model is presented in Section 2. At this point we stress that d will always denote the constant degree bound of this model, and so all graphs we considered will have maximal degree at most d. In contrast (to this constant d), the number of vertices, denoted n, is the main varying parameter. Another varying parameter is the proximity parameter ϵ; actually, in lower bounds we shall set it to some small positive constant, whereas in upper bounds we shall state its effect on the complexity.

The main focus of this paper is on the query complexity of testing expansion. We also consider the query complexity of testing girth. Both these properties are very natural graph properties, which we consider independent of the connectivity property.

1.1 Expansion without connectivity

The common and naive notion of expansion views it as a property of connected graphs. For example, a connected graph has edge expansion c if the number of edges that cross a cut over the size of the smaller vertex set (in the cut) is at least c. We say that such a graph is an expander. Clearly, an unconnected graph cannot be an expander, since no edge goes out of each of its connected component (i.e., the corresponding cut has no edges). Yet, intuitively, a graph that consists of connected components that are each an expander (w.r.t., the same constant c) is "expanding".

Hence, *we distinguish between being an expander and having the property of expansion.* The property of expansion postulates that each of the connected components of the graph is an expander (with the same expansion parameter), whereas the property of being an expander augments this requirement by insisting that the graph itself is connected. The latter augmentation is not inherent to the expansion property.

Definition 1.1 (the expansion property): *For a constant $c > 0$, we let $\Xi_c(n)$ denote the set of n-vertex graphs that consist of connected components that are each expanding in the sense of having* edge expansion *at least c. Let $\Xi_c = \{\Xi_c(n)\}_{n\in\mathbb{N}}$.*

We stress that the graphs in Ξ_c (as well as those in Γ_g below) are not necessarily connected. We mention that expansion without connectivity arises naturally in the context of high dimensional expansion of the homological kind; namely, coboundary expansion [14, 10]. Indeed, Ξ_c equals the set of all *c-coboundary expanders*.[1] Expansion without connectivity also arises naturally in a striking

[1] The coboundary expansion of a graph may be defined by viewing it as a system of 2-variable equalities, where

result of Oppenheim [15] that states that if all links of a graph G (i.e., the subgraphs of a graph G that are induced by the neighbors of each vertex of G) are sufficiently strong expanders, then G itself is expanding (albeit not necessarily connected).

One of the graph properties studied by Goldreich and Ron, when introducing the bounded-degree graph model, is the property of being an expander [6, Sec. 7.2]. They showed that testing whether an n-vertex graph is an expander requires $\Omega(n^{1/2})$ queries. Their proof is pivoted on the fact that an n-vertex graph consisting of two $n/2$-vertex expanders is far from being an n-vertex expander. Arguably, this is a violation of the connectivity condition and not of the intuitive notion of graph expansion, which may be applied meaningfully to each connected component.

Hence, rather than testing whether a graph is an expander, one may wish to test whether a graph consists of connected components that are each an expander (i.e., the foregoing property Ξ_c). We first prove that the $\Omega(n^{0.5})$ query lower bound extends also to this case.

Theorem 1.2 (a lower bound on the query complexity of testing expansion): *For every $d \geq 3$, there exists $c > 0$ such that testing $\Xi_c(n)$ requires $\Omega(\sqrt{n})$ queries.*

As usual, our proof uses the indistinguishability method [4, Sec. 7.2]; that is, we present two distributions of d-regular n-vertex graphs such that, on the one hand, these two distributions cannot be distinguished by a $o(\sqrt{n})$-query algorithm, while, on the other hand, a tester for $\Xi_c(n)$ must distinguish them (because the first distribution is concentrated on expanders, whereas the second distribution is far from graphs that are expanding). Both distributions are based on a distribution of 3-regular m-vertex graphs that consist of an m-cycle and a random perfect matching. In the first distribution three such 3-regular $n/3$-vertex graphs are connected by three perfect matchings, whereas in the second distribution, for a large but constant ℓ, we connect ℓ graphs (in a cycle) by ℓ perfect matchings. (Lastly, we perform degree reduction.)

We stress that, while it is not *a priori* clear whether the lower bound on n-vertex expanders extends to $\Xi_c(n)$, it is also unclear whether the upper bound extends. (Indeed, like deciding, testing may become either easier or harder when the property is extended.) Actually, the known upper bound of [11], which asserts an $n^{0.501}$-time algorithm, applies to distinguishing c-expanders from graphs that are far from being c'-expanders, where $c' = \Omega(c^2)$. Both the lower bound of [6] and the lower bound of Theorem 1.2 apply also to this case (see Theorem 3.1), hereafter referred to as *pseudo-testing*.

We do not know whether the known upper bound of $n^{0.501}$ queries on pseudo-testing expander graphs [11] extends to pseudo-testing expansion, but rather prove the following result.

Theorem 1.3 (an upper bound on the query complexity of pseudo-testing expansion): *For every $d \geq 3$ and $c > 0$, there exists an $\mathrm{poly}(1/\epsilon) \cdot n^{0.667}$-query pseudo-tester that distinguishes between graphs in $\Xi_c(n)$ and n-vertex graphs that are ϵ-far from $\Xi_{c'}(n)$, where $c' = \Omega(c^2)$.*

The algorithm combines the distribution tester of generalized uniformity (i.e., uniformity over an unspecified subset of the domain) of Batu and Canonne [2] with the pseudo-tester of expanders of

vertices are replaced and edges are replaced by equality constraints. In these terms, the coboundary expansion is the robustness of the system as a local characterization of satisfiable assignments (cf. [17]); that is, it is the minimum ratio between the number of equalities that is unsatisfied by an assignment and the distance of this assignment from any satisfying assignment. Note that this formulation makes no explicit reference to the connected components of the graph, although such a reference is implicit in the notion of a satisfying assignment (i.e., satisfying assignments are functions that are constant on each connected component).

Kale and Seshadhri [11], which in turn is based on a strategy of Goldreich and Ron [7]. Specifically, combining the ideas of [7] and [2], we first select a few random vertices and estimate the size of the "tightly connected" components in which these vertices reside by estimating the pairwise and 3-way collision probabilities of the endpoints of (lazy) random walks (of logarithmic length) from each of these sampled vertices. Next, using the expander pseudo-tester of [11], we test whether these (tightly connected) components are expanders, while using their estimated sizes as obtained in the first step.

Essentially, the *tightly connected* component of a vertex s, defined by the first step, is the set of vertices that are reached with approximately the same probability in a random walk that starts at s, where the size of this set is inversely proportional to that probability. Recall that the pseudo-tester of [11] selects a few random vertices in the graph and estimates the collision probability of random walks from each of these vertices. While this main ("test vertex") subroutine of [11, Sec. 2] is invoked in [7, 11] on few uniformly selected vertices, here (in the second step) we invoke it on vertices reached via few (logarithmically long) random walks from each vertex that was (uniformly) selected in the first step. We note that the query complexity of our entire algorithm is dominated by the estimation of the 3-way collision probabilities (performed in the first step).

In our analysis, we heavily rely on [11], but cope with the fact that the (tightly connected) components are not perfectly isolated from one another. This is a problem because random walks in one (tightly connected) component can (rarely) enter another component, whereas the analysis of [11] refers to connected graphs (i.e., the random walk never exists the relevant subgraph). Nevertheless, using a very fine estimate of collision probability (and combining [2] and [11]) we can show that if the test accepts (w.h.p.), then almost all these (potentially overlapping) components are expanders. The challenging part of the analysis is showing that a subset of these expanders are both disjoint and cover almost all the vertices of the graph.

Specifically, assuming that the endpoint of a (logarithmically long) random walk starting at s is distributed almost uniformly on a vertex-set denoted W_s, we identify in W_s a relatively large induced subgraph that is close to being an expander and has relatively few edges going out of it. We stress that this subset, denoted W'_s, contains almost all of W_s (but not necessarily all of it). Hence, we show that for almost all vertices s in the graph, the corresponding set W'_s induces a subgraph that is close to being an expander and there are relatively few edges between W'_s and the rest of the graph. Furthermore, we show that every two such sets (i.e., the W'_s's) are either disjoint or almost equal. Lastly, we identify a subset of theses vertices (i.e., s's) such that the corresponding sets (i.e., the W'_s's) are disjoint and cover almost all the vertices of the graph.

An open problem. While our results establish non-trivial lower and upper bounds on the query complexity of pseudo-testing expansion (i.e., $\Omega(n^{1/2})$ vs $O(n^{\alpha})$ for every constant $\alpha > 2/3$), they leave a significant gap. Note that obtaining an $o(n^{2/3})$-query pseudo-tester for expansion requires bypassing the lower bound on the sample complexity of generalized uniformity testing (established by [2]). On the other hand, an $n^{0.5+\Omega(1)}$-query lower bound on pseudo-testing expansion would separate this task from pseudo-testing the set of expander graphs.

1.2 Testing girth

While one often focuses on the girth of connected graphs, there is absolutely no reason to restrict the notion of a girth to connected graphs.

Recall that the girth of a graph is the length of the shortest simple cycle in it; hence, the girth of an n-vertex graph G that is not cycle-free ranges between 3 (in case G contains a triangle) and n (in case G is a n-cycle). In particular, d-regular graphs have at most logarithmic girth (see [1, 19]), assuming $d \geq 3$, and random d-regular graphs are close to having logarithmic girth (see [13, p. 2]). Although logarithmic girth is not really related to expansion, high girth does have a flavor of expansion. In any case, we consider testing the following property.

