Property Testing and Its Connection to Learning and Approximation

ODED GOLDBREICH

Weizmann Institute of Science, Rehovot, Israel

SHAFI GOLDWASSER

Massachusetts Institute of Technology, Cambridge, Massachusetts, and Weizmann Institute of Science, Rehovot, Israel

AND

DANA RON

Massachusetts Institute of Technology, Cambridge, Massachusetts

Abstract. In this paper, we consider the question of determining whether a function f has property P or is ϵ -far from any function with property P . A *property testing* algorithm is given a sample of the value of f on instances drawn according to some distribution. In some cases, it is also allowed to query f on instances of its choice. We study this question for different properties and establish some connections to problems in learning theory and approximation.

In particular, we focus our attention on testing graph properties. Given access to a graph G in the form of being able to query whether an edge exists or not between a pair of vertices, we devise algorithms to test whether the underlying graph has properties such as being bipartite, k -Colorable, or having a p -Clique (clique of density p with respect to the vertex set). Our graph property testing algorithms are probabilistic and make assertions that are correct with high probability, while making a number of queries that is *independent* of the size of the graph. Moreover, the property testing algorithms can be used to efficiently (i.e., in time linear in the number of vertices) construct partitions of the graph that correspond to the property being tested, if it holds for the input graph.

An extended abstract of this work appeared in *Proceedings of 36th Annual IEEE Symposium of Foundations of Computing Science*. IEEE, New York, 1995, pp. 339–348.

While this work was done, O. Goldreich was on sabbatical leave at Laboratory of Computer Science, Massachusetts Institute of Technology.

D. Ron would like to acknowledge the support of a National Science Foundation (NSF) postdoctoral fellowship and a Bunting Fellowship.

Authors' present addresses: O. Goldreich, Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel, e-mail: oded@wisdom.weizmann.ac.id; S. Goldwasser and D. Ron, Laboratory for Computer Science, Massachusetts Institute of Technology, 545 Technology Square, Room 326, Cambridge, MA 02139, e-mail: {shafi, danar}@theory.lcs.mit.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery (ACM), Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1998 ACM 0004-5411/98/0700-0653 \$05.00

Categories and Subject Descriptors: F.1.2 [Computation by Abstract Devices]: Modes of Computation—*probabilistic computation*; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*computations on discrete structures*

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Approximation algorithms, computational learning theory, graph algorithms

1. Introduction

Property Testing is concerned with the computational task of determining whether a given object has a predetermined property or is “far” from any object having the property. A notion of Property Testing was first formulated by Rubinfeld and Sudan [1996]. In their formulation, a property testing algorithm for property P is given oracle access to the tested function f . The algorithm must distinguish the case that f has property P from the case that f is far from any function having the property. Distance between functions is measured in terms of the fraction of arguments in the domain on which the functions disagree. Note that property testing so defined is a relaxation of the standard decision task (of distinguishing the case that f has property P from the case f does not have property P). Here we are interested in testers that are far more efficient than the corresponding decision procedure.

Property testing emerges naturally in the context of program checking and probabilistically checkable proofs (PCP). Specifically, in the context of program checking, one may choose to test that the program satisfies certain properties before checking that it computes a specified function. This paradigm has been followed both in the theory of program checking [Blum et al. 1993; Rubinfeld and Sudan 1996], and in practice where often programmers first test their programs by verifying that the programs satisfy properties that are known to be satisfied by the function they compute. In the context of probabilistically checkable proofs, the property tested is being a codeword with respect to a specific code. This paradigm, explicitly introduced in Babai et al. [1991a], has shifted from testing codes defined by low-degree polynomials¹ to testing Hadamard codes,² and recently to testing the “long code.”³ All of these papers have focused on property testing in the sense of Rubinfeld and Sudan [1996], and on testing algebraic properties such as linearity, multi-linearity and being a low-degree polynomial.

