



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/139624/>

Version: Published Version

Proceedings Paper:

Fichtenberger, H., Peng, P. and Sohler, C. (2019) Every testable (infinite) property of bounded-degree graphs contains an infinite hyperfinite subproperty. In: Chan, T., (ed.) Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA19). 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA19), 06-09 Jan 2019, San Diego, California. SIAM. ISBN: 978-1-61197-548-2.

<https://doi.org/10.1137/1.9781611975482.45>

© 2019 The Authors. Reproduced in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Every Testable (Infinite) Property of Bounded-Degree Graphs Contains an Infinite Hyperfinite Subproperty *

Hendrik Fichtenberger [†]

Pan Peng [‡]

Christian Sohler [§]

Abstract

One of the most fundamental questions in graph property testing is to characterize the combinatorial structure of properties that are testable with a constant number of queries. We work towards an answer to this question for the bounded-degree graph model introduced in [GR02], where the input graphs have maximum degree bounded by a constant d . In this model, it is known (among other results) that every *hyperfinite* property is constant-query testable [NS13], where, informally, a graph property is hyperfinite, if for every $\delta > 0$ every graph in the property can be partitioned into small connected components by removing δn edges.

In this paper we show that hyperfiniteness plays a role in *every* testable property, i.e. we show that every testable property is either finite (which trivially implies hyperfiniteness and testability) or contains an infinite hyperfinite subproperty. A simple consequence of our result is that no infinite graph property that only consists of expander graphs is constant-query testable.

Based on the above findings, one could ask if every infinite testable non-hyperfinite property might contain an infinite family of expander (or near-expander) graphs. We show that this is not true. Motivated by our counterexample we develop a theorem that shows that we can partition the set of vertices of every bounded degree graph into a constant number of subsets and a separator set, such that the separator set is small and the distribution of k -discs on every subset of a partition class, is roughly the same as that of the partition class if the subset has small expansion.

1 Introduction

Understanding the structure of very large graphs like social networks or the webgraph is a challenging task. Given the size of these networks, it is often hopeless to compute structural information exactly. A feasible approach is to design random sampling algorithms that only inspect a small portion of the graph and derive conclusions about the structure of the whole

graph from this random sample. However, there are different ways to sample from graphs (random induced subgraphs, random sets of edges, random walks, random BFS, etc.) and also many structural graph properties. This raises the question, which sampling approaches (if any) are suitable to detect or approximate which structural properties.

Graph property testing provides a formal algorithmic framework that allows us to study the above setting from a complexity theory point of view. In this framework, given oracle access to an input graph, our goal is to distinguish between the case that the graph satisfies some property or that it is “far from” having the property by randomly sampling from the graph. Here, a graph property denotes a set of graphs that is invariant under graph isomorphism. Both oracle access and the notion “far from” depend on the representation of the graph. Several models have been proposed in the past two decades for dealing with different types of graphs (see the recent book [Gol17]).

For dense graphs, Goldreich et al. [GGR98] introduced the *adjacency matrix* model, in which the algorithm can perform any *vertex-pair* query to the oracle. That is, upon an input vertex pair u, v , the oracle returns 1 if there is an edge between u, v and 0 otherwise. A graph is called ε -far from having a property Π if one has to modify more than εn^2 vertices to make it satisfy Π for any small constant ε . Since the time when the model was introduced, many properties Π were found to be *testable* in the sense that there exists an algorithm, called *tester*, that can distinguish if a graph satisfies Π or is ε -far from having Π while only making a *constant* number of queries. The research in this model has culminated in the seminal work by Alon et al. [AFNS09], who gave a full characterization of constant-query testable properties by the regularity lemma.

Our understanding of property testing for sparse graphs (e.g., bounded degree graphs) is much more limited. Goldreich and Ron [GR02] initiated the study of property testing for bounded degree graphs in the *adjacency list* model. A graph G is called a *d-bounded* graph if its maximum degree is at most

*First and third author acknowledge the support by ERC grant No. 307696.

[†]Department of Computer Science, TU Dortmund. Email: hendrik.fichtenberger@tu-dortmund.de.

[‡]Department of Computer Science, University of Sheffield. Email: p.peng@sheffield.ac.uk.

[§]Department of Computer Science, TU Dortmund. Email: christian.sohler@tu-dortmund.de.

d , which is assumed to be a constant. The property tester for a d -bounded graph is given oracle access to the adjacency list of the graph, that is, upon an input (u, i) such that $i \leq d$, the oracle returns the i -th neighbor of u if such a neighbor exists, and a special symbol otherwise. A d -bounded graph is said to be ε -far from having the property Π if one needs to modify more than εdn edges to obtain a graph that satisfies Π . In this model, there exist several properties that are known to be testable with a constant number of queries (see discussion below). There also exist a number of properties that require $\tilde{O}(\sqrt{n})$ or $\tilde{O}(n^{1/2+c})$ queries, including bipartiteness [GR99], expansion [GR00, CS10, NS10, KS11], k -clusterability [CPS15] and one-sided error minor-freeness [CGR⁺14, FLVW18, KSS18]. For the property of being 3-colorable there is a known $\Omega(n)$ lower bound on the number of queries needed to test the property [BOT02].

One of the most important questions in this area is to give a purely combinatorial characterization of which graph properties are testable with a constant number of queries. Goldreich and Ron were the first to show that a number of fundamental graph properties including connectivity, k -edge connectivity, subgraph-freeness, cycle-freeness, Eulerian and degree regularity can be tested with constant queries in bounded degree graphs [GR02]. A number of properties with small separators are now known to be testable in a constant number of queries, such as minor closed properties [BSS10, HKNO09], and hyperfinite properties [NS13]. In particular, in the latter work it is proved that every property is constant-query testable in hyperfinite graphs. There are also constant-query properties that are closed under edge insertions, including k -vertex connectivity [YI12], perfect matching [YYI12], sparsity matroid [ITY12] and the supermodular-cut condition [TY15]. Furthermore, there exist global monotone properties¹ that contain expander graphs and can be tested with constant queries, including the property of being subdivision-free [KY13]. There also exist some work on testable properties in some special classes of bounded degree graphs. For example, it is known that every hereditary property² is testable with a constant number of queries in non-expanding d -bounded graphs [CSS09]. A property called δ -robust spectral property is constant-query testable in the class of high-girth graphs [CKSV18].

¹A graph property is called monotone if it is closed under edge deletions.

²A graph property is called hereditary if it is closed under vertex deletions.

However, very little is known about characteristics of all testable properties in general.

1.1 Our Results Although many properties are known to be constant-query testable in bounded degree graphs, our knowledge on characteristics of *all* testable properties is fairly restricted. One prominent example of testable properties is the family of hyperfinite properties [NS13], which includes planar graphs and graphs that exclude any fixed minor (see e.g., [BSS10, HKNO09]). For the statement of our results and the discussion of techniques, we state the definition of hyperfinite graphs at this place.

DEFINITION 1.1. *Let $\varepsilon \in (0, 1]$ and $k \geq 1$. A graph G with maximum degree bounded by d is called (ε, k) -hyperfinite if one can remove at most $\varepsilon d|V(G)|$ edges from G so that each connected component of the resulting graph has at most k vertices. For a function $\rho : \mathbb{R}^+ \rightarrow \mathbb{N}^+$, a graph G is called ρ -hyperfinite if G is $(\varepsilon, \rho(\varepsilon))$ -hyperfinite for every $\varepsilon > 0$. A set (or property) Π of graphs is called ρ -hyperfinite if every graph in Π is ρ -hyperfinite. A set (or property) Π of graphs is called hyperfinite if it is ρ -hyperfinite for some function ρ .*

Also, many testable properties are known that are not hyperfinite. Our main result is that, nevertheless, for infinite properties the existence of an infinite set of hyperfinite graphs in the property is a necessary condition for its constant-query testability (finite properties are trivially hyperfinite). Since some of these testable properties, e.g., subdivision-freeness, contain expander graphs, a hyperfinite subproperty might seem somewhat surprising. (A subproperty of a property Π is a subset of graphs in Π that is also invariant under graph isomorphism.) Indeed, the complement of every non-trivially constant-query testable property also contains hyperfinite graphs, where a property is non-trivially testable if it is testable and there exists an $\varepsilon > 0$ such that there is an infinite number of graphs that are ε -far from Π .

THEOREM 1.1. *Every constant-query testable property Π of bounded-degree graphs is either finite or contains an infinite hyperfinite subproperty. Also, the complement of every non-trivially constant-query testable graph property contains an infinite hyperfinite subproperty.*

To our best knowledge, our theorem gives the first non-trivial result on the combinatorial structure of *every* constant-query testable property in bounded-degree graphs. A direct corollary from our main result is that expansion and the k -clusterability

property are not constant-query testable, as any hyperfinite graph will have many small subsets with small expansion and thus does not satisfy the properties. Indeed, a much stronger lower bound of $\Omega(\sqrt{n})$ on the query complexity for testing these two properties was already known prior to this work [GR00]. However, our result further implies that every infinite intersection of a family of expander graphs with any other property is also not testable.

