

Faster Property Testers in a Variation of the Bounded Degree Model

Isolde Adler 

University of Leeds, School of Computing, UK
i.m.adler@leeds.ac.uk

Polly Fahey 

University of Leeds, School of Computing, UK
mm11pf@leeds.ac.uk

Abstract

Property testing algorithms are highly efficient algorithms, that come with probabilistic accuracy guarantees. For a property P , the goal is to distinguish inputs that have P from those that are *far* from having P with high probability correctly, by querying only a small number of local parts of the input. In property testing on graphs, the *distance* is measured by the number of edge modifications (additions or deletions), that are necessary to transform a graph into one with property P . Much research has focussed on the *query complexity* of such algorithms, i. e. the number of queries the algorithm makes to the input, but in view of applications, the *running time* of the algorithm is equally relevant.

In (Adler, Harwath STACS 2018), a natural extension of the bounded degree graph model of property testing to relational databases of bounded degree was introduced, and it was shown that on databases of bounded degree and bounded tree-width, every property that is expressible in monadic second-order logic with counting (CMSO) is testable with constant query complexity and *sublinear* running time. It remains open whether this can be improved to constant running time.

In this paper we introduce a new model, which is based on the bounded degree model, but the distance measure allows both edge (tuple) modifications and vertex (element) modifications. Our main theorem shows that on databases of bounded degree and bounded tree-width, every property that is expressible in CMSO is testable with constant query complexity and *constant* running time in the new model. We also show that every property that is testable in the classical model is testable in our model with the same query complexity and running time, but the converse is not true.

We argue that our model is natural and our meta-theorem showing constant-time CMSO testability supports this.

2012 ACM Subject Classification Theory of computation → Streaming, sublinear and near linear time algorithms; Theory of computation → Database query processing and optimization (theory)

Keywords and phrases Constant Time Algorithms, Logic and Databases, Property Testing, Bounded Degree Model

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2020.7

1 Introduction

Extracting information from large amounts of data and understanding its global structure can be an immensely challenging and time consuming task. When the input data is huge, many traditionally “efficient” algorithms are no longer practical. The framework of property testing aims at addressing this problem. Property testing algorithms (*testers*, for short) are given oracle access to the inputs, and their goal is to distinguish between inputs which have a given property \mathbf{P} or are structurally *far* from having \mathbf{P} with high probability correctly. This can be seen as a relaxation of the classical yes/no decision problem for \mathbf{P} . Testers make these decisions by exploring only a small number of local parts of the input which are randomly chosen. They come with probabilistic guarantees on the quality of the answer.



© Isolde Adler and Polly Fahey;

licensed under Creative Commons License CC-BY

40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2020).

Editors: Nitin Saxena and Sunil Simon; Article No. 7; pp. 7:1–7:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Typically, only a constant number of small local parts are explored and the algorithms often run in constant or sublinear time. This speed up in running time, whilst sacrificing some accuracy, can be crucial for dealing with large inputs. In particular it can be useful for a quick exploration of newly obtained data (e.g. biological networks). Based on the outcome of the exploration, a decision can then be taken whether to use a more time consuming exact algorithm in a second step.

A *property* is simply an isomorphism-closed class of graphs or relational databases. For example, each Boolean database query q defines a property \mathbf{P}_q , the class of all databases satisfying q . In the bounded degree graph model [16], a uniform upper bound d on the degree of the graphs is assumed. For a small $\epsilon \in (0, 1]$, two graphs \mathcal{G} and \mathcal{H} , both on n vertices, are ϵ -close, if at most ϵdn edge modifications (deletions or insertions in \mathcal{G} or \mathcal{H}) are necessary to make \mathcal{G} and \mathcal{H} isomorphic. If \mathcal{G} and \mathcal{H} are not ϵ -close, then they are called ϵ -far. A graph \mathcal{G} is called ϵ -close to a property \mathbf{P} , if \mathcal{G} is ϵ -close to a member of \mathbf{P} , and \mathcal{G} is ϵ -far from \mathbf{P} otherwise. The natural generalisation of this model to relational databases of bounded degree (where a database has degree at most d if each element in its domain appears in at most d tuples) was studied in [1], where two databases \mathcal{D} and \mathcal{D}' , both with n elements in the domain, are ϵ -close, if at most ϵdn tuple modifications (deletions from relations or insertions to relations) are necessary to make \mathcal{D} and \mathcal{D}' isomorphic, and \mathcal{D} and \mathcal{D}' are ϵ -far otherwise. We call this model for bounded degree relational databases the BDRD model.

Our contributions. In this paper we propose a new model for property testing on bounded degree relational databases, which we call the $\text{BDRD}_{+/-}$ model, with a distance measure that allows both tuple deletions and insertions, and *deletion and insertion of elements of the domain*. On graphs, this translates to edge insertions and deletions, and *vertex insertions and deletions*. We argue that this yields a natural distance measure. Indeed, take any (sufficiently large) graph \mathcal{G} , and let \mathcal{H} be obtained from \mathcal{G} by adding an isolated vertex. Then \mathcal{G} and \mathcal{H} are ϵ -far for every $\epsilon \in (0, 1]$ under the classical distance measure, although they only differ in one vertex. In contrast, our distance measure allows for a small number of vertex modifications. While comparing graphs on different numbers of vertices by adding isolated vertices was done implicitly as part of the study the testability of outerplanar graphs [4], to the best of our knowledge, such a distance measure has not been considered before as part of a model in property testing, which seems surprising to us.

Formally, in the $\text{BDRD}_{+/-}$ model, two databases \mathcal{D} and \mathcal{D}' are ϵ -close, if they can be made isomorphic by at most ϵdn modifications, where a modification is either, (1) removing a tuple from a relation, (2) inserting a tuple to a relation, (3) removing an element from the domain (and, as a consequence, any tuple containing that element is removed), or (4) inserting an element into the domain. Here n is the minimum of the sizes of the domains of \mathcal{D} and \mathcal{D}' . In Section 3 we give the full details of our model. We note that the $\text{BDRD}_{+/-}$ model differs from the BDRD model only in the choice of the distance measure. While we work in the setting of relational databases, we would like to emphasize that our results carry over to (undirected and directed) graphs, as these can be seen as special instances of relational databases.

It is known that in the bounded degree graph model, every minor-closed property is testable [6], and, more generally, every hyperfinite graph property is testable [23] with constant query complexity. However, no bound on the running time can be obtained in these general settings. Indeed, there exist hyperfinite properties (of edgeless graphs) that are uncomputable. In [1], Adler and Harwath ask which conditions guarantee both low query complexity *and* efficient running time. They prove a meta-theorem stating that, on classes of

databases (or graphs) of bounded degree and bounded tree-width, every property that can be expressed by a sentence of monadic second-order logic with counting (CMSO) is testable with *constant* query complexity and *polylogarithmic* running time in the BDRD model. Treating many algorithmic problems simultaneously, this can be seen as an algorithmic *meta-theorem* within the line of research inspired by Courcelle’s famous theorem [9] that states that each property of relational databases which is definable in CMSO is decidable in linear time on relational databases of bounded tree-width. CMSO extends first-order logic (FO) and hence properties expressible in FO (e.g. subgraph/sub-database freeness) are also expressible in CMSO. Other examples of graph properties expressible in CMSO include bipartiteness, colourability, even-hole-freeness and Hamiltonicity. Rigidity (i. e. the absence of a non-trivial automorphism) cannot be expressed in CMSO (cf. [10] for more details).

Our main theorem (Theorem 17) shows that in the $\text{BDRD}_{+/-}$ model, on classes of databases (or graphs) of bounded degree and bounded tree-width, every property that can be expressed by a sentence of monadic second-order logic with counting (CMSO) is testable with *constant* query complexity and *constant* running time. The question whether constant running time can also be achieved in the BDRD model remains open.

We show that the $\text{BDRD}_{+/-}$ model is in fact stronger than the BDRD model: Any property testable in the BDRD model is also testable in the $\text{BDRD}_{+/-}$ model with the same query complexity and running time (Lemma 4), but there are examples that show that the converse is not true (Lemma 6).

In the future, it would be interesting to obtain a characterisation of the properties that are (efficiently) testable in the $\text{BDRD}_{+/-}$ model.