In this work we extend the scope of property testing in two ways:

- (1) Working within the above framework, we venture into the domain of *combinatorial objects*. In particular, we study property testing as applied to graph properties, demonstrating its relevance to standard and dual notions of approximation, and derive extremely fast algorithms for testing several natural graph properties. These in turn yield results such as a constant-time approximation scheme for Max-Cut in dense graphs. We believe that, in general, property testing offers a new perspective on approximation problems.

¹ See, for example, Babai et al. [1991a; 1991b], Feige et al. [1991], Arora and Safra [1992], and Arora et al. [1992].

² See, for example, Arora et al. [1992], Bellare et al. [1993; 1995a], Bellare and Sudan [1994], Kiwi [1996], and Trevisan [1997].

³ See, for example, Bellare et al. [1995b], Håstad [1996a; 1997], and Trevisan [1998].

- (2) We generalize the above definition so as to allow an arbitrary probability distribution D over arguments to the function f , as well as the consideration of algorithms that only obtain random labeled examples, of the form $(x, f(x))$, where x is selected according to D . Distance between functions is measured, accordingly, with respect to D . This formulation is inspired by the PAC learning model [Valiant 1984], and we indeed relate property testing so formulated to various variants of the PAC model. We believe that property testing offers new perspective with respect to computational learning theory.

We start with the second item.

1.1. GENERAL PROPERTY TESTING. We are interested in the following general question of *Property Testing*:

Let P be a fixed property of functions, and f be an unknown function. Our goal is to determine (possibly probabilistically) if f has property P or if it is far from any function that has property P , where distance between functions is measured with respect to some distribution D on the domain of f . Towards this end, we are given examples of the form $(x, f(x))$, where x is distributed according to D . We may also be allowed to query f on instances of our choice.

Let $\mathcal{F}$ be the class of functions that satisfy property P . Then, testing property P corresponds to testing membership in the class $\mathcal{F}$. The two most relevant parameters to property testing are the permitted distance, hereafter denoted ϵ , and the desired confidence, denoted δ . We require the tester to accept each function in $\mathcal{F}$ and reject every function that is further than ϵ away from any function in $\mathcal{F}$. We allow the tester to be probabilistic, and make incorrect positive and negative assertions with probability at most δ . The complexity measures we focus on are the *sample complexity* (the number of examples of the function's values that the tester requires), the *query complexity* (the number of function queries made—if at all), and the *running time* of the tester.

We believe that property testing is a natural notion whose relevance to applications goes beyond program checking, and whose scope goes beyond the realm of testing algebraic properties. Firstly, in some cases, one may be merely interested in whether a given function, modeling an environment (respectively, a given program) possess a certain property rather than be interested in learning the function (respectively, checking that the program computes a specific function correctly). In such cases, learning the function (respectively, checking the program) as a means of ensuring that it satisfies the property may be an over-kill. Secondly, theoretical analysis of learning algorithms typically works under the postulation that the function (representing the environment) belongs to a particular class.⁴ It may be more efficient to test this postulation first before trying to learn the function (and possibly failing when the postulation is wrong).⁵ We stress that the generalization of property testing to arbitrary distributions and

⁴ Here and throughout the introduction we refer to the standard PAC learning model. We shortly discuss agnostic learning [Kearns et al. 1992], where no assumption is made on the target function, towards the end of Section 3.

⁵ Similarly, in the context of program checking, one may choose to test that the program satisfies certain properties before checking that it computes a specified function. As mentioned above, this paradigm has been followed both in the theory and practice of program checking.

to algorithm that only obtain random labeled examples is essential to some of the potential applications mentioned above.