Definition 1.4 (the high girth property): *For a function $g : \mathbb{N} \to \mathbb{N}$ such that $g(n) \in [3, \log n]$, we let $\Gamma_g(n)$ denote the set of n-vertex graphs that have girth at least $g(n)$, and let $\Gamma_g = \{\Gamma_g(n)\}_{n\in\mathbb{N}}$.*

A straightforward algorithm that picks a few random vertices and performs a BFS to depth $\lceil g/2 \rceil$ from each of them yields a tester for Γ_{g+1} (i.e., the set of graphs that have girth greater than g).

Theorem 1.5 (a straightforward tester of girth): *For every $g : \mathbb{N} \to \mathbb{N}$ and $d \in \mathbb{N}$, the property Γ_{g+1} can be tested using $O(d^{\lceil g/2 \rceil}/\epsilon)$ queries. Furthermore, Γ_{g+1} has a $O(d^{\lceil g/2 \rceil})$-time proximity-oblivious tester of linear detection probability.*[2]

We comment that an exponential dependence of the query complexity on the girth is unavoidable (see Theorem 1.6).

Proof: We rely on the fact that Γ_{g+1} can be characterized in terms of subgraph-freeness; that is, *a graph is in Γ_{g+1} if and only if it is C_ℓ-free for every $\ell \in \{3, .., g\}$*, where C_ℓ denotes the ℓ-vertex cycle. The furthermore clause follows by using the straightforward proximity-oblivious tester for subgraph-freeness (see [8, Sec. 5] and [4, Sec. 9.2.1]). ■

Theorem 1.6 (lower bound on testing girth): *For every $d \geq 3$ and $g : \mathbb{N} \to \mathbb{N}$ such that $g(n) \in [3, 0.499 \cdot \log_2 n]$, the query complexity of testing Γ_g is $\exp(\Omega(g))$.*

The proof is based on the fact that a random 3-regular k-vertex graph is extremely close to having girth greater than $0.499 \cdot \log_2 k$ but is far from having girth at least $2.001 \cdot \log_2 k$. In contrast, a $o(\sqrt{k})$-query algorithm that explores a random 3-regular k-vertex graph is unlikely to see a cycle. This means that an algorithm that makes $o(\sqrt{h})$ queries cannot distinguish between a random 3-regular k-vertex graph and a k-vertex graph consisting of k/h connected components each being a random 3-regular h-vertex graph. The theorem follows by setting $k = 2^{g(n)/0.499}$ and $h = 2^{g(n)/2.001}$.

Open problems. We mention two related open problems. The first is determining the query complexity of testing whether a graph is C_ℓ-free (i.e., contains no simple cycle of length ℓ). The second open problem is determining the query complexity of testing whether a graph contains no simple path of length ℓ. Both properties have straightforward $O(d^{\lceil \ell/2 \rceil}/\epsilon)$-time testers, but the question is whether one can do significantly better. (Initial studies of the second problem appeared in [16, 18].)[3]

[2] See [8, Sec. 2] and [4, Sec. 1.3.3].

[3] In particular, combining [18] with [12], one can obtain a poly$(d\ell/\epsilon)$-query tester for both problems, *in case the input graph is promised to belong to a minor-free class.*

1.3 Organization

In Section 2 we briefly review the definitions that underlie the study of testing graph properties in the bounded-degree graph model. As stated upfront, the main focus of this work is on the query complexity of testing expansion. In Section 3 we prove both the lower bound claimed in Theorem 1.2 and the upper bound (i.e., tester) claimed in Theorem 1.3. In Section 4 we study the query complexity of testing girth and prove Theorem 1.6. These two sections can be read independently of one another.

2 Preliminaries: Testing in the Bounded-Degree Graph Model

(Testing graph properties in the bounded-degree graph model was initiated in [6]. An extensive introduction to this line of research can be found in [4, Chap. 9].)

The bounded-degree graph model refers to a fixed (constant) degree bound, denoted $d \geq 2$. In this model, a graph $G = (V, E)$ of maximum degree d is represented by the incidence function $g : V \times [d] \to V \cup \{\perp\}$ such that $g(v, j) = u \in V$ if u is the j^{th} neighbor of v and $g(v, j) = \perp \notin V$ if v has less than j neighbors.[4] Distance between graphs is measured in terms of their foregoing representation; that is, as the fraction of (the number of) different array entries (over $d \cdot |V|$).

The tester is given oracle access to the representation of the input graph (i.e., to the incidence function g), where for simplicity we assume that $V = [n]$ for $n \in \mathbb{N}$. In addition, the tester is also given a proximity parameter ϵ and a size parameter (i.e., n). Recall that graph properties are sets of graphs that are closed under isomorphism.

Definition 2.1 (property testing in the bounded-degree graph model): *For a fixed $d \in \mathbb{N}$, a* tester for the graph property Π *is a probabilistic oracle machine T that, on input a* proximity parameter $\epsilon > 0$ *and* size parameter $n \in \mathbb{N}$*, and when given oracle access to an incidence function $g : [n] \times [d] \to [n] \cup \{\perp\}$, outputs a binary verdict that satisfies the following two conditions:*

1. *The tester accepts each graph $G = ([n], E)$ in Π with probability at least $2/3$; that is, for every $g : [n] \times [d] \to [n] \cup \{\perp\}$ representing a graph in Π* (and every $\epsilon > 0$)*, it holds that* $\Pr[T^g(n, \epsilon) = 1] \geq 2/3$.

2. *Given $\epsilon > 0$ and oracle access to any graph G that is ϵ-far from Π, the tester rejects with probability at least $2/3$; that is, for every $g : [n] \times [d] \to [n] \cup \{\perp\}$ that represents a graph that is ϵ-far from Π, it holds that* $\Pr[T^g(n, \epsilon) = 0] \geq 2/3$*, where the graph represented by g is* ϵ-far from Π *if for every $g' : [n] \times [d] \to [n] \cup \{\perp\}$ that represents a graph in Π it holds that*

$$\delta(g, g') \stackrel{\text{def}}{=} \frac{|\{(v, j) \in V \times [d] : g(v, j) \neq g'(v, j)\}|}{dn} > \epsilon. \tag{1}$$

The tester is said to have one-sided error probability *if it always accepts graphs in Π; that is, for every $g : [n] \times [d] \to [n] \cup \{\perp\}$ representing a graph in Π* (and every $\epsilon > 0$)*, it holds that* $\Pr[T^g(n, \epsilon) = 1] = 1$.

The query complexity of a tester for Π is a function (of the parameters d, n and ϵ) that represents the number of queries made by the tester on the worst-case n-vertex graph of maximum degree d, when

[4] For simplicity, we adopt the standard convention by which the neighbors of v appear in arbitrary order in the sequence $(g(v, 1), ..., g(v, \deg(v)))$, where $\deg(v) \stackrel{\text{def}}{=} |\{j \in [d] : g(v, j) \neq \perp\}|$.

given the proximity parameter ϵ. Fixing d, we typically ignore its effect on the complexity (equiv., treat d as a hidden constant); that is, we say that the tester has query complexity $q : \mathbb{N} \times [0,1] \to \mathbb{N}$ if on input parameters n and ϵ it makes $q(n,\epsilon)$ queries (on the worst-case n-vertex graph).

3 On the query complexity of testing expansion (i.e., Ξ)

In this section we study the query complexity of testing expansion, proving both the lower bound claimed in Theorem 1.2 and the upper bound (i.e., tester) claimed in Theorem 1.3. The following two proofs (equiv., subsections) can be read independently of one another.

3.1 Proof of Theorem 1.2

We start by restating Theorem 1.2 in a more general form.

Theorem 3.1 (a lower bound on the query complexity of pseudo-testing expansion): *For every $d \geq 3$ there exists $b > 0$ such that the following holds. For every $c \in (0,b]$ there exists $\epsilon > 0$ such that no $o(n^{1/2})$-query algorithm can distinguish between connected n-vertex graphs in $\Xi_b(n)$ and n-vertex graphs that are ϵ-far from $\Xi_c(n)$.*

Overview of the proof. We first prove the result for $d = 5$. Once this is achieved, we perform degree reduction to support the case of $d = 3$ (which is optimal [9]). The constant $b > 0$ is determined such that, with probability at least 0.999, a random n-cycle augmented by a random perfect matching is a connected graph with edge-expansion at least $10b$ (see [4, Clm. 9.18.1]).[5] The constant ϵ is proportional to c (e.g., $\epsilon = c/100$ would do).

As stated in the introduction, our proof uses the indistinguishability method [4, Sec. 7.2]; that is, we present two distributions of d-regular n-vertex graphs such that, on the one hand, these two distributions cannot be distinguished by a $o(\sqrt{n})$-query algorithm, while, on the other hand, a pseudo-tester for expansion must distinguish them. Specifically, the first distribution is concentrated on expanders (with edge-expansion at least b), whereas the second distribution is far from graphs that are expanding (i.e., $\Omega(c)$-far from Ξ_c). Both distributions are based on the aforementioned basic distribution of 3-regular m-vertex graphs, obtained by augmenting an m-cycle with a random perfect matching.

In the first distribution three 3-regular $n/3$-vertex basic graphs are connected by three perfect matchings, whereas in the second distribution, for a large but constant ℓ, we connect ℓ graphs (in a cycle) by ℓ perfect matchings. On the one hand, the first distribution is concentrated on expander graphs (i.e., on connected graphs that are b-edge expanding), whereas the second distribution is far from Ξ_c. (The first fact follows easily from [4, Clm. 9.18.1], whereas the second fact is proved in Claim 3.2.) On the other hand, as shown in Claim 3.3, a $o(\sqrt{n}$-query algorithm cannot distinguish these two distributions. The latter fact is proved by relying on [6, Lem. 7.4], which shows that any q-query algorithm that explores an m-vertex graph selected from the basic distribution sees a (simple) cycle with probability at most $O(q^2/m)$.