Our techniques. To prove our main theorem, we give a general condition under which properties are testable in constant time in the $\text{BDRD}_{+/-}$ model whereas the fastest known testers for such properties in the BDRD model run in polylogarithmic time. To describe this condition let us first briefly introduce some definitions. A property \mathbf{P} is *hyperfinite* on a class of databases \mathbf{C} if every database in \mathbf{P} can be partitioned into connected components of constant size by removing only a constant fraction of the tuples such that the resulting partitioned database is in \mathbf{C} . Let $r \in \mathbb{N}$, given an element a in the domain of a database \mathcal{D} the *r -neighbourhood type* of a in \mathcal{D} is the isomorphism type of the sub-database of \mathcal{D} induced by all elements that are at distance at most r from a in the underlying graph of \mathcal{D} , expanded by a . The *r -histogram* of a bounded degree database \mathcal{D} , denoted by $h_r(\mathcal{D})$, is a vector indexed by the r -neighbourhood types, where the component corresponding to the r -neighbourhood type τ contains the number of elements in \mathcal{D} that realise τ . The *r -neighbourhood distribution* of \mathcal{D} is the vector $h_r(\mathcal{D})/n$ where \mathcal{D} is on n elements. We show that for any property \mathbf{P} and input class \mathbf{C} , if \mathbf{P} is hyperfinite on \mathbf{C} and the set of r -histograms of the databases in \mathbf{P} are semilinear, then \mathbf{P} is testable on \mathbf{C} in constant time (Theorem 16). As a corollary we then obtain our main theorem, that every property definable by a CMSO sentence is testable on the class of databases with bounded degree and bounded tree-width in constant time (Theorem 17).

Alon [22, Proposition 19.10] proved that for every bounded degree graph \mathcal{G} there exists a constant size graph \mathcal{H} that has a similar neighbourhood distribution to \mathcal{G} . However, the proof is based on a compactness argument and does not give an explicit upper bound on the size of \mathcal{H} . Finding such a bound was suggested by Alon as an open problem [18]. We ask under which conditions on a given property \mathbf{P} , for every member of \mathbf{P} there exists a constant size database with a similar neighbourhood distribution which is also in \mathbf{P} . We show that for any property \mathbf{P} which is hyperfinite on the input class \mathbf{C} and whose r -histograms are

semilinear, if a database \mathcal{D} is in \mathbf{P} then there exists a constant size database \mathcal{D}' in \mathbf{P} with a similar neighbourhood distribution but this is not true for databases in \mathbf{C} that are far from \mathbf{P} . Furthermore, we obtain upper and lower bounds on the size of \mathcal{D}' . We can then use this result to construct constant time testers. We first use the algorithm `EstimateFrequencies r,s` (given in [23] and adapted to databases in [1]) to approximate the neighbourhood distribution of the input database. Then we only have to check if the estimated distribution is close to the neighbourhood distribution of a constant size database in the property.

As a corollary (Corollary 14), we obtain an explicit bound on the size on graphs \mathcal{H} from Alon’s theorem for “semilinear” properties, i. e. properties, where the histogram vectors of the neighbourhood distributions form a semilinear set.

Further related work. Other than the work already mentioned in [1] there are only a handful of results on relational databases that utilise models from property testing. Chen and Yoshida [8] study a model which is close to the general graph model (cf. e. g. [2]) in which they study the testability of homomorphism inadmissibility. Ben-Moshe et al. [5] study the testability of near-sortedness (a property of relations that states that most tuples are close to their place in some desired order). Our model differs from both of these, as it relies on a degree bound and uses different types of oracle access. Explicit bounds for Alon’s theorem restricted to high-girth graphs were given in [12].

Obtaining a characterisation of constant query testable properties is a long-standing open problem. Ito et al. [19] give a characterisation of the 1-sided error constant query testable monotone and hereditary graph properties in the bounded degree (directed and undirected) graph model. Fichtenberger et al. [13] show that every constant query testable property in the bounded degree graph model is either finite or contains an infinite hyperfinite subproperty.

Organisation. In Section 2 we introduce relevant notions used throughout the paper. In Section 3 we introduce our property testing model for bounded degree relational databases and we compare it to the classical model. In Section 4 we prove our main theorems. Due to space constraints the proofs of statements labelled (*) are deferred to the appendix.

2 Preliminaries

We let \mathbb{N} be the set of natural numbers including 0, and $\mathbb{N}_{\geq 1} = \mathbb{N} \setminus \{0\}$. For each $n \in \mathbb{N}_{\geq 1}$, we let $[n] = \{1, 2, \dots, n\}$.

Databases. A *schema* is a finite set $\sigma = \{R_1, \dots, R_{|\sigma|}\}$ of relation names, where each $R \in \sigma$ has an *arity* $\text{ar}(R) \in \mathbb{N}_{\geq 1}$. A *database* \mathcal{D} of schema σ (σ -db for short) is of the form $\mathcal{D} = (D, R_1^{\mathcal{D}}, \dots, R_{|\sigma|}^{\mathcal{D}})$, where D is a finite set, the set of *elements* of \mathcal{D} , and $R_i^{\mathcal{D}}$ is an $\text{ar}(R_i)$ -ary relation on D . The set D is also called the *domain* of \mathcal{D} . An (*undirected*) *graph* \mathcal{G} is a tuple $\mathcal{G} = (V(\mathcal{G}), E(\mathcal{G}))$ where $V(\mathcal{G})$ is a set of *vertices* and $E(\mathcal{G})$ is a set of 2-element subsets of $V(\mathcal{G})$ (the *edges* of \mathcal{G}). An undirected graph can be seen as a $\{E\}$ -db, where E is a binary relation name, interpreted by a symmetric, irreflexive relation.

We assume that all databases are linearly ordered or, equivalently, that $D = [n]$ for some $n \in \mathbb{N}$ (similar to [20]). We extend this linear ordering to a linear order on the relations of \mathcal{D} via lexicographic ordering. The *Gaiifman graph* of a σ -db \mathcal{D} is the undirected graph $\mathcal{G}(\mathcal{D}) = (V, E)$, with vertex set $V := D$ and an edge between vertices a and b whenever $a \neq b$ and there is an $R \in \sigma$ and a tuple $(a_1, \dots, a_{\text{ar}(R)}) \in R^{\mathcal{D}}$ with $a, b \in \{a_1, \dots, a_{\text{ar}(R)}\}$. The

degree $\deg(a)$ of an element a in a database \mathcal{D} is the total number of tuples in all relations of \mathcal{D} that contain a . We say the *degree* $\deg(\mathcal{D})$ of a database \mathcal{D} is the maximum degree of its elements. A class of databases \mathbf{C} has *bounded degree*, if there exists a constant $d \in \mathbb{N}$ such that for all $\mathcal{D} \in \mathbf{C}$, $\deg(\mathcal{D}) \leq d$. (We always assume that classes of databases are closed under isomorphism.) Let us remark that the $\deg(\mathcal{D})$ and the (graph-theoretic) degree of $\mathcal{G}(\mathcal{D})$ only differ by at most a constant factor (cf. e.g. [11]). Hence both measures yield the same classes of relational structures of bounded degree. We define the *tree-width* of a database \mathcal{D} as the tree-width of its Gaifman graph. (See e.g. [15] for a discussion of tree-width in this context.) A class \mathbf{C} of databases has *bounded tree-width*, if there exists a constant $t \in \mathbb{N}$ such that all databases $\mathcal{D} \in \mathbf{C}$ have tree-width at most t . Let \mathcal{D} be a σ -db, and $M \subseteq D$. The sub-database of \mathcal{D} *induced by* M is the database $\mathcal{D}[M]$ with domain M and $R^{\mathcal{D}[M]} := R^{\mathcal{D}} \cap M^{\text{ar}(R)}$ for every $R \in \sigma$. An (ϵ, k) -*partition* of a σ -db \mathcal{D} on n elements is a σ -db \mathcal{D}' formed by removing at most ϵn many tuples from \mathcal{D} such that every connected component in \mathcal{D}' contains at most k elements. A class of σ -dbs $\mathbf{C} \subseteq \mathbf{D}$ is ρ -*hyperfinite* on \mathbf{D} if for every $\epsilon \in (0, 1]$ and $\mathcal{D} \in \mathbf{C}$ there exists an $(\epsilon, \rho(\epsilon))$ -partition $\mathcal{D}' \in \mathbf{D}$ of \mathcal{D} . We call \mathbf{C} *hyperfinite* on \mathbf{D} if there exists a function ρ such that \mathbf{C} is ρ -hyperfinite on \mathbf{D} .