COROLLARY 1.1. *Let Π be a property that does not contain an infinite hyperfinite subproperty, and let Π' be an arbitrary property such that $\Pi \cap \Pi'$ is an infinite set. Then, $\Pi \cap \Pi'$ is not testable.*

Note that in general, the intersection of a property that is not constant-query testable with another property may be testable. For example, the property of being planar and bipartite is testable since it is a hyperfinite property [NS13]. However, bipartiteness is not constant-query testable [GR02].

We then study the question whether a similar result can be obtained for *expander* or near-expander subproperties in testable non-hyperfinite properties. Expander graphs are those that are well connected everywhere, and thus can be thought as anti-hyperfinite graphs. Indeed, many known testable, while non-hyperfinite, properties do contain infinite expander subproperties. Typical examples include k -connectivity, subgraph-freeness and subdivision-freeness. However, this turns out to not be the case in general. We show that there exists a testable property that is not hyperfinite and every graph in the property has distance $\Omega(n)$ to being an expander graph: The property consists of all graphs that have a connected component on $\lceil |V|/2 \rceil$ vertices and all other vertices are isolated.

THEOREM 1.2. *There exists an infinite graph property Π of bounded-degree graphs such that*

- Π is testable (with 2-sided error) with query complexity $O(d/\varepsilon^2)$,
- Π is not hyperfinite,
- every graph in Π differs in $\Omega(n)$ edges from every connected graph.

Motivated by the above result we also obtain a theorem (Theorem 5.1) that shows that we can partition the set of vertices of every bounded degree graph into a constant number of subsets and a separator set, such that the separator set is small and the distribution of k -discs on every subset of a partition class, is roughly the same as that of the partition class if the subset has small expansion.

1.2 Our Techniques It is well known that constant-time property testing in the bounded-degree graph model is closely connected to the distribution of k -disc isomorphism types (see, for example, [BSS10, NS13]). The k -disc of $v \in V$ is the rooted subgraph that is induced by all vertices at distance at most k from v and has root v , i.e. the local subgraph that can be explored by running a BFS upto depth k . Thus, the distribution of k -disc isomorphism types describes the local structure of the graph. We then show (in Theorem 3.1) that every constant-query property tester can be turned into a canonical tester that is based on approximating the k -disc distribution and decides based on a net over the space of all distribution vectors. Technically, our proof for this result mostly follows an earlier construction of canonical testers introduced in [GR11] (see also [CPS16, MMPS17]).

We then exploit a result by Alon [Lov12, Proposition 19.10] that is derived from open questions in graph limits theory. Alon proved that for every bounded-degree graph G , there exists a graph of constant size H whose k -disc distribution can be made arbitrarily close (in terms of ℓ_1 norm distance) to the k -disc distribution of G . Given a graph G on n vertices from some constant-query testable property Π we can use multiple copies of H to obtain a graph that consists of connected components of constant size and whose distribution of k -discs is close to that of G . The latter implies that a canonical tester will behave similarly on H and G and thus accepts with probability at least $2/3$. Although H does not necessarily have the tested property, it must be close to it. This implies that there exists a graph H' in Π from which we can remove εdn edges to partition it into small connected components. Thus, H' is $(\varepsilon, O_\varepsilon(1))$ -hyperfinite, where $O_\varepsilon(1)$ is a constant depending on ε . However, H' may not be $(\varepsilon', O_{\varepsilon'}(1))$ -hyperfinite for $\varepsilon' < \varepsilon$. The challenge is how to construct such a graph.

In order to do so, we proceed as follows. For every suitable choice of n , we construct a series of n -vertex graphs H_i such that each H_i approximately inherits the $(\varepsilon, O_\varepsilon(1))$ -hyperfinite properties of all graphs $H_{i'}$ for all $i' < i$. The key idea is to maintain the hyperfinite properties of H_i by causing only a small perturbation of its k -disc vector. Carefully choosing the parameters of this process, at the end we obtain a graph $H^{(n)}$ that is $\rho(\varepsilon)$ -hyperfinite for a monotone function $\rho(\cdot)$ and every $\varepsilon > 0$.

In order to show that we cannot obtain a similar result for expander graphs in non-hyperfinite properties, we have designed the aforementioned property of graphs which consist of a connected component

on half of the vertices and all other vertices are isolated. Our proof of testability combines earlier ideas of testing connectivity with simple sampling based estimation of the number of isolated vertices.

1.3 Other Related Work Goldreich and Ron [GR11] gave characterizations of the graph properties that have constant-query proximity-oblivious testers for bounded-degree graphs and for dense graphs. As noted in [GR11], such a class of properties is a rather restricted subset of the class of all constant-query testable properties. Hyperfiniteness is also closely related to *graphings* that have been investigated in the theory of graph limits [Ele07, Sch08, Lov12].

2 Preliminaries

Let $G = (V, E)$ be a graph with maximum degree bounded by d , which is assumed to be a constant. We also call G a d -bounded graph.

DEFINITION 2.1. *A graph property Π is a set of graphs that is invariant under graph isomorphism. If all the graphs in Π have maximum degree upper bounded by d , then we call Π a d -bounded graph property. We let $\Pi_n \subseteq \Pi$ denote the set of graphs in Π with n vertices. Note that $\Pi = \cup_{n \geq 1} \Pi_n$.*

Let $\bar{\Pi}$ denote the complement of Π , i.e., $\bar{\Pi} = \mathcal{U} \setminus \Pi$, where \mathcal{U} denotes the set of all d -bounded graphs. Let $\bar{\Pi}_n$ denote the set of n -vertex graphs that are not in Π_n , i.e., $\bar{\Pi}_n = \mathcal{U}_n \setminus \Pi_n$, where \mathcal{U}_n denotes the set of all d -bounded n -vertex graphs.

A subset $\Pi' \subseteq \Pi$ is called a subproperty of Π if Π' is invariant under graph isomorphism.

We have the following definition on graphs that are far from having some property.

DEFINITION 2.2. *Let $\Pi = \cup_{n \geq 1} \Pi_n$ be a d -bounded graph property. An n -vertex graph is said to be ε -far from having property Π_n if one has to modify more than εdn edges to make it satisfy Π_n .*

Let $\bar{\Pi}_{n, > \varepsilon}$ denote the set of all n -vertex graphs that are ε -far from Π_n . Let $\bar{\Pi}_{> \varepsilon} \subseteq \bar{\Pi}$ be the set of all graphs that are ε -far from Π , i.e., $\bar{\Pi}_{> \varepsilon} = \cup_{n \geq 1} \bar{\Pi}_{n, > \varepsilon}$.

Given a property $\Pi = \cup_{n \geq 1} \Pi_n$, an algorithm is called a *tester* for Π , if it takes as input parameters $0 < \varepsilon \leq 1, n, d$, and has query access to the adjacency lists of an n -vertex d -bounded graph G , and with probability at least $2/3$, accepts G if $G \in \Pi_n$ and rejects G if $G \in \bar{\Pi}_{n, > \varepsilon}$. The following gives the definition of constant-query testable properties.

DEFINITION 2.3. *We call a d -bounded graph property $\Pi = \cup_{n \geq 1} \Pi_n$ (constant-query) testable, if there exists*

a tester for Π that makes at most $q_\Pi = q_\Pi(\varepsilon, d)$ queries for some function $q_\Pi(\cdot, \cdot)$ that depends only on ε, d .

k -Discs and frequency vectors. The notions of k -discs and frequency vectors play an important role for analyzing constant-query testable properties. For any vertex $v \in V$, we let $\text{disc}_k(G, v)$ denote the subgraph rooted at v that is induced by all vertices that are at distance at most k from v . For any two rooted subgraphs H_1, H_2 , we say H_1 is isomorphic to H_2 , denoted by $H_1 \simeq H_2$, if there exists a root-preserving mapping $\Phi : V(H_1) \rightarrow V(H_2)$ such that $(u, v) \in E(H_1)$ if and only if $(\Phi(u), \Phi(v)) \in E(H_2)$. Note that for constant d , the total number of possible non-isomorphic k -discs is also a constant, denoted by $N(d, k)$. Furthermore, we let $\mathcal{T}_k = \{\Delta_1, \dots, \Delta_N\}$ be the set of all isomorphism types of k -discs in any d -bounded graph, where $N = N(d, k)$. Finally, we let $\text{freq}_k(G)$ denote the *frequency vector* of G which is indexed by k -disc types in \mathcal{T}_k such that

$$\text{freq}_k(G)_\Delta = \frac{|\{v \in V : \text{disc}_k(G, v) \simeq \Delta\}|}{n}$$

for any $\Delta \in \mathcal{T}_k$, i.e., $\text{freq}_k(G)_\Delta$ denotes the fraction of vertices in G whose k -discs are isomorphic to Δ . Furthermore, for any subset S of G , we let $\text{freq}_k(S | G)$ denote the vector that is indexed by types in \mathcal{T}_k such that

$$\text{freq}_k(S | G)_\Delta = \frac{|\{v \in S : \text{disc}_k(G, v) \simeq \Delta\}|}{|S|}$$

for any $\Delta \in \mathcal{T}_k$, i.e., $\text{freq}_k(S | G)_\Delta$ denotes the fraction of vertices in S whose k -discs in G are isomorphic to Δ . Note that $\text{freq}_k(G) = \text{freq}_k(V | G)$. If S contains a single element x , we write $\text{freq}_k(x | G) = \text{freq}_k(S | G)$.