[5] For sake of simplicity, like in [6, Sec. 7] and [4, Sec 9.3], we do not disallow matching-edges that are parallel to the cycle edges. With constant probability, this may yield a non-simple graph (having parallel edges). We can regain a result regarding simple graphs by conditioning on the event in which no parallel edges were created. This event occurs with probability at least $(1/3) - o(1)$, since each matching-edge is parallel to a cycle-edge with probability $n/\binom{n}{2}$, whereas the events are almost three-wise independent (and one can use the inclusion-exclusion principle).

The actual proof. For a sufficiently large $\ell = O(1/c)$, we consider two distributions that are based on 3-regular m-vertex graphs that are selected from the basic distribution (i.e., each graph consists of an m-cycle augmented by a random perfect matching). For any n that is an integer multiple of 3ℓ, we consider the following two distributions on 5-regular n-vertex graphs.

***Distribution 1*:** *Three basic graphs connected by three perfect matchings.*

Pick three $n/3$-vertex graphs from the basic distribution and connect the vertices of each graph to the vertices of each of the other two graphs using a random perfect matching.

Note that, with high probability, the resulting graph is an expander. To see this consider any vertex-set of size at most $n/2$. If this vertex-set occupies less than 90% of each of the three parts, then this set is expanding by virtue of the expansion of the individual parts. On the other hand, if the vertex-set occupies at least 90% of one of the three parts, then the matching edges provide sufficient expansion (to the other two parts).

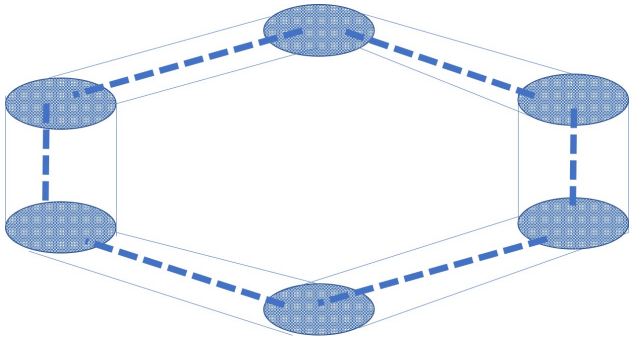

Figure 1: Illustration of Distribution 2 with $\ell = 6$.

***Distribution 2*:** *ℓ basic graphs connected (in a cycle) by ℓ perfect matchings.*

Pick ℓ graphs (each with n/ℓ vertices) from the basic distribution, arrange them in an ℓ-cycle, and connect the vertices of each graph to the vertices of each of the two neighboring graphs using a random perfect matching (see Figure 1). Indeed, Distribution 1 is a special case that corresponds to $\ell = 3$, but here we focus on larger (constant) ℓ.

In other words, each graph in this distribution is obtained by starting from an ℓ-cycle, replacing each vertex by a random graph drawn from the basic distribution, and replacing each edge by a random n/ℓ-edge matching.

We shall show (see Claim 3.2) that the resulting graph is $\Omega(1/\ell)$-far from being expanding (regardless of connectivity). For starters, note that the graph itself is far from being an expander (i.e., a connected expanding graph); this can be seen by considering a cut that places $\ell/2$ consecutive parts (i.e., n/ℓ-vertex graphs) on one side.

On the other hand, we shall show (see Claim 3.3) that an q-query algorithm exploring a random n-vertex graph drawn for any of the two distributions sees a (simple) cycle with probability $O(q^2/n)$. It follows that a $o(\sqrt{n})$-query algorithm, cannot distinguish the two distributions.

We mention that the specific structure of the basic distribution will be used in the proof of Claim 3.3, which relies on [6, Lem. 7.4], which in turn refers to the basic distribution. In contrast,

Claim 3.2 holds regardless of the structure of the graphs in the basic distribution and only relies on the fact that (whp) these graphs are expanders.

Claim 3.2 (the support of Distribution 2): *Let $G_0, ..., G_{\ell-1}$ be n/ℓ-vertex graphs drawn from the basic distribution, and let G be a graph that consists of these graphs that are connected as in Distribution 2; that is, for every $i \in \mathbb{Z}_\ell \stackrel{\text{def}}{=} \{0, 1, ..., \ell-1\}$, the vertices of G_i are connected to the vertices of $G_{i+1 \bmod \ell}$ by a perfect matching. Then, for every constant $c > 0$ and sufficiently large ℓ and n, the graph G is $\Omega(1)$-far from $\Xi_c(n)$.*

Proof: Recall that we have already seen that G is far from being an expander graph. This can be seen by considering the set of $n/2$ vertices in $\bigcup_{i \in [\ell/2]} G_i$. This set neighbors only $2n/\ell$ vertices that are not in it (i.e., the vertices in $G_{(\ell/2)+1}$ and G_0). In order to show that G is far from being expanding (regardless of connectivity), we show that each set of vertices S is either far from being isolated in G (i.e., there are $\Omega(|S|)$ edges going out of S) or is far from inducing an expander (i.e., the subgraph of G induced by S is far from being an expander).

Let S be an arbitrary vertex-set representing a potential connected component in an expanding graph that is supposedly close to G, and let S_i denote the subset of S that reside in G_i.

The easy case is when $|S_i| \leq n/2\ell$ for every $i \in \mathbb{Z}_\ell$. In this case, by virtue of the expansion of the G_i's, there are $\Omega(|S_i|)$ edges of G_i that go out of S_i, which implies that S is $\Omega(1)$-far from being isolated in G. The same reasoning applies if only a majority of the vertices in S reside in i's such that $|S_i| \leq n/2\ell$.

Hence, we may assume that a majority of the vertices in S reside in i's such that $|S_i| > n/2\ell$. Let I denote the set of such i's (i.e., $I = \{i \in \mathbb{Z}_\ell : |S_i| > n/2\ell\}$), and consider a partition of I into maximal consecutive sequences. Intuitively, long sequences contribute relatively too little to the expansion of the subgraph G induced by S, whereas short sequences contribute relatively too much to the (disallowed) connectivity of this subgraph from the rest of G. Hence, either this subgraph is far from being an expander or it is far from being isolated from the rest of G. Details follow.

Consider a generic sequence, $(i+1 \bmod \ell, ..., i+k \bmod \ell)$, such that $i+1 \bmod \ell, ..., i+k \bmod \ell \in I$ but $i, i+k+1 \bmod \ell \notin I$. On the one hand, the number of edges between $\bigcup_{j \in [k]} S_{i+j \bmod \ell}$ and the rest of S is at most $2n/\ell$, and the same holds for any set $\bigcup_{j \in [k',k'']} S_{i+j \bmod \ell}$ such that $k', k'' \in [k]$. In contrast, the size of the former set is at least $k \cdot n/2\ell$ (resp., $(k''-k'+1) \cdot n/2\ell$), so if k (resp., $k''-k'$) is large, then this contribution to expansion is too small. On the other hand, the number of edges between $\bigcup_{j \in [k]} S_{i+j \bmod \ell}$ and the rest of G (i.e., $[n] \setminus S$) is at least

$$\max(\Omega((n/\ell) - |S_{i+k \bmod \ell}|)\,,\, |S_{i+k \bmod \ell}| - (n/2\ell)) = \Omega(n/\ell)$$

where the first term is due to the expansion of $G_{i+k \bmod \ell}$ and the second term is due to the matching edges (between $G_{i+k \bmod \ell}$ and $G_{i+k+1 \bmod \ell}$).[6] To finish the argument, suppose that the vertices of $\bigcup_{i \in I} S_i$ reside in t consecutive sequences. Then, the following holds.

- There exists a set $S' \subseteq S$ of size approximately $|S|/2$ such that the cut between S' and $S \setminus S'$ has at most $2t \cdot n/\ell$ edges. Indeed, this set S' consists of all the vertices that reside in some of the consecutive sequences and potentially approximately half the vertices that reside

[6] Recall that $|S_{i+k \bmod \ell}| > n/2\ell$ and $|S_{i+k+1 \bmod \ell}| \leq n/2\ell$.

in one additional sequence.[7] It follows that the subgraph of G induced by S has relative edge-expansion at most $\frac{2t \cdot n/\ell}{|S'|} = O(t/|I|)$, where the last inequality uses $|S'| > |S|/3$ and $|S| > |I| \cdot n/2\ell$.

- The number of edges between S and the rest of G (i.e., $[n] \setminus S$) is at least $\Omega(t \cdot n/\ell) = \Omega(t/|I|) \cdot |S|$, where the last inequality uses $|S| > 2 \cdot |I| \cdot n/\ell$. So the relative contribution to the violation of the isolation of the subgraph induced by S is at least $\Omega(t/|I|)$.

Hence, for every constant $c > 0$, representing a possible edge-expansion of the graph G, either the subgraph of G induced by S is far from that level of edge-expansion (in case t is sufficiently small w.r.t $c \cdot |I|$) or this subgraph is far from being isolated in G (in case t is sufficiently large w.r.t $|I|$). ■

Claim 3.3 (the indistinguishability claim): *For a fixed ℓ and variable n and q, consider an arbitrary q-query algorithm that is given oracle access to an n-vertex graph, and let p_i denote the probability that this algorithm outputs 1 when given access to a graph selected from Distribution i. Then, $|p_1 - p_2| = O(q^2/n)$.*

Proof: Following the strategy of [6, Thm. 7.1], we may think of the algorithm as exploring the input graph, and refers to the subgraph that it sees during that exploration. We shall show that, when the graph is selected from either Distribution 1 or from Distribution 2, with probability at least $1 - O(q^2/n)$, the q-query explored sees a subgraph that is a forest. Clearly, conditioned on this event, the explorer cannot distinguish Distribution 1 from Distribution 2.