Logics. We shall only briefly introduce first-order logic (FO) and monadic second-order logic with counting (CMSO). Detailed introductions can be found in [21] and [10]. Let \mathbf{var} be a countable infinite set of *variables*, and fix a relational schema σ . The set $\text{FO}[\sigma]$ is built from *atomic formulas* of the form $x_1 = x_2$ or $R(x_1, \dots, x_{\text{ar}(R)})$, where $R \in \sigma$ and $x_1, \dots, x_{\text{ar}(R)} \in \mathbf{var}$, and is closed under Boolean connectives ($\neg, \vee, \wedge, \rightarrow, \leftrightarrow$) and existential and universal quantifications (\exists, \forall). *Monadic second-order logic* (MSO) is the extension of first-order logic that also allows quantification over subsets of the domain. CMSO extends MSO by allowing first-order modular counting quantifiers \exists^m for every integer m (where $\exists^m \phi$ is true in a σ -db if the number of its elements for which ϕ is satisfied is divisible by m). A *free variable* of a formula is a (individual or set) variable that does not appear in the scope of a quantifier. A formula without free variables is called a *sentence*. For a σ -db \mathcal{D} and a sentence ϕ we write $\mathcal{D} \models \phi$ to denote that \mathcal{D} satisfies ϕ .

► **Proviso.** *For the rest of the paper, we fix a schema σ and numbers $d, t \in \mathbb{N}$ with $d \geq 2$. From now on, all databases are σ -dbs and have degree at most d , unless stated otherwise. We use \mathbf{C}_d to denote the class of all σ -dbs with degree at most d , \mathbf{C}_d^t to denote the class of all σ -dbs with degree at most d and tree-width at most t and finally we use \mathbf{C} to denote a class of σ -dbs with degree at most d .*

Property testing. Adler and Harwath [1] introduced the model of property testing for bounded degree relational databases, which is a straightforward extension of the model for bounded degree graphs [16]. We call this model the BDRD *model* for short, which we shall discuss below.

Property testing algorithms do not have access to the whole input database. Instead, they are given access via an *oracle*. Let \mathcal{D} be an input σ -db on n elements. A property testing algorithm receives the number n as input, and it can make *oracle queries*¹ of the form (R, i, j) , where $R \in \sigma$, $i \leq n$ and $j \leq \deg(\mathcal{D})$. The answer to (R, i, j) is the j^{th} tuple in $R^{\mathcal{D}}$ containing the i^{th} element² of \mathcal{D} (if such a tuple does not exist then it returns \perp). We assume oracle queries are answered in constant time.

¹ Note that an oracle query is not a database query.

² According to the assumed linear order on D .

Let $\mathcal{D}, \mathcal{D}'$ be two σ -dbs, both having n elements. In the BDRD model the *distance* between \mathcal{D} and \mathcal{D}' , denoted by $\text{dist}(\mathcal{D}, \mathcal{D}')$, is the minimum number of tuples that have to be inserted or removed from relations of \mathcal{D} and \mathcal{D}' to make \mathcal{D} and \mathcal{D}' isomorphic. For $\epsilon \in [0, 1]$, we say \mathcal{D} and \mathcal{D}' are ϵ -close if $\text{dist}(\mathcal{D}, \mathcal{D}') \leq \epsilon dn$, and \mathcal{D} and \mathcal{D}' are ϵ -far otherwise. A *property* is simply an isomorphism-closed class of databases. Note that every CMSO sentence ϕ defines a property $\mathbf{P}_\phi = \{\mathcal{D} \mid \mathcal{D} \models \phi\}$. We call $\mathbf{P}_\phi \cap \mathbf{C}$ the property *defined by ϕ on \mathbf{C}* . A σ -db \mathcal{D} is ϵ -close to a property \mathbf{P} if there exists a database $\mathcal{D}' \in \mathbf{P}$ that is ϵ -close to \mathcal{D} , otherwise \mathcal{D} is ϵ -far from \mathbf{P} .

Let $\mathbf{P} \subseteq \mathbf{C}$ be a property and $\epsilon \in (0, 1]$ be the proximity parameter. An ϵ -tester for \mathbf{P} on \mathbf{C} is a probabilistic algorithm which is given oracle access to a σ -db $\mathcal{D} \in \mathbf{C}$ and it is given $n := |\mathcal{D}|$ as auxiliary input. The algorithm does the following:


1. If $\mathcal{D} \in \mathbf{P}$, then the tester accepts with probability at least $2/3$.
2. If \mathcal{D} is ϵ -far from \mathbf{P} , then the tester rejects with probability at least $2/3$.

The *query complexity* of a tester is the maximum number of oracle queries made. A tester has *constant* query complexity, if the query complexity does not depend on the size of the input database. We say a property $\mathbf{P} \subseteq \mathbf{C}$ is *uniformly testable* in time $f(n)$ on \mathbf{C} , if for every $\epsilon \in (0, 1]$ there exists an ϵ -tester for \mathbf{P} on \mathbf{C} which has constant query complexity and whose running time on databases on n elements is $f(n)$. Note that this tester must work for all n .

Neighbourhoods. For a σ -db \mathcal{D} and $a, b \in D$, the *distance* between a and b in \mathcal{D} , denoted by $\text{dist}_{\mathcal{D}}(a, b)$, is the length of a shortest path between a and b in $\mathcal{G}(\mathcal{D})$. Let $r \in \mathbb{N}$. For an element $a \in D$, we let $N_r^{\mathcal{D}}(a)$ denote the set of all elements of \mathcal{D} that are at distance at most r from a . The r -neighbourhood of a in \mathcal{D} , denoted by $\mathcal{N}_r^{\mathcal{D}}(a)$, is the tuple $(\mathcal{D}[N_r^{\mathcal{D}}(a)], a)$ where a is called the *centre*. We omit the superscript and write $N_r(a)$ and $\mathcal{N}_r(a)$, if \mathcal{D} is clear from the context. Two r -neighbourhoods, $\mathcal{N}_r(a)$ and $\mathcal{N}_r(b)$, are *isomorphic* (written $\mathcal{N}_r(a) \cong \mathcal{N}_r(b)$) if there is an isomorphism between $\mathcal{D}[N_r(a)]$ and $\mathcal{D}[N_r(b)]$ which maps a to b . An \cong -equivalence-class of r -neighbourhoods is called an r -neighbourhood type (or r -type for short). We let $T_r^{\sigma, d}$ denote the set of all r -types with degree at most d , over schema σ . Note that for fixed d and σ , the cardinality $|T_r^{\sigma, d}| =: c(r)$ is a constant, only depending on r and d . We say that an element $a \in D$ has r -type τ , if $\mathcal{N}_r^{\mathcal{D}}(a) \in \tau$. For $r \in \mathbb{N}$, the r -histogram of a database \mathcal{D} , denoted by $h_r(\mathcal{D})$, is the vector with $c(r)$ components, indexed by the r -types, where the component corresponding to type τ contains the number of elements of \mathcal{D} of r -type τ . The r -neighbourhood distribution of \mathcal{D} , denoted by $\text{dv}_r(\mathcal{D})$, is the vector $h_r(\mathcal{D})/n$ where $|D| = n$. For a class of σ -dbs \mathbf{C} and $r \in \mathbb{N}$, we let $h_r(\mathbf{C}) := \{h_r(\mathcal{D}) \mid \mathcal{D} \in \mathbf{C}\}$. A set is *semilinear* if it is a finite union of linear sets. A set $M \subseteq \mathbb{N}^c$ is linear if $M = \{\bar{v}_0 + a_1 \bar{v}_1 + \dots + a_k \bar{v}_k \mid a_1, \dots, a_k \in \mathbb{N}\}$, for some $\bar{v}_0, \dots, \bar{v}_k \in \mathbb{N}^c$. From a result in [14] about many-sorted spectra of CMSO sentences it can be derived that that the set of r -histograms of properties defined by a CMSO sentence on \mathbf{C}_d^t are semilinear.

► **Lemma 1** ([1, 14]). *For each $r \in \mathbb{N}$ and each property $\mathbf{P} \subseteq \mathbf{C}_d^t$ definable by a CMSO sentence on \mathbf{C}_d^t , the set $h_r(\mathbf{P})$ is semilinear.*

Model of computation. We use Random Access Machines (RAMs) and a uniform cost measure when analysing our algorithms, i.e. we assume all basic arithmetic operations including random sampling can be done in constant time, regardless of the size of the numbers involved.



■ **Figure 1** The graphs $\mathcal{G}_{n,m}$ and $\mathcal{H}_{n,m}$ (respectively) of Example 3.

3 The Model

We shall now introduce our property testing model for bounded degree relational databases, which is an extension of the BDRD model discussed in Section 2. The notions of oracle queries, properties, ϵ -tester, query complexity and uniform testability remain the same but we have an alternative definition of distance and ϵ -closeness. In our model, which we shall call the BDRD $_{+/-}$ model for short, we can add and remove elements as well as tuples and can therefore compare databases that are on a different number of elements.