For any vector f , we let $\|f\|_1$ denote its ℓ_1 -norm. We have the following simple lemma on the ℓ_1 -norm distance of the frequency vectors of two graphs that are ε -close to each other. The proof follows from the proof of Corollary 3 in [FPS15], while we provide a proof here for the sake of completeness.

LEMMA 2.1. *Let $\varepsilon > 0$ and $k \geq 1$. Let G_1, G_2 be d -bounded graphs such that G_1 is ε -close to G_2 . Then, $\|\text{freq}_k(G_1) - \text{freq}_k(G_2)\|_1 < 6\varepsilon d^{k+1}$.*

Proof. Let $F := E(G_1) \Delta E(G_2)$ denote the set of edges that appear only in one of the two graphs G_1, G_2 . Since G_1 is ε -close to G_2 , it holds that $|F| \leq \varepsilon dn$. Note that for any $e \in F$, the total number of vertices that are within distance at most k to either of its endpoint is at most $2(1 + d + d(d-1) + \dots + d(d-$

$1)^{k-1}) \leq 3d^k$. This further implies that the total number of vertices that may have different k -disc types in G_1 and G_2 is at most $|F| \cdot 3d^k \leq 3\epsilon d^{k+1}n$. Finally, we note that each vertex with different k -disc types in G_1 and G_2 contributes at most $\frac{2}{n}$ to the ℓ_1 -norm distance of $\text{freq}_k(G_1)$ and $\text{freq}_k(G_2)$, which implies that

$$\|\text{freq}_k(G_1) - \text{freq}_k(G_2)\|_1 < 3\epsilon d^{k+1}n \cdot \frac{2}{n} = 6\epsilon d^{k+1}.$$

This completes the proof of the lemma.

The converse to the above lemma is not true in general, that is, it is not true that the closeness of the frequency vectors of two graphs implies the closeness of these two graphs. However, Benjamini et al. [BSS10] showed that the converse somehow still holds for hyperfinite graphs. More precisely, they proved the following result.

LEMMA 2.2. (THEOREM 2.2 FROM [BSS10]) *Let $d, s \geq 1$ and $\epsilon > 0$. Let Λ_1 be the set of (ϵ, s) -hyperfinite d -bounded graphs, and let Λ_2 be the set of d -bounded graphs that are not $(4\epsilon \log(4d/\epsilon), s)$ -hyperfinite. Then it holds that for any graph $G_1 \in \Lambda_1$ and graph $G_2 \in \Lambda_2$,*

$$\|\text{freq}_k(G_1) - \text{freq}_k(G_2)\|_1 > \frac{8\epsilon}{d} \log(4/3),$$

where $k = 10sd^{2s+1}/\epsilon$.

Frequency preservers and blow-up graphs.

The following lemma is due to Alon, and it roughly says that for any n -vertex d -bounded graph, there always exists a “small” graph whose size is independent of n that preserves the local structure well, i.e., its k -disc frequencies.

LEMMA 2.3. (PROPOSITION 19.10 IN [LOV12]) *For any $\delta > 0$ and $d, k \geq 1$, there exists a function $M_d(\delta, k)$ such that for every n -vertex graph G , there exists a graph H of size at most $M_d(\delta, k)$ such that $\|\text{freq}_k(G) - \text{freq}_k(H)\|_1 < \delta$.*

DEFINITION 2.4. ((δ, k)-DFP) *We call the small graph H obtained from Lemma 2.3 a (δ, k) -disc frequency preserver (abbreviated as (δ, k) -DFP) of G .*

We remark that though we know the existence of the function $M_d(\delta, k)$ that upper bounds the size of some (δ, k) -DFP, there is no known explicit bound on $M_d(\delta, k)$ for arbitrary d -bound graphs (see [FPS15] for explicit bounds of $M_d(\delta, k)$ for some special classes of graphs).

We use DFPs as a building block to construct n -vertex graphs that have constant-size connected

components and approximately preserve the k -disc frequencies of a given n -vertex graph G . More precisely, we have the following definition.

DEFINITION 2.5. (BLOW-UP GRAPH) *Let $\delta, k > 0$, and let G be a d -bounded n -vertex graph. Let H be a (δ, k) -DFP of G graph of size $h \leq M_d(\delta, k)$. Let H' be the n -vertex graph that is composed of $\lfloor n/h \rfloor$ disjoint copies of H and $n - h \cdot \lfloor n/h \rfloor$ isolated vertices. We call H' the (δ, k) -blow-up graph of G .*

The following lemma follows directly from the above definition of blow-up graphs and the fact that the blow-up graph contains at most $h \leq M_d(\delta, k)$ isolated vertices.

LEMMA 2.4. *Let $\delta, d, k > 0$. Let $n \geq n_0(\delta, d, k) := 20M_d(\delta, k)/\delta$. Let G be any d -bounded n -vertex graph and let H be the (δ, k) -blow-up graph of G . We have $\|\text{freq}_k(G) - \text{freq}_k(H)\|_1 < 1.1\delta$.*

Expansion and expander graphs. Let $G = (V, E)$ be a d -bounded graph. Let $S \subset V$ be a subset such that $|S| \leq |V|/2$. The *expansion* or *conductance* of set S is defined to be $\phi_G(S) = \frac{e(S, V \setminus S)}{d|S|}$, where $e(S, V \setminus S)$ denotes the number of crossing edges from S to $V \setminus S$. The expansion of G is defined as $\phi(G) := \min_{S: |S| \leq |V|/2} \phi_G(S)$. We call G a ϕ -*expander* if $\phi(G) \geq \phi$. We simply call G an *expander* if G is a ϕ -expander for some universal constant ϕ .

3 Constant-Query Testable Properties and Hyperfinite Properties

In this section, we give the proof of main theorem, i.e., Theorem 1.1. We first give the necessary tools in Section 3.1, and then give the proof of the first part and second part of Theorem 1.1 in Section 3.2 and 3.3, respectively.

3.1 Basic Tools The following is a direct corollary of Lemma 2.2 by Benjamini et al. [BSS10].

LEMMA 3.1. *Let $\epsilon, s > 0$. Let Π be a testable graph property. Suppose there exists a graph $G \in \Pi_n$ that is (ϵ, s) -hyperfinite. Then, every graph $G' \in \Pi_n$ such that $\|\text{freq}_k(G) - \text{freq}_k(G')\|_1 < \frac{8\epsilon}{d} \log(4/3)$ is $(4\epsilon \log \frac{4d}{\epsilon}, s)$ -hyperfinite, where $k = 10sd^{2s+1}/\epsilon$.*

Our second tool is the following characterization of constant-query testable properties by the so-called *canonical tester*. Such a characterization is similar to the previous ones given in [GR11, CPS16] for bounded-degree testable graph properties. The main difference here is that our canonical tester makes decisions based on the frequency vectors, instead of the forbidden subgraphs as considered in the previous

work. We have the following theorem, whose proof is deferred to Section 3.4.

THEOREM 3.1. (CANONICAL TESTER) *Let $\Pi = (\Pi_n)_{n \in \mathbb{N}}$ be a graph property that can be tested with query complexity $q_\Pi(\varepsilon, d)$. Then there exists $t := c \cdot q_\Pi(\varepsilon, d)$ for some constant $c > 1$, $n_1 := n_1(\varepsilon, d)$ such that for any $\varepsilon > 0$, d , $n \geq n_1$, there exists a tester \mathcal{T}_C that*

1. *accepts any d -bounded n -vertex graph G with probability at least $2/3$, if $\min_{G' \in \Pi_n} \|\text{freq}_t(G) - \text{freq}_t(G')\|_1 \leq \frac{1}{12t}$,*
2. *rejects any d -bounded n -vertex graph G with probability at least $2/3$, if $\min_{G' \in \overline{\Pi}_n, > \varepsilon} \|\text{freq}_t(G) - \text{freq}_t(G')\|_1 \leq \frac{1}{12t}$.*

The canonical tester has query complexity $\hat{q}_\Pi(\varepsilon, d) \leq t \cdot d^{t+2}$.