1.1.1. Property Testing and Learning Theory. Our formulation of testing mimics the standard frameworks of learning theory. In both cases one is given access to an unknown *target* function (either in the form of random instances accompanied by the function values or in the form of oracle access to the function). (An insignificant semantic difference is that, for sake of uniformity, even in case the functions are Boolean, we refer to them as functions rather than concepts.) However, there are two important differences between property testing and learning. Firstly, the goal of a learning algorithm is to *find* a good approximation to the target function $f \in \mathcal{F}$, whereas a testing algorithm should only *determine* whether the target function is in $\mathcal{F}$ or is far away from it. This makes the task of the testing seem easier than that of learning. On the other hand, a learning algorithm should perform well only when the target function belongs to $\mathcal{F}$ whereas a testing algorithm must perform well also on functions far away from $\mathcal{F}$.

We show that the relation between learning and testing is nontrivial. On one hand, *proper* learning (i.e., when the hypothesis of the learning algorithm must belong to the same class as the target function) implies testing. On the other hand, there are function classes for which testing is harder than (*nonproper*) learning (i.e., when the hypothesis is not required to belong to the same class as the target function), provided $\mathcal{NP} \not\subseteq \mathcal{BPP}$. Nonetheless, there are also function classes for which testing is much easier than learning. In addition, the graph properties discussed below provide a case where testing (with queries) is much easier than learning (also with queries).

The above results as well as additional results regarding the relations between property testing and learning appear in Section 3.

1.2. TESTING GRAPH PROPERTIES. Property testing is a natural notion of approximation, and furthermore it is related to standard notions of approximation. This holds even with respect to the restricted notion of property testing where one considers only the uniform distribution and allows the algorithm to make queries of its choice. The above assertion is demonstrated within one of the most basic domains of approximation algorithms—the one of graph algorithms. But let us start with a general high level discussion.

1.2.1. Motivating Discussion. Throughout the rest of the introduction, we refer to property testing in the restricted sense of Rubinfeld and Sudan [1996]. Recall that the definition of property testing is a relaxation of the standard definition of a decision task: The tester is allowed arbitrary behavior when the object does not have the property, and yet is “close” to an object having the property. Thus, a property tester may be far more efficient than a standard decision procedure (for the same property).

In case the object is huge, as when one thinks of a function and algorithms which operate in time polynomial in the length of the arguments to the function, there is actually no other alternative to approximation. That is, it is typically infeasible (i.e., requires exponential time in the length of the arguments) to decide whether such a function has the desired property. A property testing algorithm which operates in time polynomial in the length of the arguments thus

offers a feasible approximation to a problem that is intractable in the exact formulation.

Property testers are valuable also in case one deals with objects of feasible size (i.e., size for which scanning the entire object is feasible): If a property tester is much faster than the exact decision procedure, then it makes sense to run it before running the decision procedure. In case the object is far from having the property, we may obtain an indication towards this fact, and save the time we might have used running the decision procedure. In case the tester supplies proofs of violation of the property (as in some of the testers discussed below), we obtain an absolutely correct answer without running the decision procedure at all. Thus, we may only need to run the decision procedure on objects which are close to having the property. In some setting where *typical* objects are either good (i.e., have the property) or very bad (i.e., are very far from objects having the property), we may gain a lot. Furthermore, *if it is guaranteed* that objects are either good or very bad, then we may not even need to run the decision procedure at all. The gain in such a setting is enormous.

Being close to an object which has the property is a notion of approximation which, in certain applications, may be of great value. In some cases, being close to an object having the property translates to a standard notion of approximation. In other cases, it translates to a notion of “dual approximation”. This point is clarified and exemplified below (by referring to specific properties). In both cases, a fast property tester that is more efficient than the decision procedure is of value, both if the decision procedure is feasible and more so if it is not.

Alternatively, we may be forced to take action, without having time to run a decision procedure, while given the option of modifying the object in the future, at a cost proportional to the number of modifications of the object. For example, suppose you are given a graph which represents some design problem, where Bipartite graphs correspond to a good design and changes in the design correspond to edge additions/omissions. Using a Bipartiteness tester, you always accept a good design and reject with high probability designs that will cost a lot to modify. You may still accept bad designs, but then you know that it will not cost you much to modify them later.