► **Definition 2** (Distance and ϵ -closeness). *Let $\mathcal{D}, \mathcal{D}' \in \mathbf{C}_d$ and $\epsilon \in [0, 1]$. The distance between \mathcal{D} and \mathcal{D}' (denoted by $\text{dist}_{+/-}(\mathcal{D}, \mathcal{D}')$) is the minimum number of modifications we need to make to \mathcal{D} and \mathcal{D}' to make them isomorphic where a modification is either (1) inserting a new element, (2) deleting an element (and as a result deleting any tuple that contains that element), (3) inserting a tuple, or (4) deleting a tuple. We then say \mathcal{D} and \mathcal{D}' are ϵ -close if $\text{dist}_{+/-}(\mathcal{D}, \mathcal{D}') \leq \epsilon d \min\{|\mathcal{D}|, |\mathcal{D}'|\}$ and are ϵ -far otherwise.*

The following example illustrates the difference between the distance measure of the BDRD and the distance measure of the BDRD $_{+/-}$ model.

► **Example 3.** Let $\mathbf{P} = \{\mathcal{G}_{n,m} \mid n, m \in \mathbb{N}_{>1}\}$ where $\mathcal{G}_{n,m}$ is an n by m grid graph as shown in Figure 1. Let us consider the graph $\mathcal{H}_{n,m}$ for some $n, m \in \mathbb{N}$ which is formed from $\mathcal{G}_{n,m}$ by removing a corner vertex. In the BDRD $_{+/-}$ model the distance between $\mathcal{H}_{n,m}$ and $\mathcal{G}_{n,m}$ is 1 (we remove a corner vertex from $\mathcal{G}_{n,m}$ to get $\mathcal{H}_{n,m}$) and therefore $\mathcal{H}_{n,m}$ is at distance 1 from \mathbf{P} in the BDRD $_{+/-}$ model. In the BDRD model if two graphs are on a different number of vertices then the distance between them is infinity. Therefore if $nm - 1$ is a prime number then $\mathcal{H}_{n,m}$ is at distance infinity from \mathbf{P} in the BDRD model.

We now show that if a property is testable in the BDRD model then it is also testable in the BDRD $_{+/-}$ model but the converse is not true. This allows for more testable properties in the BDRD $_{+/-}$ model.

► **Lemma 4** (*). *Let $\mathbf{P} \subseteq \mathbf{C}$. If \mathbf{P} is uniformly testable on \mathbf{C} in time $f(n)$ in the BDRD model then \mathbf{P} is also uniformly testable on \mathbf{C} in time $f(n)$ in the BDRD $_{+/-}$ model.*

► **Theorem 5** ([16]). *In the bounded degree model, bipartiteness cannot be tested with query complexity $o(\sqrt{n})$, where n is the number of vertices of the input graph.*

► **Lemma 6.** *There exists a class \mathbf{C} of σ -dbs and a property $\mathbf{P} \subseteq \mathbf{C}$ such that \mathbf{P} is trivially testable on \mathbf{C} in the BDRD $_{+/-}$ model but is not testable on \mathbf{C} in the BDRD model.*

Proof. Let \mathbf{C} be the class of all graphs with degree at most d . Let $\mathbf{P} = \mathbf{P}_1 \cup \mathbf{P}_2 \subseteq \mathbf{C}$ be the property where \mathbf{P}_1 contains all bipartite graphs in \mathbf{C} and \mathbf{P}_2 contains all graphs in \mathbf{C} that have an odd number of vertices. In the $\text{BDRD}_{+/-}$ model every $\mathcal{G} \in \mathbf{C}$ is ϵ -close to \mathbf{P} if $|V(\mathcal{G})| \geq 1/(\epsilon d)$ and hence \mathbf{P} is trivially testable on \mathbf{C} in the $\text{BDRD}_{+/-}$ model (the tester accepts if $|V(\mathcal{G})| \geq 1/(\epsilon d)$ and does a full check of the input otherwise). In the BDRD model, if the input graph has an even number of vertices then it is far from \mathbf{P}_2 and so we have to test for \mathbf{P}_1 . By Theorem 5, bipartiteness is not testable (with constant query complexity) in the BDRD model. In particular, in the proof of Theorem 5, Goldreich and Ron show that for any even n there exists two families, $\mathcal{G}_1 \subseteq \mathbf{C}$ and $\mathcal{G}_2 \subseteq \mathbf{C}$, of n -vertex graphs such that every graph in \mathcal{G}_1 is bipartite and almost all graphs in \mathcal{G}_2 are far from being bipartite but any algorithm that performs $o(\sqrt{n})$ queries cannot distinguish between a graph chosen randomly from \mathcal{G}_1 and a graph chosen randomly from \mathcal{G}_2 . Therefore \mathbf{P} is not testable on \mathbf{C} in the BDRD model. \blacktriangleleft

Note that the underlying general principle of the above proof can be applied to obtain further examples of properties that are testable in the $\text{BDRD}_{+/-}$ model but not testable in the BDRD model.

It is known that every hyperfinite property is “local” (Theorem 7), where “local” means that if a σ -db \mathcal{D} has a similar r -histogram to some σ -db (with the same domain size) that has the (hyperfinite) property, then \mathcal{D} must be ϵ -close to the property [23, 1]. This is summarised in Theorem 7 below. We use Theorem 7 to prove a similar result in the $\text{BDRD}_{+/-}$ model (Lemma 8). Lemma 8 is essential for the proof of Theorem 9.

► **Theorem 7** ([23, 1]). *Let $\epsilon \in (0, 1]$ and let \mathbf{C} be closed under removing tuples. If a property $\mathbf{P} \subseteq \mathbf{C}$ is hyperfinite on \mathbf{C} then there exists $\lambda_7 := \lambda_7(\epsilon) \in (0, 1]$ and $r_7 := r_7(\epsilon) \in \mathbb{N}$ such that for each $\mathcal{D} \in \mathbf{P}$ and $\mathcal{D}' \in \mathbf{C}$ with the same number n of elements, if $\|h_{r_7}(\mathcal{D}) - h_{r_7}(\mathcal{D}')\|_1 \leq \lambda_7 n$, then \mathcal{D}' is ϵ -close to \mathbf{P} in the BDRD model.*

► **Lemma 8.** *Let $\epsilon \in (0, 1]$ and let \mathbf{C} be closed under removing tuples. If a property $\mathbf{P} \subseteq \mathbf{C}$ is hyperfinite on \mathbf{C} then there exists $\lambda := \lambda(\epsilon) \in (0, 1]$ and $r := r(\epsilon) \in \mathbb{N}$ such that for each $\mathcal{D} \in \mathbf{P}$ and $\mathcal{D}' \in \mathbf{C}$, on $|D|$ and $|D'|$ elements respectively, if $\|h_r(\mathcal{D}) - h_r(\mathcal{D}')\|_1 \leq \lambda \min\{|D|, |D'|\}$, then \mathcal{D}' is ϵ -close to \mathbf{P} in the $\text{BDRD}_{+/-}$ model.*

Proof. Let $r = r_7(\epsilon/4)$ and let $\lambda = \frac{\epsilon \lambda_7(\epsilon/4)}{1 + d^{r+1}}$. Let us assume that $\|h_r(\mathcal{D}) - h_r(\mathcal{D}')\|_1 \leq \lambda \min\{|D|, |D'|\}$ and \mathbf{P} is hyperfinite on \mathbf{C} . If $|D| = |D'|$ then by Theorem 7 and the choice of λ , \mathcal{D}' is ϵ -close to \mathbf{P} . So let us assume that $|D| \neq |D'|$. Let \mathcal{D}_1 be the σ -db on $|D|$ elements formed from \mathcal{D}' by either removing $|D'| - |D|$ elements if $|D| < |D'|$ or adding $|D| - |D'|$ new elements if $|D'| < |D|$. Note that as $\|h_r(\mathcal{D}) - h_r(\mathcal{D}')\|_1 \leq \lambda \min\{|D|, |D'|\}$ and by definition $\|h_r(\mathcal{D}) - h_r(\mathcal{D}')\|_1 = \sum_{i=1}^{c(r)} |h_r(\mathcal{D}) - h_r(\mathcal{D}')|$ we have $||D| - |D'|| \leq \lambda \min\{|D|, |D'|\}$. When an element a is removed, the r -type of any element in $N_r(a)$ will change. As $|N_r(a)| \leq d^{r+1}$ (cf. e.g. Lemma 3.2 (a) of [7]) and $||D| - |D'|| \leq \lambda \min\{|D|, |D'|\}$, we have $\|h_r(\mathcal{D}') - h_r(\mathcal{D}_1)\|_1 \leq \lambda \min\{|D|, |D'|\} d^{r+1}$. Therefore

$$\|h_r(\mathcal{D}) - h_r(\mathcal{D}_1)\|_1 \leq \lambda \min\{|D|, |D'|\} (1 + d^{r+1}) \leq \lambda_7(\epsilon/4) |D|$$