3.2 Infinite Testable Property Contain Infinite Hyperfinite Subproperties We now prove the first part of Theorem 1.1, i. e., every infinite testable property contains an infinite hyperfinite subproperty.

We start by showing that for any fixed ε , and any graph G in a testable property Π , we can find another graph G' such that G' is (ε, s) -hyperfinite and the frequency vectors of G and G' are close.

LEMMA 3.2. *Let $\delta, \varepsilon, k > 0$. Let $\varepsilon' = \min\{\varepsilon, \frac{\delta}{18d^{k+1}}\}$. Let $n \geq n_2(\varepsilon, \delta, d, k)$. Let Π be a testable graph property with query complexity $q_\Pi = q_\Pi(\varepsilon, d)$ and let $G \in \Pi_n$. Then, there exists $G' \in \Pi_n$ such that*

- *G' is $(\varepsilon, M_d(\frac{\delta'}{3}, k'))$ -hyperfinite, and*
- $\|\text{freq}_k(G) - \text{freq}_k(G')\|_1 < \delta$,

where $\delta' = \min\{\delta, \frac{1}{5c \cdot q_\Pi(\varepsilon', d)}\}$ and $k' = \max\{k, c \cdot q_\Pi(\varepsilon', d)\}$ for some constant $c > 1$.

Proof. Let $n_2(\varepsilon, \delta, d, k) = \max\{n_0(\frac{\delta'}{3}, d, k), n_1(\varepsilon', d)\}$, where n_0, n_1 are the numbers in the statements of Lemma 2.4 and Theorem 3.1, respectively. Let $t = c \cdot q_\Pi(\varepsilon', d)$ for the constant $c > 1$ from Theorem 3.1. By definition, it holds that $t \leq k'$. Let H be the $(\frac{\delta'}{3}, k')$ -blow-up graph of G . By Lemma 2.4 and our assumption that $n \geq n_2$, it holds that $\|\text{freq}_{k'}(G) - \text{freq}_{k'}(H)\|_1 \leq \frac{1.1\delta'}{3}$, which implies that

$$(3.1) \|\text{freq}_t(G) - \text{freq}_t(H)\|_1 \leq \frac{1.1\delta'}{3} \leq \frac{1}{12t},$$

as t satisfies that $\frac{1}{5t} \geq \delta'$ and that $t \leq k'$.

Let \mathcal{T}_C be the canonical tester for Π with parameter ε' with corresponding query complexity $t = \hat{q}_\Pi(\varepsilon', d)$. Then by Theorem 3.1, \mathcal{T}_C will accept H with probability at least $2/3$. This implies that H is ε' -close to Π . Let $G' \in \Pi$ such that H is ε' -close to G' . We claim that G' is the graph we are looking for.

First, we show that G' is $(\varepsilon, M_d(\frac{\delta'}{3}, k'))$ -hyperfinite. Recall that by definition, H is composed of $\lfloor n/h \rfloor$ disjoint copies of a graph of size h and $n - h \cdot \lfloor n/h \rfloor$ isolated vertices, where $h \leq M_d(\frac{\delta'}{3}, k')$. This implies that H is $(0, M_d(\frac{\delta'}{3}, k'))$ -hyperfinite. It follows that G' is $(\varepsilon, M_d(\frac{\delta'}{3}, k'))$ -hyperfinite because we can remove at most $\varepsilon' dn \leq \varepsilon dn$ edges from G' to obtain a graph of which all connected components have size at most $M_d(\frac{\delta'}{3}, k')$.

Second, we prove that $\|\text{freq}_k(G) - \text{freq}_k(G')\|_1 \leq \delta$. Note that the bound given by inequality (3.1) implies

$$\|\text{freq}_k(G) - \text{freq}_k(H)\|_1 \leq \frac{1.1\delta'}{3} \leq \frac{1.1\delta}{3},$$

as $k \leq k'$ and $\delta \geq \delta'$. Now since H and G' are ε' -close to each other, by Lemma 2.1, we have that

$$\|\text{freq}_k(H) - \text{freq}_k(G')\|_1 < 6\varepsilon' d^{k+1} \leq \frac{\delta}{3},$$

where the last inequality follows from our setting of parameters. The claim then follows by applying the triangle inequality. This completes the proof of the lemma.

The above lemma only guarantees that for every fixed $\varepsilon > 0$, and graph $G \in \Pi_n$, one can find a graph $G_\varepsilon \in \Pi_n$ that is $(\varepsilon, M_d(\delta', k'))$ -hyperfinite (for δ' and k' as in Lemma 3.2). However, we cannot directly use G_ε to construct an infinite hyperfinite subproperty. Recall that a set Π of graphs is called to be a hyperfinite property if there exists a function $\rho : (0, 1] \rightarrow \mathbb{N}$ such that Π is $(\varepsilon, \rho(\varepsilon))$ -hyperfinite for every $\varepsilon > 0$. Now, for any $\varepsilon' < \varepsilon$, we cannot guarantee that after removing $\varepsilon' dn$ edge from G_ε , one can obtain a graph that is the union of connected components of constant size. Furthermore, it is not guaranteed that $G_{\varepsilon_1} \simeq G_{\varepsilon_2}$ if $\varepsilon_1 \neq \varepsilon_2$.

Our idea of overcoming the above difficulty is to start with the above hyperfinite graph $G_0 := G_\varepsilon \in \Pi_n$ for some fixed $\varepsilon > 0$, and then iteratively construct a sequence of graphs $G_i \in \Pi_n$ with $i \geq 1$ from G_{i-1} . The constructed graph G_{i+1} is guaranteed to inherit hyperfinite properties from G_i . The key idea is to maintain the hyperfinite properties of G_i by causing only a small perturbation of its k -disc vector. Choosing the parameters in this process carefully, we can maintain these hyperfinite properties for

the whole sequence of graphs. Now we give the details in the following lemma. Note that the first part of Theorem 1.1 follows from this lemma.

LEMMA 3.3. *Let Π be an infinite d -bounded graph property that is testable with query complexity $q_{\Pi}(\varepsilon, d)$. Then, there exists $\Pi' \subseteq \Pi$ such that*

- Π' is an infinite subproperty of Π , and
- there exists a monotonically decreasing function $\rho : (0, 1] \rightarrow \mathbb{N}$ such that Π' is $(\varepsilon, \rho(\varepsilon))$ -hyperfinite for every $\varepsilon > 0$.

Proof. Let $X := \{|V| : G = (V, E) \in \Pi\}$ be the set of sizes $|V(G)|$ of graphs G in Π . Since Π is an infinite graph property, it holds that X is also an infinite set. We show there exists a monotonically decreasing function $\rho : (0, 1] \rightarrow \mathbb{N}$ such that for each $n \in X$, we can find a graph $H^{(n)} \in \Pi_n$ that is $(\varepsilon, \rho(\varepsilon))$ -hyperfinite for every $\varepsilon > 0$. This will imply that the set $\Pi' = \{H^{(n)} : n \in X\}$ is an infinite ρ -hyperfinite property, which will then prove the lemma.

Let us now fix an arbitrary $n \in X$ and let $G \in \Pi_n$ be an arbitrary graph in Π_n . We let $\text{FINDHYPER}(G, \delta, \varepsilon, k, \Pi_n)$ denote the graph G' that is obtained by applying Lemma 3.2 on $G \in \Pi_n$ with parameters δ, ε, k . Now we construct $H^{(n)}$ as follows.

Let $\varepsilon_1 = \frac{1}{10}$. Let $\delta_1 = 4\varepsilon_1/d \log(4/3)$ and let $k_1 = 1$. If $n < n_3(d) := n_2(\varepsilon_1, \delta_1, d, k_1) = n_2(\frac{1}{10}, \frac{2}{5d} \log(4/3), d, 1)$, where n_2 is the number given in Lemma 3.2, then we simply let $H^{(n)} = G$, which is a finite graph of size at most n_3 . In the following, we assume that $n \geq n_3$.

Let $G_0 = G$. We start by applying Lemma 3.2 to G_0 with parameters $\delta = \delta_1, \varepsilon = \varepsilon_1$ and $k = k_1$ to obtain a graph G_1 that is (ε_1, s_1) -hyperfinite, where $s_1 := M_d(\frac{\delta_1}{3}, k_1')$ and $\delta_1' = \min\{\delta_1, \frac{1}{5c \cdot q_{\Pi}(\varepsilon_1', d)}\}, k_1' = \max\{k_1, c \cdot q_{\Pi}(\varepsilon_1', d)\}, \varepsilon_1' = \min\{\varepsilon_1, \frac{\delta_1}{18d^{k_1+1}}\}$.