It was shown in [6, Lem. 7.4] that, with probability at least $1 - O(q^2/n)$, an algorithm that makes q queries to a graph selected from the basic distribution sees a forest. Furthermore, it was actually shown that none of its queries is answered by a vertex seen before (unless the edge was traversed in the opposite direction).[8]

Rather than extending the proof of [6, Lem. 7.4] to Distribution i, for both $i \in \{1, 2\}$, we reduce the analysis of the exploration of Distribution i to the analysis of the exploration of the basic distribution. For simplicity, we augment the model such that the five edges incident at each vertex are ordered so that the edges of the basic 3-regular graphs appear as the first three edges and an edge connecting a vertex of G_i to a vertex of $G_{i+1 \bmod \ell}$ appears as the 4th (resp., 5th) edge of the first (resp., second) vertex. We may also assume that the vertices of G_i are labeled $\langle i, \cdot \rangle$. Clearly, these conventions may only make the task of the explorer easier, whereas we may now assume that the explorer never tries to traverse an edge in both directions (see Footnote 8). We may also focus on Distribution 2, since Distribution 1 is a special case (i.e., $\ell = 3$).

We first reduce the analysis of an exploration of G, which is selected from Distribution 2, to the exploration of the graph G' that is obtained from G by removing all matching edges (i.e., the edges between G_i and $G_{i+1 \bmod \ell}$ for all i's). In other words, G' is generated by taking ℓ samples of the basic distribution (of n/ℓ-vertex graphs) and considering the graph that consists of these ℓ

[7] Suppose that, for every $j \in [t]$, the j^{th} sequence contains $m_j > n/2\ell$ vertices and recall that $\sum_{j \in [t]} m_j > |S|/2$. Then, there exists $t' \in [t]$ such that $\sum_{j \in [t']} m_j > |S|/2$ but $\sum_{j \in [t'-1]} m_j \leq |S|/2$. Note that the requirement is satisfied by either the t' first sequences or the first $t'-1$ sequences or the first $t'-1$ sequences augmented by half of the t'^{th} sequence.

[8] In fact, assuming that the edges incident at each vertex are ordered such that the clockwise (resp., anti-clockwise) cycle-edge appears as the first (resp., second) edge, we can avoid traversing an edge in both directions. We note that [6, Lem. 7.4] is proved in this model, which provides the explorer with additional information (for free).

isolated graphs. The key observation is that exploring the 4th (resp., 5th) edge of a vertex of G that resides in G_i is equivalent to selecting a random vertex in $G_{i+1 \bmod \ell}$ (resp., $G_{i-1 \bmod \ell}$). The point is that the latter operation can be emulated in G'. Hence, the probability that a q-query explorer encounters a vertex for the second time when exploring G is upper-bounded by the corresponding probability that refers to G', where G is selected from Distribution 2 and G' is obtained from G as defined above.

Lastly, the probability of that a q-query explorer encounters a vertex for the second time when exploring G' (distributed as above) is upper-bounded by the probability of this happening when a q-query algorithm explores a single graph taken from the basic distribution (on n/ℓ-vertex graphs).[9] Hence, we get an upper bound of $O(q^2/(n/\ell)) = O(q^2/n)$. ■

Conclusion: Recalling that Distribution 1 is concentrated on expander graphs and using Claim 3.2, we conclude that a tester for Ξ_c must distinguish between Distribution 1 and Distribution 2 with gap at least $0.66 - 0.34 > 0.3$; that is, in terms of Claim 3.3, we must have $|p_1 - p_2| \geq 0.3$. On the other hand, Claim 3.3 asserts that $|p_1 - p_2| = O(q^2/n)$, where q is the query complexity of the tester. Hence, $q = \Omega(\sqrt{n})$ must hold.

Degree reduction: As stated upfront, the result was proved for $d = 5$, which falls short of the stated theorem, which asserts the same lower bound for any $d \geq 3$. Using a local reduction (cf. [4, Sec. 7.4]), we derive the result for $d = 3$ (and the claim for any constant $d \geq 3$ follows triviality). Specifically, we capitalized on the fact that the lower bound was established for 5-regular graphs and that expansion is preserved under a standard degree reduction (which replaces each vertex by a 5-cycle and uses a single edge of each new vertex in order to connect to some other 5-cycle). That is, given oracle access to a 5-regular n-vertex graph G, we emulate oracle access to a 3-regular $5n$-vertex graph G' such that each vertex v of G is replaced by a 5-cycle C_v in G' and the i^{th} edge of v (in G) is replaced by an edge that is incident at the i^{th} vertex of the 5-cycle C_v. Note that this transformation preserves expansion and connectivity as well as a constant factor of the distance from being expanding (where the specific (constant) level of expansion may vary).

3.2 Proof of Theorem 1.3

We start by restating Theorem 1.3 in a more accurate form.

Theorem 3.4 (an upper bound on the complexity of pseudo-testing expansion): *For every $d \geq 3$, $c > 0$ and $\alpha > 2/3$, there exist $c' = \Omega(c^2)$ and an $\text{poly}(1/\epsilon) \cdot n^\alpha$-query algorithm that distinguishes between n-vertex graphs in $\Xi_c(n)$ and n-vertex graphs that are ϵ-far from $\Xi_{c'}(n)$. Furthermore, the time complexity of this algorithm is linear in its query complexity.*

We start with an overview of the proof. It will be followed by a presentation of the actual algorithm. The actual analysis of the algorithm copes with issues that are hardly mentioned in the overview; see the second part of the analysis (following Definition 3.4.3).

[9] Letting q_i denotes the number of queries made to G_i (when exploring G'), we get a probability bound of $\sum_i O(q_i^2/(n/\ell)) \leq O((\sum_i q_i)^2 \ell/n)$.

Overview. We follow the strategy of Goldreich and Ron [7], which reduces testing whether an n-vertex graph is an expander to testing whether an $O(\log n)$-step lazy random walk that starts at a fixed (random) vertex ends in a distribution that is almost uniform over the vertices of the graph. (Using lazy random walks implies that the stationary distribution is uniform on the vertices of each connected component, also in case the graph is not regular.)[10] Recall that testing the uniform distribution is performed by approximating the collision probability using $O(\sqrt{n})$ samples and comparing it to the value $1/n$ (i.e., the collision probability of the uniform distribution over n elements).

The problem we face is that we also have to accept n-vertex graphs that consists of connected components (of unknown sizes) that are each an expander. Hence, selecting a random start vertex, we would like to test whether the endpoint of a random walk starting at this vertex is distributed uniformly over the vertices of the connected component to which this start vertex belongs. That is, testing expansion is reduced to testing generalized uniformity (i.e., uniformity over an unspecified subset). Fortunately, a study of the latter testing problem was initiated by Batu and Canonne [2], and its sample complexity was shown to be $\Theta(n^{2/3}) \cdot \mathrm{poly}(1/\epsilon)$. Specifically, the tester consists of comparing the pairwise collision probability to the 3-way collision probability (and accepting iff the square of the first value approximately equals the second value).[11]

The latter reduction also provides us with a very good estimate of the "effective size" of the relevant connected component; that is, a number m such that there are at least $(1-\epsilon)\cdot m$ vertices in this connected component that are each reached by the random walk with probability $(1 \pm \epsilon)/m$.[12] For simplicity, let us assume that we get an $1 \pm \epsilon$ factor estimate of the actual size of the connected component. As we shall see, such an estimate is extremely valuable towards invoking the expander pseudo-tester of Kale and Seshadhri [11]. Specifically, while [11, Thm. 3.4] is stated for a known graph size, little changes if this size is known up to a factor of $1 \pm \sigma$, where $\sigma = m^{-\Omega(1)}$ is a key parameter that reflects the distance between the distribution of the endpoint of a random walk and uniform distribution.[13] This means that we have to invoke the generalized uniformity tester of [2] with proximity parameter set to σ (or rather $\min(\epsilon, \sigma)$), which implies that our query complexity will be $O(n^{2/3}) \cdot \mathrm{poly}(1/\sigma)$.

Unfortunately, the reduction to testing generalized uniformity does not provide an approximation to the size of the connected component but rather to its the "effective size" (as defined above). Yet, as shown below (see Claim 3.4.1), the number of edges in the cut between the counted vertices (i.e., those counted in the "effective size") and the uncounted ones is very small. Hence, using random walks that start at a random vertex chosen among the counted vertices, the distribution of walks that avoid the cut edges is very close to the real distribution, and so we can apply the result of [11, Thm. 3.4].

This may be a good time to explicitly (re)state [11, Thm. 3.4]. Loosely speaking, it asserts that, for $\ell' = O(\log m)$, *if for at least* $1 - \gamma$ *fraction of the vertices* v *in a* m*-vertex graph it holds that the collision probability of the endpoint of an* ℓ'*-step random walk starting at* v *does not exceed*

[10] By a lazy random walk we mean that at at each step, the edge leading to each neighbor is traversed with probability $1/2d$, and otherwise the walk stays in the current vertex.

[11] For simplicity, we do not use the subsequent tester of [3], which achieves optimal dependence on ϵ.

[12] Loosely speaking, m equals the reciprocal of the collision probability of the random walk. Indeed, it is not *a priori* clear that there exists a set of at least $(1-\epsilon)\cdot m$ vertices that are each reached with probability $(1 \pm \epsilon)/m$, but this can be easily proved (see the proof of Claim 3.4.1).