by the choice of λ . By Theorem 7, in the BDRD model \mathcal{D}_1 is $\epsilon/4$ -close to \mathbf{P} . Hence there exists a σ -db $\mathcal{D}_2 \in \mathbf{P}$ such that $|D_2| = |D|$ and $\text{dist}(\mathcal{D}_1, \mathcal{D}_2) \leq \epsilon d |D| / 4$. By the definition of the two distance measures dist and $\text{dist}_{+/-}$, we have $\text{dist}_{+/-}(\mathcal{D}_1, \mathcal{D}_2) \leq \text{dist}(\mathcal{D}_1, \mathcal{D}_2) \leq \epsilon d |D| / 4$ and by the choice of \mathcal{D}_1 we have $\text{dist}_{+/-}(\mathcal{D}', \mathcal{D}_1) \leq \lambda \min\{|D|, |D'|\}$. Therefore

$$\text{dist}_{+/-}(\mathcal{D}', \mathcal{D}_2) \leq \frac{\epsilon d |D|}{4} + \lambda \min\{|D|, |D'|\} \leq \epsilon d \min\{|D|, |D'|\},$$

as $|D| \leq \min\{|D|, |D'|\} + \lambda \min\{|D|, |D'|\} \leq 2 \min\{|D|, |D'|\}$ and $\lambda \leq \epsilon d / 2$. Hence in the $\text{BDRD}_{+/-}$ model \mathcal{D}' is ϵ -close to \mathbf{P} . \blacktriangleleft

4 Main Results

We begin this section with the first of our main theorems (Theorem 9). We show that for any property \mathbf{P} which is hyperfinite on the input class \mathbf{C} , if the set of r -histograms of \mathbf{P} is semilinear, then for every σ -db \mathcal{D} in \mathbf{P} there exists a constant size σ -db in \mathbf{P} with a neighbourhood distribution similar to that of \mathcal{D} , but this is not true for σ -dbs in \mathbf{C} that are far from \mathbf{P} . We then use this result to prove that such properties are testable in constant time in the $\text{BDRD}_{+/-}$ model (Theorem 16). As a corollary we obtain that CMSO definable properties on σ -dbs of bounded tree-width and bounded degree are testable in constant time (Theorem 17).

► **Theorem 9.** *Let $\epsilon \in (0, 1]$ and let $r := r(\epsilon)$ be as in Lemma 8. Let \mathbf{C} be closed under removing tuples and let $\mathbf{P} \subseteq \mathbf{C}$ be a property that is hyperfinite on \mathbf{C} such that the set $h_r(\mathbf{P})$ is semilinear. There exist $n_{\min} := n_{\min}(\epsilon), n_{\max} := n_{\max}(\epsilon) \in \mathbb{N}$ and $f := f(\epsilon), \mu := \mu(\epsilon) \in (0, 1)$ such that for every $\mathcal{D} \in \mathbf{C}$ with $|D| > n_{\max}$,*

1. *if $\mathcal{D} \in \mathbf{P}$, then there exists a $\mathcal{D}' \in \mathbf{P}$ such that $n_{\min} \leq |D'| \leq n_{\max}$ and $\|\text{dv}_r(\mathcal{D}) - \text{dv}_r(\mathcal{D}')\|_1 \leq f - \mu$, and*
2. *if \mathcal{D} is ϵ -far from \mathbf{P} (in the $\text{BDRD}_{+/-}$ model), then for every $\mathcal{D}' \in \mathbf{P}$ such that $n_{\min} \leq |D'| \leq n_{\max}$, we have $\|\text{dv}_r(\mathcal{D}) - \text{dv}_r(\mathcal{D}')\|_1 > f + \mu$.*

Proof. Let $\lambda := \lambda(\epsilon)$ be as in Lemma 8 and $c := c(r)$ (the number of r -types). First note that if \mathbf{P} is empty then for any choice of n_{\min}, n_{\max}, f and μ , both 1. and 2. in the theorem statement are true and hence we shall assume that \mathbf{P} is non-empty. As $h_r(\mathbf{P})$ is a semilinear set we can write it as follows, $h_r(\mathbf{P}) = M_1 \cup M_2 \cup \dots \cup M_m$ where $m \in \mathbb{N}$ and for each $i \in [m]$, $M_i = \{\bar{v}_0^i + a_1 \bar{v}_1^i + \dots + a_{k_i} \bar{v}_{k_i}^i \mid a_1, \dots, a_{k_i} \in \mathbb{N}\}$ is a linear set where $\bar{v}_0^i, \dots, \bar{v}_{k_i}^i \in \mathbb{N}^c$ and for each $j \in [k_i]$, $\|\bar{v}_j^i\|_1 \neq 0$. Let $k := \max_{i \in [m]} k_i + 1$ and $v := \max_{i \in [m]} \left(\max_{j \in [0, k_i]} \|\bar{v}_j^i\|_1 \right)$ (note that $v > 0$ as \mathbf{P} is non-empty). Let $n_{\min} := n_0 - kv, n_{\max} := n_0 + kv, f := \frac{\lambda}{3c}$, and $\mu := \frac{\lambda}{6c}$ where

$$n_0 := kv \left(\frac{3ckv}{f - \mu} + 1 \right).$$

Note that $n_{\min} > 0$ by the choice of n_0, f and μ .

(Proof of 1.) Assume $\mathcal{D} \in \mathbf{P}$ and $|D| = n > n_{\max}$. Then there exists some $i \in [m]$ and $a_1^{\mathcal{D}}, \dots, a_{k_i}^{\mathcal{D}} \in \mathbb{N}$ such that $h_r(\mathcal{D}) = \bar{v}_0^i + a_1^{\mathcal{D}} \bar{v}_1^i + \dots + a_{k_i}^{\mathcal{D}} \bar{v}_{k_i}^i$ (note that $n = \|\bar{v}_0^i\|_1 + \sum_{j \in [k_i]} a_j^{\mathcal{D}} \|\bar{v}_j^i\|_1$). Let \mathcal{D}' be the σ -db with r -histogram $\bar{v}_0^i + a_1^{\mathcal{D}'} \bar{v}_1^i + \dots + a_{k_i}^{\mathcal{D}'} \bar{v}_{k_i}^i \in M_i$ where $a_j^{\mathcal{D}'}$ is the nearest integer to $a_j^{\mathcal{D}} n_0 / n$, and hence $a_j^{\mathcal{D}'} n_0 / n - 1/2 \leq a_j^{\mathcal{D}} n_0 / n \leq a_j^{\mathcal{D}'} n_0 / n + 1/2$. Note that since $\bar{v}_0^i + a_1^{\mathcal{D}'} \bar{v}_1^i + \dots + a_{k_i}^{\mathcal{D}'} \bar{v}_{k_i}^i \in h_r(\mathbf{P})$, \mathcal{D}' exists and $\mathcal{D}' \in \mathbf{P}$. We need to show that $n_{\min} \leq |D'| \leq n_{\max}$ and $\|\text{dv}_r(\mathcal{D}) - \text{dv}_r(\mathcal{D}')\|_1 \leq f - \mu$.

▷ Claim 10 (*). $|D'| \geq n_{\min}$.

▷ Claim 11 (*). $|D'| \leq n_{\max}$.

▷ Claim 12. $\|\text{dv}_r(\mathcal{D}) - \text{dv}_r(\mathcal{D}')\|_1 \leq f - \mu$.

Proof. By definition, $\|\text{dv}_r(\mathcal{D}) - \text{dv}_r(\mathcal{D}')\|_1 = \sum_{j \in [c]} |\text{dv}_r(\mathcal{D})[j] - \text{dv}_r(\mathcal{D}')[j]|$. First recall that $0 < n_0 - kv \leq |D'| \leq n_0 + kv < n$ and note that for every $\ell \in [k_i]$, $a_\ell^{\mathcal{D}} \leq n$ (since

7:10 Faster Property Testers in a Variation of the Bounded Degree Model

$\|\bar{v}_\ell^i\|_1 \neq 0$). Then for every $j \in [c]$, by the choice of $a_\ell^{\mathcal{D}'}$ for $\ell \in [k_i]$,

$$\begin{aligned} \text{dv}_r(\mathcal{D})[j] - \text{dv}_r(\mathcal{D}')[j] &= \bar{v}_0^i[j] \left(\frac{1}{n} - \frac{1}{|D'|} \right) + \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{a_\ell^{\mathcal{D}}}{n} - \frac{a_\ell^{\mathcal{D}'}}{|D'|} \right) \\ &\leq \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{a_\ell^{\mathcal{D}}}{n} - \frac{a_\ell^{\mathcal{D}} n_0}{n|D'|} + \frac{1}{2|D'|} \right) = \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{a_\ell^{\mathcal{D}}}{n} \left(\frac{|D'| - n_0}{|D'|} \right) + \frac{1}{2|D'|} \right) \\ &\leq \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{n}{n} \left(\frac{kv + n_0 - n_0}{n_0 - kv} \right) + \frac{1}{2(n_0 - kv)} \right) = \left(\frac{2kv + 1}{2(n_0 - kv)} \right) \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \\ &\leq \frac{kv(2kv + 1)}{n_0 - kv}. \end{aligned}$$