We now iteratively construct a new n -vertex graph G_{i+1} from a graph G_i that is (ε_i, s_i) -hyperfinite, where $s_i := M_d(\frac{\delta_i'}{3}, k_i')$. Let

$$\begin{aligned} \delta_{i+1} &:= \delta_i/2, & \varepsilon_{i+1} &:= \varepsilon_i/2, \\ k_{i+1} &:= \max\{k_i, 10s_i d^{2s_i+1}/\varepsilon_i\}. \end{aligned}$$

We apply Lemma 3.2 to G_i with parameters $\varepsilon = \varepsilon_{i+1}, \delta = \delta_{i+1}$ and $k = k_{i+1}$ to obtain a graph G_{i+1} that is $(\varepsilon_{i+1}, s_{i+1})$ -hyperfinite, where $s_{i+1} := M_d(\frac{\delta_{i+1}'}{3}, k_{i+1}')$, and

$$\begin{aligned} \delta_{i+1}' &= \min\{\delta_{i+1}, \frac{1}{5c \cdot q_{\Pi}(\varepsilon_{i+1}', d)}\}, \\ k_{i+1}' &= \max\{k_{i+1}, c \cdot q_{\Pi}(\varepsilon_{i+1}', d)\}, \end{aligned}$$

$$\varepsilon_{i+1}' = \min\{\varepsilon_{i+1}, \frac{\delta_{i+1}}{18d^{k_{i+1}+1}}\}.$$

Finally, we stop the process after the i' -th iteration such that $\varepsilon_{i'} dn < 1$. We set $H^{(n)} = G_{i'}$. The pseudo-code of the whole process is given in Algorithm 1 (which invokes Algorithm 2 for setting the parameters as a subroutine).

Algorithm 1 Construction of $H^{(n)}$

```

1: procedure CONSTRUCT( $G, \Pi_n$ )
2:    $G_0 \leftarrow G, \varepsilon_1 \leftarrow \frac{1}{10}$ 
3:    $\delta_1 \leftarrow 4\varepsilon_1/d \log(4/3), k_1 \leftarrow 1$ 
4:    $G_1 \leftarrow \text{FINDHYPER}(G_0, \delta_1, \varepsilon_1, k_1, \Pi_n)$ 
5:    $s_1 \leftarrow \text{SETSIZE}(\varepsilon_1, \delta_1, k_1, \Pi_n)$ 
6:    $i \leftarrow 1$ 
7:   while  $\varepsilon_i dn \geq 1$  do
8:      $\varepsilon_{i+1} \leftarrow \varepsilon_i/2, \delta_{i+1} \leftarrow \delta_i/2$ 
9:      $k_{i+1} \leftarrow \max\{k_i, 10s_i d^{2s_i+1}/\varepsilon_i\}$ 
10:     $G_{i+1} \leftarrow \text{FINDHYPER}(G_i, \delta_{i+1}, \varepsilon_{i+1}, k_{i+1}, \Pi_n)$ 
11:     $s_{i+1} \leftarrow \text{SETSIZE}(\varepsilon_{i+1}, \delta_{i+1}, k_{i+1}, \Pi_n)$ 
12:     $i \leftarrow i + 1$ 
13:  return  $H^{(n)} \leftarrow G_i$ 
14: end procedure

```

Algorithm 2 Set the value of s

```

1: procedure SETSIZE( $\varepsilon, \delta, k, \Pi_n$ )
2:    $\varepsilon' \leftarrow \min\{\varepsilon, \frac{\delta}{18d^{k+1}}\}$ 
3:    $\delta' \leftarrow \min\{\delta, \frac{1}{5c \cdot q_{\Pi}(\varepsilon', d)}\}$ 
4:    $k' \leftarrow \max\{k, c \cdot q_{\Pi}(\varepsilon', d)\}$ 
5:    $s \leftarrow M_d(\frac{\delta'}{3}, k')$ 
6:  return  $s$ 
7: end procedure

```

Now we also note that by the construction and Lemma 3.2, it holds that for any $i \geq 0$,

$$\|\text{freq}_{k_{i+1}}(G_{i+1}) - \text{freq}_{k_{i+1}}(G_i)\|_1 < \delta_{i+1}.$$

By noting that $k_j \leq k_{i+1}$ for any $j \leq i + 1$, we have that

$$\begin{aligned} &\|\text{freq}_{k_j}(G_{i+1}) - \text{freq}_{k_j}(G_i)\|_1 \\ &\leq \|\text{freq}_{k_{i+1}}(G_{i+1}) - \text{freq}_{k_{i+1}}(G_i)\|_1 < \delta_{i+1}. \end{aligned}$$

Furthermore, we have the following claim.

CLAIM 3.1. *It holds that $\|\text{freq}_{k_j}(G_{i+1}) - \text{freq}_{k_j}(G_j)\|_1 < 8\varepsilon_j/d \log(4/3)$ for all $j \leq i + 1$.*

Proof. Recall that $\varepsilon_{i+1} = \varepsilon_i/2$ and $\delta_{i+1} = \delta_i/2$ for all $i > 1$. We have

$$\|\text{freq}_{k_j}(G_{i+1}) - \text{freq}_{k_j}(G_j)\|_1$$

$$\begin{aligned} &\leq \sum_{\ell=j}^i \|\text{freq}_{k_j}(G_{\ell+1}) - \text{freq}_{k_j}(G_\ell)\|_1 \\ &\leq \delta_j \sum_{\ell=j}^i \frac{1}{2^{\ell-j}} \\ &\leq 2\delta_j, \end{aligned}$$

where the first inequality follows from the triangle inequality and the second inequality follows from the convergence of the geometric series $\sum_{\ell=0}^{\infty} 2^{-\ell} = 2$. Since $\delta_j = \frac{\delta_1}{2^{j-1}}$, $\varepsilon_j = \frac{\varepsilon_1}{2^{j-1}}$ and $\delta_1 = 4\varepsilon_1/d \log(4/3)$, it holds that $\delta_j = \frac{4\varepsilon_j}{d \log(4/3)}$. This completes the proof of the claim.

Now by the fact that $G_j \in \Pi_n$ is (ε_j, s_j) -hyperfinite, Claim 3.1 and Lemma 3.1, it follows that G_{i+1} is $(4\varepsilon_j \log \frac{4d}{\varepsilon_j}, s_j)$ -hyperfinite, for any $j \leq i + 1$.

In particular, let i' denote the index such that our algorithm outputs $G_{i'}$, i.e., $H^{(n)} = G_{i'}$. For any $\varepsilon > 0$, $j_\varepsilon = \min\{i \mid 1 \leq i \leq i', 4\varepsilon_i \log \frac{4d}{\varepsilon_i} \leq \varepsilon\}$. It is important to note that even though i' might depend on n , the index j_ε is always independent of n , and depends only on ε .

Then we define

$$\rho(\varepsilon) := \max\{n_3(d), s_{j_\varepsilon}\}.$$

By the above analysis, for any $n \in X$ with $n \geq n_3(d)$, we find an n -vertex graph $H^{(n)} \in \Pi_n$ satisfying the following: for any $\varepsilon > 0$, there exists j_ε such that by removing $(4\varepsilon_{j_\varepsilon} \log \frac{4d}{\varepsilon_{j_\varepsilon}}) \cdot dn \leq \varepsilon dn$ edges, one can decompose $H^{(n)}$ into connected components each of which has size at most $s_{j_\varepsilon} \leq \rho(\varepsilon)$. Thus, it holds that $H^{(n)}$ is $(\varepsilon, \rho(\varepsilon))$ -hyperfinite for any $\varepsilon > 0$.

This completes the proof of the lemma.

3.3 Every Complement of a Non-Trivially Testable Property Contains a Hyperfinite Subproperty We now prove the second part of Theorem 1.1, i. e., the complement of every non-trivially testable property contains a hyperfinite subproperty. The formal definition of non-trivially testable property is given as follows.

DEFINITION 3.1. (NON-TRIVIALY TESTABLE) *A graph property Π is non-trivially testable if it is testable and there exists $\varepsilon > 0$ such that the set of graphs that is ε -far from Π is infinite.*

Note that for a property that is *not* non-trivially testable, for any $\varepsilon > 0$, we can always accept all graphs of size $n \geq n_4$, where $n_4 := n_4(\varepsilon)$ is a finite number (that might not be computable) such that there are at most n_4 graphs that are ε -far from having

the property. For graphs of size smaller than n_4 , one can simply read the whole graph to test if the graph satisfies the property or not.

The second part of Theorem 1.1 will follow from the following lemma.