1.2.2. Representing graphs as functions. We view graphs as Boolean functions on pairs of vertices, the value of the function representing the existence of an edge. We mainly consider testing algorithms which use queries and work under the uniform distribution. That is, a testing algorithm for graph property P makes queries of the form “is there an edge between vertices u and v ” in an unknown graph G . Accordingly, distance between two N -vertex graphs is defined as the fraction (over N^2) of vertex-pairs that are adjacent in one graph but not in the other. A testing algorithm for property P is required to decide whether G has property P or is “ ϵ -away” from any graph with property P , and is allowed to err with probability, say, $\delta = 1/3$. (For simplicity, we assume throughout the introduction that $\delta = 1/3$. Standard error-reduction techniques are obviously applicable to property testing. However, better dependencies on δ may be obtained in special cases as can be seen in the rest of the paper.)

A few comments are in place. Firstly, we note that this representation, as well as our results, are applicable both in case the graph is huge and in case it is of feasible size. The reader may thus choose whether to think of the graph as being

huge and so accessible only by such queries, or as being explicitly given and inspected in a random-access manner. Secondly, the above adjacency predicate representation is most appropriate for dense graphs, and so the reader may think of the graph as being dense (e.g., having at least ϵN^2 edges). An alternative representation, appropriate for bounded-degree graphs, has been subsequently considered in Goldreich and Ron [1997a].

1.2.3. Our Algorithms. We present algorithms of $\text{poly}(1/\epsilon)$ query-complexity and running-time⁶ at most $\exp(\tilde{O}(1/\epsilon^3))$ for testing the following natural graph properties:

- k-Colorability*, for any fixed $k \geq 2$. (Here the query-complexity is $\text{poly}(k/\epsilon)$, and for $k = 2$ the running-time is $\tilde{O}(1/\epsilon^3)$.)
- *ρ -Clique*, for any $\rho > 0$. That is, does the N -vertex graph have a clique of size ρN ?
- *ρ -Cut*, for any $\rho > 0$. That is, does the N -vertex graph have a cut of size at least ρN^2 ? A generalization to k -way cuts works with query-complexity $\text{poly}((\log k)/\epsilon)$, and has running time $\exp(\text{poly}(\log k/\epsilon))$.
- *ρ -Bisection*, for any $\rho > 0$. That is, does the N -vertex graph have a bisection of size at most ρN^2 ?

Furthermore:

- (1) For all the above properties, in case the graph has the desired property, the testing algorithm outputs some auxiliary information which allows to construct, in $\text{poly}(1/\epsilon) \cdot N$ -time, a partition that approximately obeys the property. For example, for ρ -Cut, we can construct a partition with at least $(\rho - \epsilon)N^2$ crossing edges.
- (2) The k -Colorability tester has one-sided error: it always accepts k -Colorable graphs. Furthermore, when rejecting a graph, this tester always supplies a $\text{poly}(1/\epsilon)$ -size subgraph that is not k -Colorable. All other algorithms have two-sided error, and this is unavoidable within $o(N)$ query-complexity.
- (3) Our algorithms for k -Colorability, ρ -Clique and ρ -Cut can be easily extended to provide testers with respect to *product distributions*: that is, distributions $\Psi: V(G)^2 \mapsto [0, 1]$ of the form $\Psi(u, v) = \psi(u) \cdot \psi(v)$, where $\psi: V(G) \mapsto [0, 1]$ is a distribution on the vertices, which is poly-time sampleable (by the tester). In contrast, it is not possible to test any of the graph properties discussed above in a distribution-free manner.

We comment that, except for Bipartite (2-Colorability) testing, running-time of $\text{poly}(1/\epsilon)$ is unlikely, as it will imply $\mathcal{NP} \subseteq \mathcal{BPP}$. Also, none of these properties can be tested without queries when using $o(\sqrt{N})$ random examples.