On the other hand,

$$\begin{aligned} \text{dv}_r(\mathcal{D})[j] - \text{dv}_r(\mathcal{D}')[j] &\geq -\frac{\bar{v}_0^i[j]}{|D'|} + \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{a_\ell^{\mathcal{D}}}{n} \left(\frac{|D'| - n_0}{|D'|} \right) - \frac{1}{2|D'|} \right) \\ &\geq -\frac{\bar{v}_0^i[j]}{|D'|} + \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{a_\ell^{\mathcal{D}}}{n} \left(\frac{-kv + n_0 - n_0}{|D'|} \right) - \frac{1}{2|D'|} \right) \\ &= -\frac{\bar{v}_0^i[j]}{|D'|} - \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{a_\ell^{\mathcal{D}} kv}{n|D'|} + \frac{1}{2|D'|} \right) \\ &\geq -\frac{\bar{v}_0^i[j]}{n_0 - kv} - \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{nk v}{n(n_0 - kv)} + \frac{1}{2(n_0 - kv)} \right) \\ &= -\frac{\bar{v}_0^i[j]}{n_0 - kv} - \left(\frac{2kv + 1}{2(n_0 - kv)} \right) \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \geq -\frac{kv(2kv + 1)}{n_0 - kv}. \end{aligned}$$

Hence,

$$|\text{dv}_r(\mathcal{D})[j] - \text{dv}_r(\mathcal{D}')[j]| \leq \frac{kv(2kv + 1)}{n_0 - kv} \leq \frac{3(kv)^2}{n_0 - kv} = \frac{f - \mu}{c}$$

by the choice of n_0 . Therefore,

$$\|\text{dv}_r(\mathcal{D}) - \text{dv}_r(\mathcal{D}')\|_1 = \sum_{j \in [c]} |\text{dv}_r(\mathcal{D})[j] - \text{dv}_r(\mathcal{D}')[j]| \leq f - \mu$$

as required. \triangleleft

(Proof of 2.) Assume \mathcal{D} is ϵ -far from \mathbf{P} and $|D| = n > n_{\max}$. For a contradiction let us assume there does exist a σ -db $\mathcal{D}' \in \mathbf{P}$ such that $n_{\min} \leq |D'| \leq n_{\max}$ and $\|\text{dv}_r(\mathcal{D}) - \text{dv}_r(\mathcal{D}')\|_1 \leq f + \mu$. As $\mathcal{D}' \in \mathbf{P}$ there exists some $i \in [m]$ and $a_1^{\mathcal{D}'}, \dots, a_{k_i}^{\mathcal{D}'} \in \mathbb{N}$ such that $\text{h}_r(\mathcal{D}') = \bar{v}_0^i + a_1^{\mathcal{D}'} \bar{v}_1^i + \dots + a_{k_i}^{\mathcal{D}'} \bar{v}_{k_i}^i$. Let \mathcal{D}'' be the σ -db with r -histogram $\bar{v}_0^i + a_1^{\mathcal{D}''} \bar{v}_1^i + \dots + a_{k_i}^{\mathcal{D}''} \bar{v}_{k_i}^i \in M_i$ where $a_j^{\mathcal{D}''}$ is the nearest integer to $a_j^{\mathcal{D}'} n / |D'|$. Note as $\bar{v}_0^i + a_1^{\mathcal{D}''} \bar{v}_1^i + \dots + a_{k_i}^{\mathcal{D}''} \bar{v}_{k_i}^i \in \text{h}_r(\mathbf{P})$, \mathcal{D}'' exists and $\mathcal{D}'' \in \mathbf{P}$.

\triangleright **Claim 13.** \mathcal{D} is ϵ -close to \mathbf{P} .

Proof. First note that as $\|\text{dv}_r(\mathcal{D}) - \text{dv}_r(\mathcal{D}')\|_1 \leq f + \mu$ and $\text{h}_r(\mathcal{D}') = \bar{v}_0^i + a_1^{\mathcal{D}'} \bar{v}_1^i + \dots + a_{k_i}^{\mathcal{D}'} \bar{v}_{k_i}^i$, for every $j \in [c]$

$$\frac{\bar{v}_0^i[j] + \sum_{\ell \in [k_i]} a_\ell^{\mathcal{D}'} \bar{v}_\ell^i[j]}{|D'|} - f - \mu \leq \text{dv}_r(\mathcal{D})[j] \leq \frac{\bar{v}_0^i[j] + \sum_{\ell \in [k_i]} a_\ell^{\mathcal{D}'} \bar{v}_\ell^i[j]}{|D'|} + f + \mu$$

and therefore

$$n \left(\frac{\bar{v}_0^i[j] + \sum_{\ell \in [k_i]} a_\ell^{\mathcal{D}'} \bar{v}_\ell^i[j]}{|D'|} - f - \mu \right) \leq h_r(\mathcal{D})[j] \leq n \left(\frac{\bar{v}_0^i[j] + \sum_{\ell \in [k_i]} a_\ell^{\mathcal{D}'} \bar{v}_\ell^i[j]}{|D'|} + f + \mu \right).$$

Hence, by the choice of $a_\ell^{\mathcal{D}''}$ for $\ell \in [k_i]$,

$$\begin{aligned} h_r(\mathcal{D})[j] - h_r(\mathcal{D}'')[j] &\leq \bar{v}_0^i[j] \left(\frac{n}{|D'|} - 1 \right) + \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{a_\ell^{\mathcal{D}'} n}{|D'|} - a_\ell^{\mathcal{D}''} \right) + fn + \mu n \\ &\leq \bar{v}_0^i[j] \frac{n}{|D'|} + \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{a_\ell^{\mathcal{D}'} n}{|D'|} - \left(\frac{a_\ell^{\mathcal{D}'} n}{|D'|} - \frac{1}{2} \right) \right) + fn + \mu n \\ &= \bar{v}_0^i[j] \frac{n}{|D'|} + \frac{1}{2} \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] + fn + \mu n. \end{aligned}$$

Similarly, by the choice of $a_\ell^{\mathcal{D}''}$ for $\ell \in [k_i]$ and as $n > |D'|$,

$$\begin{aligned} h_r(\mathcal{D})[j] - h_r(\mathcal{D}'')[j] &\geq \bar{v}_0^i[j] \left(\frac{n}{|D'|} - 1 \right) + \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{a_\ell^{\mathcal{D}'} n}{|D'|} - a_\ell^{\mathcal{D}''} \right) - fn - \mu n \\ &\geq -\bar{v}_0^i[j] \frac{n}{|D'|} + \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] \left(\frac{a_\ell^{\mathcal{D}'} n}{|D'|} - \left(\frac{a_\ell^{\mathcal{D}'} n}{|D'|} + \frac{1}{2} \right) \right) - fn - \mu n \\ &= -\bar{v}_0^i[j] \frac{n}{|D'|} - \frac{1}{2} \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] - fn - \mu n. \end{aligned}$$

Therefore,

$$\begin{aligned} |h_r(\mathcal{D})[j] - h_r(\mathcal{D}'')[j]| &\leq \bar{v}_0^i[j] \frac{n}{|D'|} + \frac{1}{2} \sum_{\ell \in [k_i]} \bar{v}_\ell^i[j] + fn + \mu n \\ &\leq \frac{n}{|D'|} \sum_{0 \leq \ell \leq k_i} \bar{v}_\ell^i[j] + fn + \mu n \leq \frac{nk v}{|D'|} + fn + \mu n \\ &= n \left(\frac{k v}{|D'|} + \frac{\lambda}{3c} + \frac{\lambda}{6c} \right) \leq n \left(\frac{\lambda}{18c} + \frac{\lambda}{3c} + \frac{\lambda}{6c} \right) = \frac{5\lambda n}{9c} \end{aligned}$$

by the choice of f and μ and as

$$|D'| \geq n_{\min} = \frac{3c(kv)^2}{f - \mu} = \frac{18(ckv)^2}{\lambda} \geq \frac{18ckv}{\lambda}.$$