LEMMA 3.4. *The complement of every non-trivially testable d -bounded graph property Π contains an infinite $(0, c)$ -hyperfinite subproperty, where c depends only on Π .*

Proof. Since Π is non-trivially testable, by Definition 3.1, there exists $\varepsilon > 0$ and an infinite set $N \subseteq \mathbb{N}$ such that for every $n \in N$, $\bar{\Pi}_{n, > \varepsilon}$ is non-empty. Let $\varepsilon > 0$ be the largest value such that $\bar{\Pi}_{> \varepsilon}$ contains an infinite number of graphs. Let $\delta = \frac{1}{13t}$, where $t := q_\Pi(\varepsilon, d)$ denotes the query complexity of Π . Let $k = \hat{q}_\Pi(\varepsilon, d) = t^{2t}$. Fix an arbitrary $n \in N$ such that $n \geq n_0$, where $n_0 = n_0(\delta, d, k)$ is the number given in Lemma 2.4. Let $G_n \in \bar{\Pi}_{n, > \varepsilon}$ be an arbitrary graph in $\bar{\Pi}_{n, > \varepsilon}$. Let $H^{(n)}$ be the (δ, k) -blow-up graph of G_n . Note that $H^{(n)}$ is $(0, k)$ -hyperfinite. Now we claim that $H^{(n)} \notin \Pi$.

Assume on the contrary that $H^{(n)} \in \Pi$. By Lemma 2.4, $\|\text{freq}_k(G_n) - \text{freq}_k(H^{(n)})\|_1 \leq 1.1\delta$. Therefore, by Theorem 3.1, the canonical tester for Π accepts G_n with probability at least $2/3$, which is a contradiction to the fact that $G_n \in \bar{\Pi}_{n, > \varepsilon}$. The lemma follows by defining the set $\Pi' := \{H^{(n)} : n \in N\}$ and $c = k = q_\Pi(\varepsilon, d)^{2q_\Pi(\varepsilon, d)}$.

3.4 Proof of Theorem 3.1 In this section, we give the proof sketch of Theorem 3.1. The first part (i.e., the transformations from the original tester \mathcal{T} to the canonical tester \mathcal{T}_C) of the proof follows from the proof of the canonical testers in [GR11, CPS16], and we sketch the main ideas for the sake of completeness. The last part (i.e., how the behaviour of tester \mathcal{T}_C relates to the frequency vector) of the proof differs from previous work and it is tailored to obtain the characterization as stated in the theorem, which in turn will be suitable for our analysis of the structures of constant-query properties.

Proof. [Proof Sketch of Theorem 3.1] Let \mathcal{T} be a tester for Π_n on n -vertex graphs with error probability (reduced to) at most $\frac{1}{24}$. The query complexity of the tester \mathcal{T} will be $t := c \cdot q_\Pi(\varepsilon, d)$ for some constant $c > 1$, where $q_\Pi(\varepsilon, d)$ is the query complexity of the tester for Π with error probability at most $\frac{1}{3}$. We will then transform \mathcal{T} to a canonical tester \mathcal{T}_C in the same way as in the proof of Lemma 3.1 in [CPS16] (see also [GR11]).

Slightly more precisely, we first convert \mathcal{T} into a tester \mathcal{T}_1 that samples random t -discs of the input

graph and answers all of \mathcal{T} 's queries using the corresponding subgraph H . That is, it samples a set S of t vertices and then makes its decision on the basis of the t -discs rooted at vertices in S by using uniformly random ordering of vertices and emulating the execution of \mathcal{T} accordingly on the permuted graph.

Then, we convert \mathcal{T}_1 into a tester \mathcal{T}_2 whose output depends only on the edges and non-edges in the explored subgraph, the ordering of all explored vertices and its own random coins. This can be done by letting \mathcal{T}_2 accept the input graph G with the average probability that \mathcal{T}_1 accepts G over all possible labellings of H with corresponding sequences of queries and answers.

Next, we convert \mathcal{T}_2 into the final tester \mathcal{T}_3 whose output is independent of the ordering of all explored vertices. This can be done by letting \mathcal{T}_3 accept with probability that is equal to the average of all acceptance probabilities of \mathcal{T}_2 over all possible relabellings of vertices in H .

Finally, we convert \mathcal{T}_3 into a tester \mathcal{T}_C that returns the output deterministically according to the unlabeled version of the explored subgraph and its roots. This can be done by letting \mathcal{T}_C accept the input graph if and only if the probability associated with the explored subgraph H is at least $1/2$.

By similar arguments in the proof of Lemma 3.1 in [CPS16], we can show that \mathcal{T}_C is a tester for Π that has error probability at most $1/12$. That is, for each $G \in \Pi_n$, \mathcal{T}_C accepts G with probability at least $1 - \frac{1}{12}$. For any graph $G \in \bar{\Pi}_{n_1; > \varepsilon}$, \mathcal{T}_C rejects G with probability at least $1 - \frac{1}{12}$. Furthermore, note that the query complexity of \mathcal{T}_C is at most $t \cdot d^{t+2}$.

Now if we let $n_1 := 12d^{2t}t^2$, then for any $n \geq n_1$, it holds that with probability at least $1 - \frac{d^{2t}t^2}{n} \geq 1 - \frac{1}{12}$, none of the t sampled t -discs will intersect. That is, with probability $1 - \frac{1}{12}$, the decision of the tester \mathcal{T}_C will only depend on the structure (or the isomorphic types) of the explored t disjoint t -discs.

Let $\delta_C = \frac{1}{12t}$. We now consider the input graph G satisfying that $\min_{G' \in \Pi_n} \|\text{freq}_t(G) - \text{freq}_t(G')\|_1 \leq \delta_C$. Let $G' \in \Pi_n$ denote a graph for which this minimum is attained. Note that there is a bijection $\Phi : V(G) \rightarrow V(G')$ such that $\text{disc}_t(G, v) \not\cong \text{disc}_t(G', \Phi(v))$ for at most a δ_C -fraction of the vertices $v \in V(G)$. Recall that S denotes the sample set. Note that for any vertex v that is sampled independently and uniformly at random, the probability that $\text{disc}_t(G, v) \not\cong \text{disc}_t(G', v)$ is bounded by the total variation distance of $\text{freq}_t(G)$ and $\text{freq}_t(G')$, which is at most $\delta_C/2$ by our assumption. By the union bound, the probability that there exists some vertex $v \in S$ with $\text{disc}_t(G, v) \not\cong \text{disc}_t(G', \Phi(v))$ is at most $|S| \cdot \delta_C \leq t \cdot \frac{1}{12t} \leq \frac{1}{12}$. Since \mathcal{T}_C rejects G' with

probability at most $\frac{1}{12}$ and the probability that there exists some pair of all t sampled t -discs intersecting is at most $\frac{1}{12}$, \mathcal{T}_C rejects G with probability at most $\frac{1}{12} + \frac{1}{12} + \frac{1}{12} = \frac{1}{4}$.

The case when G satisfying that $\min_{G' \in \bar{\Pi}_{n_1; > \varepsilon}} \|\text{freq}_t(G) - \text{freq}_t(G')\|_1 \leq \delta_C$ can be analyzed analogously. In particular, if G satisfies this condition, then \mathcal{T}_C accepts G with probability at most $\frac{1}{12} + \frac{1}{12} + \frac{1}{12} = \frac{1}{4}$.

Therefore, \mathcal{T}_C accepts (resp. rejects) G with probability at least $1 - \frac{1}{4} > \frac{2}{3}$, if $\min_{G' \in \Pi_n} \|\text{freq}_t(G) - \text{freq}_t(G')\|_1 \leq \delta_C$ (resp. if $\min_{G' \in \bar{\Pi}_{n_1; > \varepsilon}} \|\text{freq}_t(G) - \text{freq}_t(G')\|_1 \leq \delta_C$).

This completes the proof of the theorem.

4 Do Testable Non-Hyperfinite Properties Contain Infinitely Many Expanders?

In the light of the previous result, a natural question is whether every testable infinite property that is not hyperfinite must contain an infinite subproperty that consists only of expander graphs or graphs that are close to an expander graph. Unfortunately, such a statement is not true as the aforementioned Theorem 1.2 shows. In the following, we present the proof of Theorem 1.2.

Proof. [Proof of Theorem 1.2] We start by defining the graph property. Π consists of all graphs $G = (V, E)$ with maximum degree d that have a single connected component with $\lceil |V|/2 \rceil$ vertices and the remaining $\lfloor |V|/2 \rfloor$ connected components are isolated vertices. We observe that Π is not hyperfinite as the big connected component may be an expander graph and so it requires to remove $\Omega(n)$ edges to partition it into small connected components. Furthermore, it requires to insert $\Omega(n)$ edges to make the graph connected, which is a necessary condition for having expansion greater than 0. Finally, we show that the property can be tested with query complexity $O(d/\varepsilon^2)$.

The algorithm consists of two stages. In the first stage, we sample $O(1/\varepsilon^2)$ vertices uniformly at random and estimate the number of isolated vertices. We reject, if this number differs from $\lfloor |V|/2 \rfloor$ by more than $\varepsilon|V|/8$. In the second stage, we sample another $O(1/\varepsilon)$ vertices and perform, for every sampled vertex v , a BFS until we have explored the whole connected component of v or we have explored more than $12/\varepsilon$ vertices. We may assume that the graph contains more than, say, $100/\varepsilon$ vertices as otherwise, we can simply query the whole graph. The tester rejects, if it finds a connected component that is not an isolated vertex.