[13] In [11], σ has the form $n^{-\mu/4}$, where n denotes the size of the graph, the query complexity is essentially $n^{(1+\mu)/2}$ and $c' = \mu \cdot c^2/400$.

$\tau = (1 + m^{-\Omega(1)})/m$, *then the graph is* $O(\gamma)$*-close to being a* c'*-edge expander.* Needless to say, the hidden constants in ℓ' and τ are related to the constants c and c'. Most importantly, the hidden constant in τ, hereafter denoted β (i.e, $\tau = (1 + m^{-\beta})/m$), is proportional to the ratio c'/c^2. (The constant in $\ell' = O(\log m)$ is set such that if the graph is a c-edge expander, then an ℓ-step random walk on the graph ends at a vertex that is $1/m^2$-close to the uniform distribution.)

The actual algorithm. Let $\beta = \Theta(\alpha - (2/3))$, where $\alpha > 2/3$ is the constant claimed in Theorem 3.4. On input a size parameter n and proximity parameter $\epsilon > 0$, and given oracle access to an n-vertex graph G (of maximal degree d), we perform the following steps $O(1/\epsilon^2)$ times.

***Step 1*:** *Testing generalized uniformity of the distribution generated by a random walk.*

Select a vertex $s \in [n]$ uniformly at random, and invoke the generalized uniformity tester of [2], while setting its proximity parameter to $\epsilon' = \epsilon^2 \cdot n^{-3\beta}$ and providing it with $\mathrm{poly}(1/\epsilon') \cdot n^{2/3}$ samples, each generated by taking an $O(\log n)$-step (lazy) random walk starting at s. Specifically, at each step of the random walk, the edge leading to each neighbor is traversed with probability $1/2d$, and otherwise the walk stays in the current vertex; the generated sample is the vertex reached at the end of the walk.

(Recall that the generalized uniformity tester distinguishes between distributions that are uniform over some subset of $[n]$ and distributions that are ϵ'-far from the class of generalized uniform distributions, where the distance here is the total variation distance.)

Using error reduction (for the same start vertex s), we may assume that the error probability of the generalized uniformity tester is $o(\epsilon^2)$ (rather than at most 1/3).

If the generalized uniformity tester rejects, then we reject. Otherwise, let m denote the reciprocal of the (empirical) collision probability of the tested distribution; that is, the sampled distribution was deemed ϵ'-close to the uniform distribution over a set of m elements.

***Step 2*:** *Pseudo-testing the expansion of a subgraph reached from* s*.*

Invoke the expander pseudo-tester of [11, Sec. 2], on the graph G (or rather on the connected component of s), while setting the size (of the graph) parameter to m, setting the proximity parameter to $\epsilon'' \stackrel{\text{def}}{=} \sqrt{\epsilon'}$, and selecting the $O(1/\epsilon'')$ starting vertices by taking $O(\log n)$-step random walks from s. If this pseudo-tester rejects, then we reject.

Recall that (when using proximity parameter ϵ'') the expander pseudo-tester consists of $O(1/\epsilon'')$ invocations of a "vertex tester" (see [11, Sec. 2]), where in [11, Sec. 2] each invocation is provided with a uniformly selected vertex (as its explicit input). Here, each of these $O(1/\epsilon'')$ invocations is provided with the endpoint of an $O(\log n)$-step random walk starting at s. (This pseudo-tester distinguishes c-edge expanders from graphs that are ϵ''-far from being c'-edge expanders.)

We set the deviation parameter (denoted $\sigma = n^{-\mu/4}$ in [11, Sec. 2]) to $\delta = O(\epsilon'')$; recall that the "vertex tester" estimates the collision probability of the (endpoints) of random walks and accepts if the estimated value does not exceed $(1 + \delta)/m$. Recall that the query complexity of the "vertex tester" (of [11, Sec. 2]) is $\mathrm{poly}(1/\delta) \cdot \widetilde{O}(n^{1/2})$, which is dominated by the query complexity of Step 1, which in turn is $\mathrm{poly}(1/\epsilon'') \cdot n^{2/3}$.

Using error reduction, we may assume that the error probability of the expander pseudo-tester is $o(\epsilon^2)$ (rather than at most 1/3).

If none of the steps rejected, then we accept.

Comment: Step 2 vs [11]. When seeking to pseudo-test whether a graph G is an expander, Kale and Seshadhri invoke the "vertex tester" of [11, Sec. 2] for $t = O(1/\epsilon)$ times at t uniformly and independent selected vertices of G, where ϵ is the proximity parameter of this tester. Instead, Step 2 invokes this "vertex tester" for $t'' = O(1/\epsilon'')$ times at t'' vertices that are each generated by a $O(\log n)$-step random walk starting at s, where our aim is to pseudo-test the expansion of an "effective connected component" (denoted W_s below) in which s resides. This pseudo-tester uses proximity parameter $\epsilon'' = \epsilon \cdot n^{-3\beta/2}$, where ϵ is the proximity parameter of the entire algorithm.

The analysis. The query complexity of the foregoing algorithm is $\mathrm{poly}(1/\epsilon') \cdot n^{2/3} = \mathrm{poly}(n^\beta/\epsilon) \cdot n^{2/3}$, which is $\mathrm{poly}(1/\epsilon) \cdot n^\alpha$. (Hence we may obtain any constant $\alpha > 2/3$, but this comes at the cost of decreasing the ratio c'/c^2 (by a factor of $\alpha - (2/3)$).) Next, observe that if G is expanding (i.e., $G \in \Xi_c$), then the foregoing algorithm accepts (w.h.p). Specifically, in this case, each invocation of Step 1 provides the generalized uniformity tester with samples that are distributed almost uniformly in the relevant connected component, and (w.h.p.) the tester will accept while providing us with a $1 \pm \epsilon'$ factor approximation of the size of the connected component. Hence, in Step 2, the expander pseudo-tester will accept (w.h.p).

As usual, the challenging part is the analysis of graphs that are ϵ-far from $\Xi_{c'}$. Typically, one proves the contrapositive; that is, if the tester accepts G with high probability, then G is ϵ-close to $\Xi_{c'}$. We focus on showing that if, for at least $1 - \gamma$ fraction of the choices of a vertex s in a connected component C, with high probability Steps 1–2 do not reject, then C is $O(\gamma)$-close to $\Xi_{c'}$. (The rest of the argument is quite standard.)[14] Note that *if Step 1 returned a very good estimate of the actual size of* C, then [11, Thm. 3.4] would imply that C is $O(\gamma)$-close to a c'-edge expander (which is obviously in $\Xi_{c'}$).

Unfortunately, we are only guaranteed that if Step 1 rejects with small probability, then it returns the size of am arbitrary set $W_s \subseteq C$ such that the distribution of the endpoint of a random walk on C that starts in s is close to the uniform distribution on W_s. Specifically, the distribution of the endpoint of a random walk that starts in s, denoted X_s, is ϵ'-close to the uniform distribution on W_s, since Step 1 invokes the generalized uniformity tester with proximity parameter ϵ'. A key observation is that the relative size of the cut between W_s and $C \setminus W_s$ is related to the total variation distance between X_s and the uniform distribution on W_s. We actually prove the following.

Claim 3.4.1 (the cut $E(W_s, C \setminus W_s)$ is relatively small): *Let $\ell = O(\log n)$ and X_s denote the distribution of the endpoint of an ℓ-step random walk on C that starts in s. If X_s is η-close to being uniform over the vertex set W_s, then there exists $W'_s \subseteq W_s$ such that $\Pr[X_s \in W'_s] \geq 1 - 2\eta^{1/2}$ and the number of edges between W'_s and the rest of C is $O(\eta^{1/2} \cdot |W'_s|)$. Furthermore, for each $v \in W'_s$ it holds that $\Pr[X_s = v] = (1 \pm O(\eta^{1/2}))/|W_s|$.*

It follows that, for any ℓ', an ℓ'-step random walk starting at a vertex selected uniformly in W_s (alternatively, starting at X_s) stays in W'_s with probability $1 - O(\ell' \cdot \eta^{1/2})$.

[14] Specifically, for $\gamma = \epsilon/O(1)$, at most γ^2 of the vertices s of G lead Steps 1–2 to reject with non-small probability. Hence, at most a γ fraction of the vertices reside in connected components C such that at most γ fraction of the vertices $s \in C$ lead Steps 1–2 to reject with non-small probability.

Proof: Letting $W'_s \stackrel{\text{def}}{=} \{v \in W_s : \Pr[X_s = v] = (1 \pm \eta^{1/2})/|W_s|\}$, we observe that $\Pr[X_s \notin W'_s] < 2\eta^{1/2}$, because otherwise at least $\eta^{1/2}$ of the probability mass resides on vertices that are either each reached with probability larger than $(1+\eta^{1/2})/|W_s|$ or each reached with probability smaller than $(1-\eta^{1/2})/|W_s|$. It follows that $|W'_s| > (1-3\eta^{1/2}) \cdot |W_s|$, because $\Pr[X_s \in W'_s] \leq |W'_s| \cdot (1+\eta^{1/2})/|W_s|$. Observe that each of the vertices in W'_s is reached by an $(\ell-1)$-step (lazy) random walk (from s) with probability at least $p = (1-o(1)) \cdot (1-\eta^{1/2})/|W_s| > 0.9/|W_s|$, since ℓ-step *lazy* random walks stay in place (at each step) with probability at least $1/2$. Hence, each cut edge (i.e., an edge in $E(W'_s, C \setminus W'_s)$) is traversed in the ℓ^{th} step with probability at least $p/2d = \Omega(1/|W'_s|)$, and it follows that the walk ends outside of W'_s with probability $\Omega(|E(W'_s, C \setminus W'_s)|/|W'_s|)$. Recalling that $\Pr[X_s \notin W'_s] \leq 2\eta^{1/2}$, the claim follows (i.e., $|E(W'_s, C \setminus W'_s)| = O(\eta^{1/2} \cdot |W'_s|)$). ■

Definition 3.4.2 (good vertices): *A vertex $s \in C$ is called* good *if performing Step 1 starting at s makes the generalized uniformity tester accept with probability $1 - o(\epsilon)$.*

Recall that (given that the test accepted w.h.p.) we shall assume that a $1-\gamma$ fraction of the vertices in C are good, and that every good vertex s defines a set of vertices W_s such that the distribution of the endpoint of an $O(\log n)$-step random walk on C that starts in s is ϵ'-close to being uniform over the vertex set W_s.