To apply Lemma 8 we need to show that $\|h_r(\mathcal{D}) - h_r(\mathcal{D}'')\|_1 \leq \lambda \min\{n, |D''|\}$. If $|h_r(\mathcal{D})[j] - h_r(\mathcal{D}'')[j]| \leq \frac{\lambda}{c} \min\{n, |D''|\}$ then $\|h_r(\mathcal{D}) - h_r(\mathcal{D}'')\|_1 \leq \lambda \min\{n, |D''|\}$. Clearly, $\frac{5\lambda n}{9c} < \frac{\lambda n}{c}$. We also have

$$\begin{aligned} |D''| &= \|\bar{v}_0^i\|_1 + \sum_{\ell \in [k_i]} a_\ell^{\mathcal{D}''} \|\bar{v}_\ell^i\|_1 \geq \|\bar{v}_0^i\|_1 + \sum_{\ell \in [k_i]} \left(\frac{a_\ell^{\mathcal{D}'} n}{|D'|} - \frac{1}{2} \right) \|\bar{v}_\ell^i\|_1 \\ &= \|\bar{v}_0^i\|_1 - \frac{1}{2} \sum_{\ell \in [k_i]} \|\bar{v}_\ell^i\|_1 + \frac{n}{|D'|} \sum_{\ell \in [k_i]} a_\ell^{\mathcal{D}'} \|\bar{v}_\ell^i\|_1 \geq -kv + \frac{n}{|D'|} (|D'| - \|\bar{v}_0^i\|_1) \\ &\geq -\frac{n}{18} + \frac{17}{18}n \geq \frac{5n}{9} \end{aligned}$$

7:12 Faster Property Testers in a Variation of the Bounded Degree Model

as

$$|D'| \geq \frac{18ckv}{\lambda} \geq 18v \geq 18\|\bar{v}_0^i\|_1 \text{ and } kv \leq \frac{(ckv)^2}{\lambda} = \frac{n_{\min}}{18} \leq \frac{n}{18}.$$

Therefore, $\frac{5\lambda n}{9c} \leq \frac{\lambda|D'|}{c}$ and hence $\|\mathbf{h}_r(\mathcal{D}) - \mathbf{h}_r(\mathcal{D}')\|_1 \leq \lambda \min\{n, |D'|\}$. By Lemma 8, \mathcal{D} is ϵ -close to \mathbf{P} . \triangleleft

Claim 13 gives us a contradiction and therefore for every $\mathcal{D}' \in \mathbf{P}$ such that $n_{\min} \leq |D'| \leq n_{\max}$, we have $\|\mathbf{d}_{v_r}(\mathcal{D}) - \mathbf{d}_{v_r}(\mathcal{D}')\|_1 > f + \mu$ as required. \blacktriangleleft

As mentioned in the introduction, Alon [22, Proposition 19.10] proved that on bounded degree graphs, for any graph \mathcal{G} , radius r and $\epsilon > 0$ there always exists a graph \mathcal{H} whose size is independent of $|V(\mathcal{G})|$ and whose r -neighbourhood distribution vector satisfies $\|\mathbf{d}_{v_r}(\mathcal{G}) - \mathbf{d}_{v_r}(\mathcal{H})\|_1 \leq \epsilon$. However, the proof is only existential and does not provide an explicit bound on the size of \mathcal{H} . As a corollary to the proof of Theorem 9, we immediately obtain explicit bounds for classes of graphs and relational databases of bounded degree whose histogram vectors form a semilinear set.

► Corollary 14. *Let $\epsilon \in (0, 1]$, $r \in \mathbb{N}$ and \mathcal{D} be a σ -db that belongs to a class of σ -dbs \mathbf{C} such that the set $\mathbf{h}_r(\mathbf{C})$ is semilinear, i.e. $\mathbf{h}_r(\mathbf{C}) = M_1 \cup M_2 \cup \dots \cup M_m$ where $m \in \mathbb{N}$ and for each $i \in [m]$, $M_i = \{\bar{v}_0^i + a_1\bar{v}_1^i + \dots + a_{k_i}\bar{v}_{k_i}^i \mid a_1, \dots, a_{k_i} \in \mathbb{N}\}$ is a linear set where $\bar{v}_0^i, \dots, \bar{v}_{k_i}^i \in \mathbb{N}^{c(r)}$. Then there exists a σ -db \mathcal{D}_0 such that*

$$\|\mathbf{d}_{v_r}(\mathcal{D}) - \mathbf{d}_{v_r}(\mathcal{D}_0)\|_1 \leq \epsilon \text{ and } |D_0| \leq kv \left(\frac{3ckv}{\epsilon} + 2 \right)$$

where $c := c(r)$, $k := \max_{i \in [m]} k_i + 1$ and $v := \max_{i \in [m]} \left(\max_{j \in [0, k_i]} \|\bar{v}_j^i\|_1 \right)$.

Our aim is to construct constant time testers for hyperfinite properties whose set of r -histograms are semilinear. If we can approximate the r -neighbourhood distribution of a σ -db then by Theorem 9 we only need to check whether this distribution is close or not to the r -neighbourhood distribution of some small constant size σ -db. We let $\text{EstimateFrequencies}_{r,s}$ be the algorithm that, given oracle access to an input σ -db \mathcal{D} , samples s many elements uniformly and independently from D and computes their r -type. The algorithm then returns the r -neighbourhood distribution vector of the sample.

► Lemma 15 ([1]). *Let $\mathcal{D} \in \mathbf{C}_d$ be a σ -db on n elements, $\mu \in (0, 1)$ and $r \in \mathbb{N}$. If $s \geq c(r)^2/\mu^2 \cdot \ln(20c(r))$, with probability at least $9/10$ the vector \bar{v} returned by the algorithm $\text{EstimateFrequencies}_{r,s}$ on input \mathcal{D} satisfies $\|\bar{v} - \mathbf{d}_{v_r}(\mathcal{D})\|_1 \leq \mu$.*

► Theorem 16. *Let \mathbf{C} be closed under removing tuples and let $\mathbf{P} \subseteq \mathbf{C}$ be a property that is hyperfinite on \mathbf{C} . If for each $r \in \mathbb{N}$ the set $\mathbf{h}_r(\mathbf{P})$ is semilinear, then \mathbf{P} is uniformly testable on \mathbf{C} in constant time in the $\text{BDRD}_{+/-}$ model.*

Proof. Let $\epsilon \in (0, 1]$. Let $r := r(\epsilon)$ be as in Lemma 8, let $n_{\min} := n_{\min}(\epsilon)$, $n_{\max} := n_{\max}(\epsilon)$, $f := f(\epsilon)$ and $\mu := \mu(\epsilon)$ be as in Theorem 9 and let $s = c(r)^2/\mu^2 \cdot \ln(20c(r))$. Assume that the set $\mathbf{h}_r(\mathbf{P})$ is semilinear. Given oracle access to a σ -db $\mathcal{D} \in \mathbf{C}$ and $|D| = n$ as an input, the ϵ -tester proceeds as follows:

1. If $n \leq n_{\max}$, do a full check of \mathcal{D} and decide if $\mathcal{D} \in \mathbf{P}$.
2. Run $\text{EstimateFrequencies}_{r,s}$ and let \bar{v} be the resulting vector.
3. If there exists a $\mathcal{D}' \in \mathbf{P}$ where $n_{\min} \leq |D'| \leq n_{\max}$ and $\|\bar{v} - \mathbf{d}_{v_r}(\mathcal{D}')\|_1 \leq f$ then accept otherwise reject.

The running time and query complexity of the above tester is constant as n_{\max} is a constant (it only depends on \mathbf{P} , d and ϵ) and $\text{EstimateFrequencies}_{r,s}$ runs in constant time and makes a constant number of queries.

For correctness, first assume $\mathcal{D} \in \mathbf{P}$. By Theorem 9 there exists a σ -db $\mathcal{D}' \in \mathbf{P}$ such that $n_{\min} \leq |D'| \leq n_{\max}$ and $\|dv_r(\mathcal{D}) - dv_r(\mathcal{D}')\|_1 \leq f - \mu$. By Lemma 15 with probability at least $9/10$, $\|\bar{v} - dv_r(\mathcal{D})\|_1 \leq \mu$ and therefore $\|\bar{v} - dv_r(\mathcal{D}')\|_1 \leq f$. Hence with probability at least $9/10$ the tester will accept.

Now assume \mathcal{D} is ϵ -far from \mathbf{P} . By Theorem 9 for every $\mathcal{D}' \in \mathbf{P}$ with $n_{\min} \leq |D'| \leq n_{\max}$, we have $\|dv_r(\mathcal{D}) - dv_r(\mathcal{D}')\|_1 > f + \mu$. By Lemma 15 with probability at least $9/10$, $\|\bar{v} - dv_r(\mathcal{D})\|_1 \leq \mu$ and therefore for every $\mathcal{D}' \in \mathbf{P}$ with $n_{\min} \leq |D'| \leq n_{\max}$, $\|\bar{v} - dv_r(\mathcal{D}')\|_1 > f$. Hence with probability at least $9/10$ the tester will reject. ◀

Combining Theorem 16 and Lemma 1 and the fact that \mathbf{C}_d^t is hyperfinite [17, 3] (and so any property is hyperfinite on \mathbf{C}_d^t) we obtain the following as a corollary.