We now prove that the above algorithm (with

proper choice of constants) is a property tester. Our analysis (in particular for the second stage) uses some ideas that were first introduced in an analysis of a connectivity tester in [GR02]. We first show that the tester accepts every $G \in \Pi$. For some sufficiently large constant in the O -notation we obtain by Chernoff bounds that the first stage of the tester approximates with probability at least $9/10$ such that the number of isolated vertices in G with an additive error of $\varepsilon|V|/8$. If this approximation succeeds, the first stage of the tester does not reject. Furthermore, the second stage never rejects a graph $G \in \Pi$. Thus, the tester accepts with probability at least $9/10$. Next consider a graph that is ε -far from Π and begin with the following claim.

CLAIM 4.1. *Let G be ε -far from Π . Then either the number of isolated vertices in G differs by more than $\varepsilon|V|/4$ from $\lfloor |V|/2 \rfloor$ or there are more than $\varepsilon|V|/12$ connected components of size at most $12/\varepsilon$ that are not isolated vertices.*

Proof. Assume that the claim is not true and there is a graph G that is ε -far from Π , the number of isolated vertices in G differs by at most $\varepsilon|V|/4$ from $\lfloor |V|/2 \rfloor$ and there are at most $\varepsilon|V|/12$ connected components of size at most $12/\varepsilon$ that are not isolated vertices. We will argue that in this case, we can modify at most εdn edges to turn G into a graph that has Π , which is a contradiction. We start with the connected components that are not isolated vertices. We can add a single edge to connect two such components. However, we must make sure that we are not violating the degree bound. If both connected components have a vertex of degree at most $d - 1$, we can simply add an edge to connect them. If all vertices of a connected component have degree $d > 1$ then the component contains a cycle. We can remove an edge from the cycle without destroying connectivity. Thus, we need to modify at most 3 edges to connect two connected components. We observe that there are at most $\varepsilon n/12$ connected components of size more than $12/\varepsilon$ and so there are at most $\varepsilon n/6$ connected components that are not isolated vertices. We can create a single connected component out of them by modifying $\varepsilon n/2$ edges. Our previous modifications did not change the number of isolated vertices in G , so it still differs by at most $\varepsilon|V|/12$ from $\lfloor |V|/2 \rfloor$. If there are too many isolated vertices, we can connect each of them to the big connected component with at most 2 edge modifications resulting in at most $\varepsilon n/2$ modifications. If there are too few isolated vertices, we need to disconnect vertices from the big connected component. For this purpose consider a spanning tree T of the connected component. We will

remove a leave of T . This can be done with d edge modifications and does not change connectivity. Thus we can create exactly $\lfloor |V|/2 \rfloor$ isolated vertices using at most $\varepsilon dn/4$ modifications. Overall, the number of modifications is at most εdn , which proves that the graph was not ε -far from Π . A contradiction.

It remains to show that our tester rejects any G that is ε -far from Π . By Claim 4.1 we know that either the number of isolated vertices in G differs by more than $\varepsilon|V|/4$ from $\lfloor |V|/2 \rfloor$ or G has at least $\varepsilon|V|/12$ connected components of size at most $12/\varepsilon$. In the first case, our algorithm rejects with probability at least $9/10$ as it approximates the number of isolated vertices with additive error $\varepsilon|V|/8$ and rejects if the estimate differs by more than $\varepsilon|V|/4$ from $\lfloor |V|/2 \rfloor$. In the second case we observe that for sufficiently large constant in the O -notation with probability at least $9/10$ we sample a connected component of size at most $12/\varepsilon$. In this case our algorithm detects the component and rejects. Thus, with probability at least $9/10$ the algorithm rejects. The query complexity and running time of the algorithm are dominated by the second stage, which can be done in $O(d/\varepsilon^2)$ time.

Since an expander graph is connected, it follows also that this property contains no graphs that are close to expander graphs. Consider the k -discs of graphs from the property Π in the proof of Theorem 1.2. Recall that the graphs from the property consist of a connected graph on $\lfloor |V|/2 \rfloor$ vertices and $\lfloor |V|/2 \rfloor$ isolated vertices. We may view graphs in Π as the union of two graphs G_1 and G_2 of roughly the same size that satisfy two different properties: G_1 is connected and the G_2 has no edges. The k -discs of these graphs have two interesting properties:

- no k -disc in G_1 occurs in G_2 and vice versa, and
- their centers cannot be adjacent in any graph.

If G_1 and G_2 have the above properties then this means that the k -discs cannot “mix” in any connected component of another graph. Thus, we know whether they are supposed to come from G_1 or G_2 , which is helpful to design a property tester. We remark that this phenomenon can also happen for other k -discs like, for example, if G_1 is 4-regular and G_2 is 6-regular. We believe that understanding this phenomenon is important for a characterization of testable properties in bounded-degree graphs as we can use it to construct other testable properties in a similar way as above. This motivates the following definition:

DEFINITION 4.1. We call two k -disc isomorphism types D_1, D_2 with roots u_1, u_2 incompatible, if there exists no graph in which two adjacent vertices u_1 and u_2 have k -disc type D_1 and D_2 , respectively.

5 Partitioning Theorem for Bounded-Degree Graphs

The fact that there are testable properties that are composed of other properties with disjoint sets of incompatible k -discs (see Definition 4.1) leads to the question if we can always decompose the vertex set of a graph into sets such that the k -disc types behave “similarly” within each set. A simple partition would be to divide the vertex set according to its k -disc isomorphism type. But such a partition is meaningless. In the light of previous work, we decided to consider the case that a partition has to have only a small fraction of the edges between the partition classes. We would like to obtain a partition into sets S_1, \dots, S_r and a set T (which is a separator), such that no edges are between S_i and S_j for any $i \neq j$ and T is of small size. The next question is to specify what it means to behave “similarly”. One such specification is to ask that the k -disc distribution inside the partition is stable for every subset. Obviously, this cannot always be the case unless there is only one k -disc isomorphism type. Instead, we are only looking at sets that do not have too many outgoing edges. For these subsets we can show that they always have roughly the same k -disc distribution as their partition. The formal theorem we prove is the following.

THEOREM 5.1. Let $G = (V, E)$ be a d -bounded graph. For every $k \geq 0$ and every $1 \geq \delta > 0$ the vertex set V can be partitioned into $r \leq f(\delta, d, k)$ subsets S_1, \dots, S_r and a set T such that

- for every $i \neq j$ there are no edges between S_i and S_j ,
- $|T| \leq \delta d|V|$,
- and for every i and every subset X of S_i with $\phi_G(X) \leq \delta^2$ it holds that

$$\|\text{freq}_k(X | G) - \text{freq}_k(S_i | G)\|_1 \leq 3\delta.$$

Proof. We will first construct a partition of V into sets A_1, \dots, A_t for some (possibly very large) value of t and a set T such that $|T| \leq \delta d|V|$ and such that there are no edges between any pair of A_i and A_j . Then we construct each set S_i as a union of some of the sets A_j . Finally, we prove that the S_i satisfy the third property (the first two follow from the construction of the A_j).

We start with $T = \emptyset$ and $W = V$. Let A be a subset of vertices of W with $\phi_G(A) \leq \delta$. We may assume that A contains no proper subset with this property (otherwise, we take this subset). We put the neighbors of vertices from A that are not in A into the set T and remove T from W . We store the set A as A_1 and remove it from W . We then repeat this process as long as possible computing the sets A_2, A_3, \dots . We observe that every vertex is removed at most once from W . Whenever we remove a set A_i we move at most $\delta d|A_i|$ neighbors into T since $\phi_G(A_i) \leq \delta$. Hence, $|T| \leq \delta d|V|$. Furthermore, we observe that by construction there are no edges between A_i and A_j for any $i \neq j$.