Using Claim 3.4.1, we observe that a random $O(\log n)$-step walk *starting at a random vertex w in W_s* is extremely unlikely to leave W'_s, and the expander pseudo-tester (applied in Step 2 with proximity parameter ϵ'') can be viewed as taking place on the subgraph of C induced by W'_s. Specifically, applying Claim 3.4.1 to a good vertex s with $\eta = n^{-3\beta}$, it follows that the collision probability of random walks starting at X_s and not leaving W'_s at least a $(1 - o(n^{-\beta}))$ factor of the value estimated in Step 2. Recall that Step 2 estimates the collision probability from $O(1/\epsilon'')$ vertices drawn from X_s, which implies that *if Step 2 does not reject* (w.h.p.), *then for at least $1-\epsilon''$ fraction of the vertices w in W'_s the collision probability of a random walk on W'_s that starts in w is at most* $(1 + o(|W'_s|^{-\beta}))/|W'_s|$. Applying [11, Thm. 1.1], it follows that in this case the subgraph of C induced by W'_s is $O(\epsilon'')$-close to a c'-edge expander. Furthermore, we can afford to omit the edges of the cut $E(W'_s, C \setminus W'_s)$, and proceed with the rest of $C \setminus W'_s$. Unfortunately, we have to do this carefully due to the affect on $C \setminus W'_s$ (see below). But let us first summarize the foregoing discussion.

Definition 3.4.3 (excellent vertices): *A good vertex $s \in C$ is called* excellent *if performing Step 2 starting at s makes the expander pseudo-tester accept with probability $1 - o(\epsilon)$.*

The foregoing discussion asserts that if s is excellent, then the subgraph induced by W'_s is $O(\epsilon'')$-close to a c'-edge expander, whereas $|E(W'_s, C \setminus W'_s)| = O(\epsilon'' \cdot |W'_s|)$, where $\epsilon'' = \sqrt{\epsilon'} \leq n^{-3\beta/2} = o(n^{-\beta}/\log n)$. As stated, we can afford to omit the edges of $E(W'_s, C \setminus W'_s)$, and proceed with the rest of $C \setminus W'_s$, but this may cause a problem later on.

The problem is that the edges that we omitted (i.e., $E(W'_s, C \setminus W'_s)$) still appear in the graph, and so random walks from other vertices of $C \setminus W'_s$ may traverse these edges. That is, the set $W'_{s'}$ defined for some excellent $s' \in C \setminus W'_s$ may include both endpoint of an edge in $E(W'_s, C \setminus W'_s)$; in this case, the omitted edge may contribute to the edge-expansion of the subgraph induced by $W'_{s'}$ (and so omitting it in service of W'_s may harm $W'_{s'}$). Hence, we need to avoid this situation. For starters, we claim that if $W'_{s'}$ contains an edge in $E(W'_s, C \setminus W'_s)$, it holds that $W'_{s'}$ and W'_s are either almost disjoint or almost overlapping. Actually, this dichotomy holds for any excellent

pair of vertices (regardless of whether the cut defined by one excellent vertex is contained in the subgraph induced by the other).

Claim 3.4.4 ($W'_{s'}$ and W'_s are either disjoint or overlapping): *Let s and s' be two excellent vertices in C. Then, $\min(|W'_s \cap W'_{s'}|, |W'_s \setminus W'_{s'}|) = O(\epsilon'' \cdot \max(|W'_s|, |W'_{s'}|))$, where $\epsilon'' = \sqrt{\epsilon'}$. Furthermore, if $|W'_s \cap W'_{s'}| = \omega(\epsilon'' \cdot \max(|W'_s|, |W'_{s'}|))$, then $|W'_s \cap W'_{s'}| \geq (1 - O(\epsilon'')) \cdot \max(|W'_s, |W'_{s'}|)$.*

Proof: We start with the main claim. The key observation is that by Claim 3.4.1 (applied to $X_{s'}$ with $\eta = \epsilon' \leq n^{-3\beta}$), the cut between $W'_s \cap W'_{s'}$ and $W'_s \setminus W'_{s'}$ contains at most $|E(W'_{s'}, C \setminus W'_{s'})| = O(\eta^{1/2} \cdot |W'_{s'}|)$ edges, whereas by the fact that the subgraph induced by W'_s is $O(\epsilon'')$-close to a c'-edge expander it follows that this cut has size at least $c' \cdot \min(|W'_s \cap W'_{s'}|, |W'_s \setminus W'_{s'}|) - O(\epsilon'') \cdot |W'_s|$. Recalling that $\eta^{1/2} = \epsilon''$ and $c' = \Omega(1)$, it follows that

$$\min(|W'_s \cap W'_{s'}|, |W'_s \setminus W'_{s'}|) = O(\epsilon'' \cdot |W'_{s'}|) + O(\epsilon'' \cdot |W'_s|).$$

Turning to the furthermore claim, observe that the hypothesis (combined with the main claim) implies that $|W'_s \setminus W'_{s'}| = O(\epsilon'') \cdot \max(|W'_s|, |W'_{s'}|)$. Likewise, $|W'_{s'} \setminus W'_s| = O(\epsilon'') \cdot \max(|W'_s|, |W'_{s'}|)$. The furthermore claim follows. ■

We actually need a more general claim.

Claim 3.4.5 (a generalization of Claim 3.4.4): *Let $s_1, ..., s_t$ and s' be $t+1$ excellent vertices in C such that the $W'_{s_1}, ..., W'_{s_t}$ are pairwise disjoint. Let D_{s_i} be a set of $O(\epsilon'' \cdot |W'_{s_i}|)$ vertex-pairs such that the subgraph induced by W'_{s_i} differs from a c'-edge expander only on D_{s_i}. Then, either for some $i \in [t]$ it holds that $|W'_{s'} \cap W'_{s_i}| \geq |W'_{s_i}| - O(\epsilon'') \cdot \max(|W'_{s_i}|, |W'_{s'}|)$ or $|W'_{s'} \cap \bigcup_{i \in [t]} W'_{s_i}| = O(\epsilon'' \cdot |W'_{s'}| + |D' \cap W'_{s'}|)$, where D' is the set of vertices incident to $\bigcup_{i \in [t]} D_{s_i}$.*

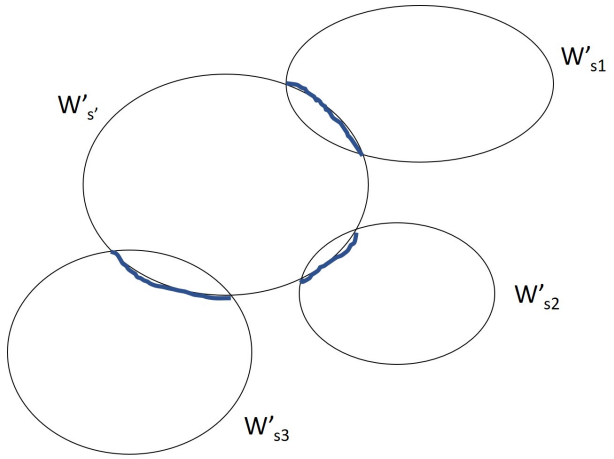

Figure 2: The cut between $W' \cap W'_{s'}$ and $W' \setminus W'_{s'}$ is marked.

Proof: Letting $W' = \bigcup_{i \in [t]} W'_{s_i}$, we consider the cut between $W' \cap W'_{s'}$ and $W' \setminus W'_{s'}$ (see illustration in Figure 2). On the one hand, applying Claim 3.4.1 to $W'_{s'}$, it follows that this cut contains at most $|E(W'_{s'}, C \setminus W'_{s'})| = O(\epsilon'' \cdot |W'_{s'}|)$ edges. On the other hand, using the hypothesis that the

subgraph induced by each W'_{s_i} differs from a c'-edge expander only on D_{s_i}, it follows that this cut (i.e., $E(W' \cap W'_{s'}, W' \setminus W'_{s'})$ has size at least

$$\sum_{i\in[t]} \Big(c' \cdot \min(|W'_{s_i} \cap W'_{s'}|, |W'_{s_i} \setminus W'_{s'}|) - |D_{s_i} \cap E(W'_{s_i} \cap W'_{s'}, W'_{s_i} \setminus W'_{s'})| \Big),$$

where we use the hypothesis that $W'_{s_1}, ..., W'_{s_t}$ are disjoint. Letting D'_{s_i} denote the set of vertices incident to D_{s_i}, we have

$$|D_{s_i} \cap E(W'_{s_i} \cap W'_{s'}, W'_{s_i} \setminus W'_{s'})| = O(|D'_{s_i} \cap W'_{s'}|).$$

Hence,

$$\sum_{i\in[t]} \Big(c' \cdot \min(|W'_{s_i} \cap W'_{s'}|, |W'_{s_i} \setminus W'_{s'}|) - O(|D'_{s_i} \cap W'_{s'}|) \Big) = O(\epsilon'' \cdot |W'_{s'}|).$$

We now consider two cases.