▶ **Theorem 17.** *Every property \mathbf{P} definable by a CMSO sentence on \mathbf{C}_d^t is uniformly testable on \mathbf{C}_d^t with constant time complexity in the $\text{BDRD}_{+/-}$ model.*

References

- 1 Isolde Adler and Frederik Harwath. Property testing for bounded degree databases. In *35th Symposium on Theoretical Aspects of Computer Science (STACS 2018)*, volume 96, page 6. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.
- 2 Noga Alon, Tali Kaufman, Michael Krivelevich, and Dana Ron. Testing triangle-freeness in general graphs. *SIAM Journal on Discrete Mathematics*, 22(2):786–819, 2008.
- 3 Noga Alon, Paul D. Seymour, and Robin Thomas. A separator theorem for graphs with an excluded minor and its applications. In Harriet Ortiz, editor, *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 293–299. ACM, 1990. doi:10.1145/100216.100254.
- 4 Jasine Babu, Areej Khoury, and Ilan Newman. Every property of outerplanar graphs is testable. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2016)*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016.
- 5 Sagi Ben-Moshe, Yaron Kanza, Eldar Fischer, Arie Matsliah, Mani Fischer, and Carl Staelin. Detecting and exploiting near-sortedness for efficient relational query evaluation. In *Proceedings of the 14th International Conference on Database Theory*, pages 256–267. ACM, 2011.
- 6 Itai Benjamini, Oded Schramm, and Asaf Shapira. Every minor-closed property of sparse graphs is testable. *Advances in mathematics*, 223(6):2200–2218, 2010.
- 7 Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering fo+ mod queries under updates on bounded degree databases. *ACM Transactions on Database Systems (TODS)*, 43(2):7, 2018.
- 8 Hubie Chen and Yuichi Yoshida. Testability of homomorphism inadmissibility: Property testing meets database theory. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 365–382. ACM, 2019.
- 9 Bruno Courcelle. Graph rewriting: An algebraic and logic approach. In *Formal Models and Semantics*, pages 193–242. Elsevier, 1990.
- 10 Bruno Courcelle and Joost Engelfriet. *Graph structure and monadic second-order logic: a language-theoretic approach*, volume 138. Cambridge University Press, 2012.
- 11 Arnaud Durand and Etienne Grandjean. First-order queries on structures of bounded degree are computable with constant delay. *ACM Transactions on Computational Logic (TOCL)*, 8(4):21, 2007.

- 12 Hendrik Fichtenberger, Pan Peng, and Christian Sohler. On constant-size graphs that preserve the local structure of high-girth graphs. In Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA*, volume 40 of *LIPICs*, pages 786–799. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICs.APPROX-RANDOM.2015.786.
- 13 Hendrik Fichtenberger, Pan Peng, and Christian Sohler. Every testable (infinite) property of bounded-degree graphs contains an infinite hyperfinite subproperty. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 714–726. SIAM, 2019.
- 14 Eldar Fischer, Johann A Makowsky, et al. On spectra of sentences of monadic second order logic with counting. *Journal of Symbolic Logic*, 69(3):617–640, 2004.
- 15 Jörg Flum and Martin Grohe. *Parameterized Complexity Theory (Texts in Theoretical Computer Science. An EATCS Series)*. Springer-Verlag, Berlin, Heidelberg, 2006.
- 16 Oded Goldreich and Dana Ron. Property testing in bounded degree graphs. *Algorithmica*, 32(2):302–343, 2002.
- 17 Avinatan Hassidim, Jonathan A. Kelner, Huy N. Nguyen, and Krzysztof Onak. Local graph partitions for approximation and testing. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*, pages 22–31. IEEE Computer Society, 2009. doi:10.1109/FOCS.2009.77.
- 18 Piotr Indyk, Andrew McGregor, Ilan Newman, and Krzysztof Onak. Open problems in data streams, property testing, and related topics. In *Bernitoto Workshop on Sublinear Algorithms*, 2011.
- 19 Hiro Ito, Areej Khoury, and Ilan Newman. On the characterization of 1-sided error strongly testable graph properties for bounded-degree graphs. *computational complexity*, 29(1):1, 2020.
- 20 Wojciech Kazana and Luc Segoufin. First-order query evaluation on structures of bounded degree. *Logical Methods in Computer Science*, 7(2), 2011. doi:10.2168/LMCS-7(2:20)2011.
- 21 Leonid Libkin. *Elements of finite model theory*. Springer, 2004.
- 22 László Lovász. *Large networks and graph limits*, volume 60. American Mathematical Soc., 2012.
- 23 Ilan Newman and Christian Sohler. Every property of hyperfinite graphs is testable. *SIAM Journal on Computing*, 42(3):1095–1112, 2013.

A Proofs of Section 3

Proof of Lemma 4. Let π be an ϵ -tester, that runs in time $f(n)$, for \mathbf{P} on \mathbf{C} in the BDRD model. We claim that π is also an ϵ -tester for \mathbf{P} on \mathbf{C} in the $\text{BDRD}_{+/-}$ model. Let $\mathcal{D} \in \mathbf{C}$ be the input σ -db. If $\mathcal{D} \in \mathbf{P}$ then π will accept with probability at least $2/3$. If \mathcal{D} is ϵ -far from \mathbf{P} in the $\text{BDRD}_{+/-}$ model then it must also be ϵ -far from \mathbf{P} in the BDRD model and therefore π will reject with probability at least $2/3$. Hence π is an ϵ -tester for \mathbf{P} on \mathbf{C} in the $\text{BDRD}_{+/-}$ model. ◀

B

 Proofs of Section 4

Proof of Claim 10. By the choice of $a_j^{\mathcal{D}'}$ for $j \in [k_i]$,

$$\begin{aligned}
 |D'| &= \|\bar{v}_0^i\|_1 + \sum_{j \in [k_i]} a_j^{\mathcal{D}'} \|\bar{v}_j^i\|_1 \geq \|\bar{v}_0^i\|_1 + \sum_{j \in [k_i]} \left(\frac{a_j^{\mathcal{D}} n_0}{n} - \frac{1}{2} \right) \|\bar{v}_j^i\|_1 \\
 &= \|\bar{v}_0^i\|_1 - \frac{1}{2} \sum_{j \in [k_i]} \|\bar{v}_j^i\|_1 + \frac{n_0}{n} \sum_{j \in [k_i]} a_j^{\mathcal{D}} \|\bar{v}_j^i\|_1 = \|\bar{v}_0^i\|_1 - \frac{1}{2} \sum_{j \in [k_i]} \|\bar{v}_j^i\|_1 + n_0 - \frac{n_0 \|\bar{v}_0^i\|_1}{n} \\
 &\geq \|\bar{v}_0^i\|_1 - \frac{1}{2} \sum_{j \in [k_i]} \|\bar{v}_j^i\|_1 + n_0 - \|\bar{v}_0^i\|_1 \geq -kv + n_0 = n_{\min},
 \end{aligned}$$

as $\sum_{j \in [k_i]} a_j^{\mathcal{D}} \|\bar{v}_j^i\|_1 = n - \|\bar{v}_0^i\|_1$ and $n > n_{\max} \geq n_0$. ◁

Proof of Claim 11. By the choice of $a_j^{\mathcal{D}'}$ for $j \in [k_i]$,

$$\begin{aligned}
 |D'| &= \|\bar{v}_0^i\|_1 + \sum_{j \in [k_i]} a_j^{\mathcal{D}'} \|\bar{v}_j^i\|_1 \leq \|\bar{v}_0^i\|_1 + \sum_{j \in [k_i]} \left(\frac{a_j^{\mathcal{D}} n_0}{n} + \frac{1}{2} \right) \|\bar{v}_j^i\|_1 \\
 &= \|\bar{v}_0^i\|_1 + \frac{1}{2} \sum_{j \in [k_i]} \|\bar{v}_j^i\|_1 + n_0 \left(1 - \frac{\|\bar{v}_0^i\|_1}{n} \right) \leq \sum_{0 \leq j \leq k_i} \|\bar{v}_j^i\|_1 + n_0 \leq kv + n_0 = n_{\max},
 \end{aligned}$$

as $\sum_{j \in [k_i]} a_j^{\mathcal{D}} \|\bar{v}_j^i\|_1 = n - \|\bar{v}_0^i\|_1$. ◁