It remains to construct the sets S_i . For this purpose, we put a δ -net over the space of all k -disc frequency vectors, i.e. we compute a smallest set $N = \{v_1, \dots, v_{|N|}\}$ of frequency vectors such that every frequency vector there exists a vector in N within l_1 distance at most δ . We observe that $|N|$ is a function of k, d and δ . We then define S_i to be the union of all A_j that have v_i as the closest vector to their frequency vector. It remains to prove that the S_i satisfy the third property for δ^2 . For this purpose consider an arbitrary subset $X \subseteq S_i$. We consider $X \cap A_j$ for the sets A_j whose union S_i is. If $X \cap A_j \neq \emptyset$ then we know that $\phi_G(X \cap A_j) > \delta$. Recall that the edges that leave $X \cap A_j$ either go to $A_j \setminus X$ or to T , where $X \cap T = \emptyset$. If $\phi_G(X) \leq \delta^2$, then it holds that at most a δ -fraction of the elements from X can be from a subset A_j with $\phi_G(X \cap A_j) > \delta$. This is true as otherwise the number of edges crossing X and $V \setminus X$ is at least $\delta|X| \cdot \delta d$, which contradicts the assumption that $\phi_G(X) \leq \delta^2$. Let J be the set of all indices j such that $A_j \cap X = \emptyset$. Hence we get

$$\begin{aligned} & \text{freq}_k(X | G) \\ &= \sum_j \sum_{x \in X \cap A_j} \frac{\text{freq}_k(x | G)}{|X|} \\ &= \frac{1}{|X|} \left(\sum_{j \in J} \sum_{x \in X \cap A_j} \text{freq}_k(x | G) \right. \\ & \quad \left. + \sum_{j \notin J} \sum_{x \in X \cap A_j} \text{freq}_k(x | G) \right). \end{aligned}$$

Now let us define $X_1 = \{x \in X | x \in A_j, j \in J\}$ and $X_2 = X \setminus X_1$. We know that $|X_2| \leq \delta|X|$. We also observe that

$$\left\| \frac{1}{|X_1|} \sum_{x \in X_1} \text{freq}_k(x | G) - \text{freq}_k(S_i | G) \right\|_1 \leq \delta$$

and

$$\left\| \frac{1}{|X_2|} \sum_{x \in X_2} \text{freq}_k(x | G) - \text{freq}_k(S_i | G) \right\|_1 \leq 2$$

since all frequency vectors have l_1 -norm 1. It follows that

$$\begin{aligned} & \left\| \text{freq}_k(X | G) - \text{freq}_k(S_i | G) \right\|_1 \\ &= \left\| \frac{1}{|X|} \cdot \left(\sum_{x \in X_1} \text{freq}_k(x | G) \right. \right. \\ &\quad \left. \left. + \sum_{x \in X_2} \text{freq}_k(x | G) \right) - \text{freq}_k(S_i | G) \right\|_1 \\ &= \left\| \frac{1}{|X|} \cdot \left(\sum_{x \in X_1} \text{freq}_k(x | G) - |X_1| \cdot \text{freq}_k(S_i | G) \right. \right. \\ &\quad \left. \left. + \sum_{x \in X_2} \text{freq}_k(x | G) - |X_2| \cdot \text{freq}_k(S_i | G) \right) \right\|_1 \\ &\leq \frac{|X_1|}{|X|} \cdot \delta + \frac{|X_2|}{|X|} \cdot 2 \\ &\leq 3\delta. \end{aligned}$$

This completes the proof of the theorem.

6 Conclusions

We have shown that every constant-time testable property in the bounded-degree graph model is either finite or contains an infinite hyperfinite subproperty. We hope that this result is a first step to obtain a full characterization of all testable properties in bounded-degree graphs. Unfortunately, a similar result cannot be derived for expander graphs, i.e. it is not true that every testable infinite property that is not hyperfinite contains an infinite family of expander graphs or graphs that are close to expander graphs. The structure of this counter-example motivated us to study partitionings of bounded-degree graphs into sets of vertices such that the distribution of k -discs on any subset with bounded expansion is close to the distribution of the set. We hope that this partitioning will be helpful to make further progress towards a characterization of all testable properties in bounded-degree graphs.

References

[AFNS09] Noga Alon, Eldar Fischer, Ilan Newman, and Asaf Shapira. A combinatorial characterization of the testable graph properties: it's all about regularity. *SIAM Journal on Computing*, 39(1):143–167, 2009.

[BOT02] Andrej Bogdanov, Kenji Obata, and Luca Trevisan. A lower bound for testing 3-colorability in bounded-degree graphs. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 93–102. IEEE, 2002.

[BSS10] Itai Benjamini, Oded Schramm, and Asaf Shapira. Every minor-closed property of sparse graphs is testable. *Advances in mathematics*, 223(6):2200–2218, 2010.

[CGR⁺14] Artur Czumaj, Oded Goldreich, Dana Ron, C Seshadhri, Asaf Shapira, and Christian Sohler. Finding cycles and trees in sublinear time. *Random Structures & Algorithms*, 45(2):139–184, 2014.

[CKSV18] David Cohen-Steiner, Weihao Kong, Christian Sohler, and Gregory Valiant. Approximating the spectrum of a graph. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19–23, 2018*, pages 1263–1271, 2018.

[CPS15] Artur Czumaj, Pan Peng, and Christian Sohler. Testing cluster structure of graphs. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC)*, pages 723–732. ACM, 2015.

[CPS16] Artur Czumaj, Pan Peng, and Christian Sohler. Relating two property testing models for bounded degree directed graphs. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1033–1045. ACM, 2016.

[CS10] Artur Czumaj and Christian Sohler. Testing expansion in bounded-degree graphs. *Combinatorics, Probability and Computing*, 19(5-6):693–709, 2010.

[CSS09] Artur Czumaj, Asaf Shapira, and Christian Sohler. Testing hereditary properties of nonexpanding bounded-degree graphs. *SIAM Journal on Computing*, 38(6):2499–2510, 2009.

[Ele07] Gábor Elek. Note on limits of finite graphs. *Combinatorica*, 27(4):503–507, 2007.

[FLVW18] Hendrik Fichtenberger, Reut Levi, Yadu Vasudev, and Maximilian Wötzel. A sublinear tester for outerplanarity (and other forbidden minors) with one-sided error. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107, pages 52:1–52:14, 2018.

[FPS15] Hendrik Fichtenberger, Pan Peng, and Christian Sohler. On Constant-Size Graphs That Preserve the Local Structure of High-Girth Graphs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015)*, volume 40, pages 786–799. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015.

[GGR98] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45(4):653–750, 1998.

[Gol17] Oded Goldreich. *Introduction to Property Testing*. Cambridge University Press, 2017.

[GR99] Oded Goldreich and Dana Ron. A sublinear bipartiteness tester for bounded degree graphs. *Combinatorica*, 19(3):335–373, 1999.

[GR00] Oded Goldreich and Dana Ron. On testing expansion in bounded-degree graphs. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 7, 2000.

- [GR02] Oded Goldreich and Dana Ron. Property testing in bounded degree graphs. *Algorithmica*, 32:302–343, 2002.
- [GR11] Oded Goldreich and Dana Ron. On proximity-oblivious testing. *SIAM Journal on Computing*, 40(2):534–566, 2011.
- [HKNO09] Avinatan Hassidim, Jonathan A Kelner, Huy N Nguyen, and Krzysztof Onak. Local graph partitions for approximation and testing. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 22–31. IEEE, 2009.
- [ITY12] Hiro Ito, Shin-Ichi Tanigawa, and Yuichi Yoshida. Constant-time algorithms for sparsity matroids. In *International Colloquium on Automata, Languages, and Programming*, pages 498–509. Springer, 2012.
- [KS11] Satyen Kale and C. Seshadhri. An expansion tester for bounded degree graphs. *SIAM Journal on Computing*, 40(3):709–720, 2011.
- [KSS18] Akash Kumar, C. Seshadhri, and Andrew Stolman. Finding forbidden minors in sublinear time: an $n^{1/2+o(1)}$ -query one-sided tester for minor closed properties. In *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2018.
- [KY13] Ken-ichi Kawarabayashi and Yuichi Yoshida. Testing subdivision-freeness: property testing meets structural graph theory. In *Proceedings of the 45th annual ACM symposium on Theory of computing (STOC)*, pages 437–446. ACM, 2013.
- [Lov12] László Lovász. *Large networks and graph limits*, volume 60. American Mathematical Society, 2012.
- [MMPS17] Morteza Monemizadeh, Shan Muthukrishnan, Pan Peng, and Christian Sohler. Testable bounded degree graph properties are random order streamable. In *44th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 131:1–131:14, 2017.
- [NS10] Asaf Nachmias and Asaf Shapira. Testing the expansion of a graph. *Information and Computation*, 208(4):309–314, 2010.
- [NS13] Ilan Newman and Christian Sohler. Every property of hyperfinite graphs is testable. *SIAM Journal on Computing*, 42(3):1095–1112, 2013.
- [Sch08] Oded Schramm. Hyperfinite graph limits. *Electronic Research Announcements*, 15:17–23, 2008.
- [TY15] Shin-Ichi Tanigawa and Yuichi Yoshida. Testing the supermodular-cut condition. *Algorithmica*, 71(4):1065–1075, 2015.
- [YI12] Yuichi Yoshida and Hiro Ito. Property testing on k -vertex-connectivity of graphs. *Algorithmica*, 62(3-4):701–712, 2012.
- [YYI12] Yuichi Yoshida, Masaki Yamamoto, and Hiro Ito. Improved constant-time approximation algorithms for maximum matchings and other optimization problems. *SIAM Journal on Computing*, 41(4):1074–1093, 2012.