Case 1: $|W'_{s_i} \cap W'_{s'}| \leq |W'_{s_i} \setminus W'_{s'}|$ ***holds for each*** $i \in [t]$**.** In this case, it follows that

$$\sum_{i\in[t]} \Big(|W'_{s_i} \cap W'_{s'}| - O(|D'_{s_i} \cap W'_{s'}|) \Big) = O(\epsilon'' \cdot |W'_{s'}|).$$

Recalling that $D'_{s_1}, ..., D'_{s_t}$ are disjoint, we get $|W'_{s'} \cap W'| = O(\epsilon'' \cdot |W'_{s'}| + |D' \cap W'_{s'}|)$ as claimed, where $D' = \bigcup_{i\in[t]} D'_{s_i}$.

Case 2: $|W'_{s_i} \cap W'_{s'}| > |W'_{s_i} \setminus W'_{s'}|$ ***holds for some*** $i \in [t]$**.** In this case (by Claim 3.4.4) it holds that $|W'_{s_i} \setminus W'_{s'}| = O(\epsilon'' \cdot \max(|W'_{s_i}|, |W'_{s'}|))$, and $|W'_{s'} \cap W'_{s_i}| \geq |W'_{s_i}| - O(\epsilon'') \cdot \max(|W'_{s_i}|, |W'_{s'}|)$ follows.

The claim follows. ■

With Claims 3.4.1 and 3.4.5 at our disposal, we can cover most of the vertices of the connected component C by disjoint sets such that the subgraph induced by each set is close to being an expander and there are few edges between these sets. These sets will be based on the sets W'_s that correspond to excellent vertices $s \in C$. Specifically, we use an iterative process. In the i^{th} iteration, having previously determined the sets $W''_{s_1}, ..., W''_{s_{i-1}}$, we seek an excellent vertex s_i such that $|W'_{s_i} \cap \bigcup_{j\in[i-1]} W''_{s_j}| = O(\epsilon'' \cdot |W'_{s_i}|)$. If no such vertex exists, then we halt, and otherwise we pick such a vertex s_i and define $W''_{s_i} \leftarrow W'_{s_i} \setminus \bigcup_{j\in[i-1]} W''_{s_j}$. Specifically, among the legitimate candidates for s_i, we pick the one that maximizes $|W'_{s_i}|$. This guarantees that $|W'_{s_i}| \leq |W'_{s_j}| \leq \frac{|W''_{s_j}|}{1-O(\epsilon'')}$ for all $j \in [i-1]$. We stress that the subgraph of C induced by W''_{s_i} is $O(\epsilon'')$-close to being a c'-edge expander and that $|E[W''_{s_i}, C \setminus W''_{s_j}| = O(\epsilon'' \cdot |W''_{s_i}|)$.

We shall show that when this iterative process terminates (after t iterations) almost all the vertices of C are covered (i.e., are included in one of the sets $W''_{s_1}, ..., W''_{s_t}$). That is, as long as there are significantly many uncovered vertices, there exists an excellent vertex s' such that $W'_{s'}$ has small intersection with the previous $i-1$ sets $W''_{s_1}, ..., W''_{s_{i-1}}$. Intuitively, this holds because, for a random excellent vertex s', the expected intersection of $W'_{s'}$ with any subset of C is linearly related to the density of this subset. This will allow us to lower-bound the intersection of $W'_{s'}$ with the set of uncovered vertices as well as upper-bound the intersection of $W'_{s'}$ with the set $\bigcup_{j\in[i-1]} D'_{s_j}$.

Claim 3.4.6 (the sampling features of excellent vertices): *For $\eta \leq \sigma/3 < 1/3$, suppose that $(1-\eta)\cdot|C|$ of the vertices in the connected component C are excellent, and let S and T be subsets of C having densities σ and τ respectively. Then, there exists a excellent vertex $s \in C$ such that both $\Pr[X_s \in S] \geq \sigma/3$ and $\Pr[X_s \in T] = O(\tau/\sigma)$ hold.*

Using Claim 3.4.1, it follows that $|W'_s \cap S| > ((\sigma/3) - O(\epsilon''))\cdot|W'_s|$, since X_s is $O(\epsilon'')$-close to be uniform on W'_s. Likewise, we get $|W'_s \cap T| = O((\tau/\sigma)+\epsilon'')\cdot|W'_s|$.

Before proving the Claim 3.4.6, we use it to complete the analysis of the foregoing iterative process. In the i^{th} iteration, we set $S = C \setminus \bigcup_{j\in[i-1]} W''_{s_j}$ and $T = \bigcup_{j\in[i-1]} D'_{s_j}$. Observe that $|T| = O(\epsilon''\cdot|C|)$, whereas we care of the case that $|S| = \Omega(\epsilon\cdot|C|)$. Furthermore, recall that the number of excellent vertices in C (i.e., $(1-\eta)\cdot|C|$) is at least $|C| - (|S|/3)$. In the latter case (i.e., $|S| = \Omega(\epsilon\cdot|C|)$), using Claim 3.4.6, we infer that there exists an excellent vertex s that satisfies $|W'_s \cap W''_{s_j}| \leq |W'_s| - |W'_s \cap S| \leq (1-\Omega(\epsilon))\cdot|W'_s|$ for every $j \in [i-1]$. Using Claim 3.4.5 (and recalling that $|W'_s \cap T| = O(\epsilon''/\epsilon)\cdot|W'_s|$), it follows that $|W'_s \cap \bigcup_{j\in[i-1]} W''_{s_j}| = O(\epsilon''\cdot|W'_s|/\epsilon)$, since $|W'_s \cap W''_{s_j}| \geq |W''_{s_j}| - O(\epsilon'')\cdot\max(|W'_s|,|W''_{s_j}|)$ is impossible (because $|W''_{s_j}| \geq (1-O(\epsilon''/\epsilon))\cdot|W'_{s_j}|$ and $|W'_s| \leq |W'_{s_j}|$ imply that $|W''_{s_j}| - O(\epsilon''/\epsilon)\cdot\max(|W'_s|,|W''_{s_j}|) \geq (1-O(\epsilon''/\epsilon))\cdot|W'_{s_j}| > (1-\Omega(\epsilon))\cdot|W'_s|$).[15]

Proof: Note that uniformly selecting $s \in C$ and outputting X_s (which is generated by a *lazy* random walk starting at s)[16] yields a uniform distribution over C, whereas selecting uniformly a excellent vertex s and outputting X_s yields a distribution, denoted Z, that is η-close to uniform. Hence,

$$\text{Exp}_s\left[\Pr[X_s \in S]\right] \geq \sigma - \eta \geq 2\sigma/3,$$

where the expectation is over uniformly distributed excellent $s \in C$. It follows that for at least a $\sigma/3$ fraction of the excellent $s \in C$, it holds that $\Pr[X_s \in A] \geq \sigma/3$. Likewise,

$$\text{Exp}_s\left[\Pr[X_s \in T]\right] \leq \frac{\tau}{1-\eta} < \tau/2,$$

which implies that for more than $1-(\sigma/4)$ fraction of the excellent $s \in C$, it holds that $\Pr[X_s \in T] \leq \tau/8\sigma$. The claim follows. ■

Conclusion. For each connected component C, if $1-\gamma$ fraction of the vertices in C are excellent, then C is $O(\gamma)$-close to some graph in $\Xi_{c'}$. This expanding graph consists of the connected components over the vertex-sets $W''_{s_1},...,W''_{s_t}$ and $|C| - \sum_{i\in[t]}|W''_{s_i}|$ isolated vertices. Recall that the i^{th} connected component (having vertex-set W''_{s_i}) is $O(\epsilon'')$-close to being a c'-edge expander. Indeed, we have omitted all edges in $E(W''_{s_i}, C\setminus W''_{s_i})$ as well as all edges incident at $R \stackrel{\text{def}}{=} C \setminus \bigcup_{i\in[t]} W''_{s_i}$, while upper-bounding that their number by

$$\sum_{i\in[t]} O(\epsilon''\cdot|W''_i|) \;\; + O(|R|) \;=\; O(\epsilon''\cdot|C|) + O(\gamma\cdot|C|) \;=\; O(\gamma\cdot|C|),$$

since $\epsilon'' = o(\epsilon)$ whereas $\gamma = \Omega(\epsilon)$. Recalling the standard counting argument outlined in Footnote 14, this completes the proof of Theorem 3.4.

[15] Note that Claim 3.4.5 is applied to $W''_{s_1},...,W''_{s_{i-1}}$ and W'_s (rather than to $W'_{s_1},...,W'_{s_{i-1}}$ and W'_s).

[16] Hence, we effectively take a (non-lazy) random walk on a $2d$-regular graph.

4 On the exponential dependence on g in testing Γ_{g+1}

Recall that Theorem 1.5 asserts a $O(d^{\lceil g/2 \rceil})$-query proximity-oblivious tester for the set of graphs having girth greater than g (i.e., Γ_{g+1}). We show that an exponential dependence of the query complexity on the girth is unavoidable.

Theorem 4.1 (lower bound on testing Γ_{g+1}): *For every $d \geq 3$, $\alpha < 1/2$, and $g : \mathbb{N} \to \mathbb{N}$ such that $g(n) \in [3, \alpha \cdot \log_2 n]$, the query complexity of testing Γ_g is $\exp(\Omega(g))$.*

We stress that Theorem 4.1 holds also for two-sided error testers, and is proved for constant proximity parameter (e.g., $\epsilon = (1 - 2\alpha)/d$).

Proof Sketch: The key observation is that a random 3-regular k-vertex graph is "extremely close" to having girth greater than $\alpha \cdot \log_2 k$ but is far from having girth at least $(1/\alpha) \cdot \log_2 k$. In contrast, a $o(\sqrt{k})$-query algorithm that explores a random 3-regular k-vertex graph is unlikely to see a cycle. This means that an algorithm that makes $o(\sqrt{h})$ queries cannot distinguish between a random 3-regular k-vertex graph and a k-vertex graph consisting of k/h connected components each being a random 3-regular h-vertex graph. Using $h = k^{\alpha^2}$, it follows that such an algorithm cannot distinguish between a distribution of graphs that have girth greater than $g = \alpha \cdot \log_2 k$ and a distribution of graphs that are far from having girth greater than $(1/\alpha) \cdot \log_2 h = \alpha \cdot \log_2 k$.

The foregoing description blurs the difference between being "extremely close" to having girth greater than $g = \alpha \cdot \log_2 k$ and having girth greater than g. Intuitively, if being extremely close means being $k^{-1/2}$-close, then omitting the relevant edges yield a graph of girth greater than g, whereas a $o(\sqrt{k})$-query algorithm is unlikely to see evidence of these omission (which result in few vertices having degree smaller than 3). Details follow.

Key facts. As evident from the foregoing overview, the proof relies on facts regarding the distance of a random d-regular k-regular graph from having certain girth (i.e., from the set of d-regular k-regular graphs having certain girth). We start by stating these facts.

***Fact 1*:** The number of (simple) ℓ-cycles in a random d-regular k-vertex graph is approximated by a Poison distribution with mean $(d-1)^\ell/2\ell$, provided that $(d-1)^{2\ell-1} = o(k)$ (cf. [13, p. 2]).

Letting $g = \alpha \cdot \log_2 k$ and noting that $2^g = k^\alpha = o(\sqrt{k})$, it follows that, with probability $1 - o(1)$, a random 3-regular k-vertex graph contains at most $\sqrt{k}$ (simple) cycles of length at most g. In other words, *a random 3-regular k-vertex graph is $k^{-1/2}$-close to having girth at least $g + 1$.*

Hence, omitting $\sqrt{k}$ edges from the graph, we obtain a graph of girth at least $g+1$ that looks like a random 3-regular k-vertex graph to any $o(\sqrt{k})$-query explorer.

***Fact 2*:** Any h-vertex graph of girth g has average degree at most $h^{2/g} + 1$ (see [1, 19]).

Letting $h = k^{\alpha^2}$ and using $g = \alpha \cdot \log_2 k$, it follows that $h^{2/g} = 2^{2\alpha^2/\alpha} < 2$, which means that an h-vertex graph of girth at least g has average degree smaller than $3 - \epsilon$, for some constant $\epsilon > 0$ (which depends on α).

Hence, *every h-vertex 3-regular graph is ϵ-far from having girth at least $g = (1/\alpha) \cdot \log_2 h$.*

Lastly, we observe that the argument in [6, Sec. 7.2] extends to showing that, for every t, distinguishing between a random 3-regular k-vertex graph and a k-vertex graph consisting of t random 3-regular k/t-vertex graphs requires $\Omega(\sqrt{k/t})$ queries. (The original argument is for the case $t = 4$.)

The actual proof. With these preliminaries in place, we turn to the actual proof. For any $g \geq 3$, we let $k = 2^{g/\alpha}$. For any $n \in \{m \cdot k : m \in \mathbb{N}\}$, we consider the following two distributions on n-vertex graphs of degree at most 3.

1. The n-vertex graph consists of n/k connected components that are each a k-vertex graph of degree at most 3 that is generated at random as follows.

   (a) We sample the conditional space of random 3-regular k-vertex graphs that contain at most $\sqrt{k}$ (simple) cycles of length at most g.
   (Recall that (by Fact 1), for $2^g = k^\alpha = o(\sqrt{k})$, with probability $1 - o(1)$, a random 3-regular k-vertex graph contains at most $\sqrt{k}$ (simple) cycles of length at most g.)

   (b) We obtain a k-vertex graph of girth at least $g + 1$ by omitting an arbitrary edge from each of the foregoing cycles.

   Hence, the resulting n-vertex graph has girth at least $g + 1$.

2. The n-vertex graph consists of n/k^{α^2} connected components that are each a random 3-regular k^{α^2}-vertex graph.

   (Recall that (by Fact 2), each connected component is ϵ-far from having girth at least g, because a k^{α^2}-vertex graph of girth g must have average degree at most $(k^{\alpha^2})^{2/g} + 1 = (2^{\alpha \cdot g})^{2/g} + 1 = 2^{2\alpha} + 1$, whereas we can set $\epsilon > 0$ such that $\epsilon \leq 2 - 2^{2\alpha}$.)

   Hence, this n-vertex graph is ϵ-far from having girth at least g.

Recall that by the foregoing observation, distinguishing between a random 3-regular k-vertex graph and a graph consisting of $k^{1-\alpha^2}$ connected components that is each a random 3-regular k^{α^2}-vertex graph requires $\Omega(\sqrt{k^{\alpha^2}}) = \exp(\Omega(g))$ queries. The same holds when trying to distinguish the two foregoing distributions, because the modifications performed (in the first distribution) cannot be noticed by an $O(\sqrt{k^{\alpha^2}})$-query exploration.[17] ■

Acknowledgments

We are grateful to Nati Linial and Merav Parter for a helpful discussions regarding girth.

References

[1] Noga Alon, Shlomo Hoory, and Nati Linial. The Moore Bound for Irregular Graphs. *Graphs and Combinatorics*, Vo. 18, pages 53–57, 2002.

[17] More precisely, we first streamline the second distribution by partitioning the n/k^{α^2} graphs of in the distribution to n/k graphs each consisting of k/k^{α^2} connected components. Letting q_i denote the number of queries made to the i^{th} graph (in one of the two distribution), we observe that the argument in [6, Sec. 7.2] implies that the distinguishability gap is upper-bounded by $\sum_i O(q_i^2/k^{\alpha^2}) = O(q^2/k^{\alpha^2})$, where $q = \sum_i q_i$.

[2] Tugkan Batu and Clement L. Canonne. Generalized Uniformity Testing. In *58th FOCS*, pages 880–889, 2017.

[3] Ilias Diakonikolas, Daniel Kane, and Alistair Stewart. Sharp Bounds for Generalized Uniformity Testing. In *ECCC*, TR17-132, 2017.

[4] Oded Goldreich. *Introduction to Property Testing*. Cambridge University Press, 2017.

[5] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property Testing and its connection to Learning and Approximation. *Journal of the ACM*, Vol. 45 (4), pages 653–750, July 1998.

[6] Oded Goldreich and Dana Ron. Property Testing in Bounded Degree Graphs. *Algorithmica*, Vol. 32 (2), pages 302–343, 2002.

[7] Oded Goldreich and Dana Ron. On Testing Expansion in Bounded-Degree Graphs. In *Studies in Complexity and Cryptography*, pages 68–75, 2011. Preliminary version in *ECCC*, TR00-020, 2000.

[8] Oded Goldreich and Dana Ron. On Proximity Oblivious Testing. *SIAM Journal on Computing*, Vol. 40, No. 2, pages 534–566, 2011.

[9] Oded Goldreich and Laliv Tauber. Testing in the bounded-degree graph model with degree bound two. *ECCC*, TR22-184, 2022.

[10] Mikhail Gromov. Singularities, Expanders and Topology of Maps. Part 2: From Combinatorics to Topology via Algebraic Isoperimetry. *Geom. Funct. Anal.*, Vol. 20 (2), pages 416–526, 2010.

[11] Satyen Kale and C. Seshadhri. An Expansion Tester for Bounded Degree Graphs. *SIAM Journal on Computing*, Vol. 40, No. 3, pages 709–720, 2011.

[12] Akash Kumar, C. Seshadhri, and Andrew Stolman. Random walks and forbidden minors III: poly(d/ϵ)-time partition oracles for minor-free graph classes. In *62nd FOCS*, pages 257–268, 2021.

[13] Brendan D. McKay, Nicholas C. Wormald, and Beata Wysocka. Short Cycles in Random Regular Graphs. *The Electronic Journal of Combinatorics*, Vol. 11, Art. R66, 2004.

[14] Nathan Linial and Roy Meshulam. Homological Connectivity of Random 2-Complexes. *Combinatorica*, Vol. 26 (4), pages 475–487, 2006.

[15] Izhar Oppenheim. Local Spectral Expansion Approach to High Dimensional Expanders Part I: Descent of Spectral Gaps. *Discrete Comput. Geom.*, Vol. 59.2, pages 293–330, 2018.

[16] Aviv Reznik. Finding k-Paths in Cycle-Free Graphs. Master thesis, Weizmann Institute of Science, 2011.
Available from `https://www.wisdom.weizmann.ac.il/~oded/msc-ar.html`.

[17] Ronitt Rubinfeld and Madhu Sudan. Robust Characterization of Polynomials with Applications to Program Testing. *SIAM Journal on Computing*, Vol. 25 (2), pages 252–271, 1996.

[18] Yavin Sabo. On Finding Small Subgraphs in Bounded-Degree Graphs. Master thesis, Tel-Aviv University, 2016.
Available from `https://www.wisdom.weizmann.ac.il/~oded/msc-ar.html`.

[19] Luca Trevisan. The Moore Bound for Irregular Graphs (an exposition).
Available from `https://lucatrevisan.wordpress.com/` (posted on February 2004).

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il