1.2.3.1. GENERAL GRAPH PARTITION. All the above property testing problems are special cases of the *General Graph Partition Testing Problem*, parameterized by a set of lower and upper bounds. In this problem, one needs to determine whether there exists a k -partition of the vertices so that the number of vertices in each component of the partition as well as the number of edges between each

⁶ Here and throughout the paper, we consider a RAM model in which trivial manipulation of vertices (e.g., reading/writing a vertex name and ordering vertices) can be done in constant time.

pair of components falls between the corresponding lower and upper bounds (in the set of parameters). We present an algorithm for solving the above problem. The algorithm uses $\tilde{O}(k^2/\epsilon)^{2k+O(1)}$ queries, runs in time exponential in its query-complexity, and makes two-sided error. Approximating partitions, if existing, can be efficiently constructed in this general case as well. We comment that the specialized algorithms perform better than the general algorithm with the appropriate parameters.

1.2.3.2. OTHER GRAPH PROPERTIES. Going beyond the general graph partition problem, we remark that there are graph properties that are very easy to test (e.g., Connectivity, Hamiltonicity, and Planarity). The reason is that for these properties either every N -vertex graph is at distance at most $O(1/N)$ from a graph having the desired property (and so for $\epsilon = \Omega(1/N)$, the trivial algorithm which always accepts will do), or the property holds only for sparse graphs (and so for $\epsilon = \Omega(1/N)$ one may reject any non-sparse graph). On the other hand, there are graph properties in $\mathcal{NP}$ that are extremely hard to test; namely, any testing algorithm must inspect at least $\Omega(N^2)$ of the vertex pairs. In view of the above, we believe that providing a characterization of graph properties, according to the complexity of testing them, may be very challenging.

1.2.4. Testing versus deciding and approximating. We shortly discuss the relation between testing graph properties and some well-known computational tasks.

1.2.4.1. RELATION TO DECIDING (RECOGNIZING) GRAPH PROPERTIES. Our notion of testing a graph property P is a *relaxation* of the notion of *deciding* (*recognizing*) *the graph property* P , which has received much attention in the last three decades [Lovász and Young 1991]. In the classical problem, there are no margins of error, and one is required to accept all graphs having property P and reject all graphs which lack it. In 1975, Rivest and Vuillemin resolved the Aanderaa–Rosenberg Conjecture [Rosenberg 1973], showing that any deterministic procedure for deciding any nontrivial monotone N -vertex graph property must examine $\Omega(N^2)$ entries in the adjacency matrix representing the graph. The query complexity of randomized decision procedures was conjectured by Yao to also be $\Omega(N^2)$. Progress towards proving this conjecture was made in Yao [1987], King [1991], and Hajnal [1991] culminating in an $\Omega(N^{4/3})$ lower bound. This stands in striking contrast to the results mentioned above, by which some nontrivial monotone graph properties can be *tested* by examining a constant number of locations in the matrix.

1.2.4.2. APPLICATION TO THE STANDARD NOTION OF APPROXIMATION. The relation of testing graph properties to the standard notions of approximation is best illustrated in the case of Max-Cut. Any tester for the class ρ -Cut, working in time $T(\rho, \epsilon, N)$, yields an algorithm for approximating the maximum cut in an N -vertex graph, up to additive error ϵN^2 , in time $1/\epsilon \cdot T(\rho, \epsilon, N)$. Thus, for any constant $\epsilon > 0$, using the abovementioned tester we can approximate the size of the max-cut to within ϵN^2 in constant time.⁷ This yields a *constant time approximation scheme* (i.e., to within any constant relative error) for dense graphs, improving over previous work of Arora et al. [1995] and de la Vega

⁷ We comment that due to the specific structure of our tester, the value of the maximum cut is actually approximated in time $T(\rho, \epsilon, N)$, rather than $1/\epsilon \cdot T(\rho, \epsilon, N)$.

[1994] who solved this problem in polynomial-time (i.e., in $O(N^{1/\epsilon^2})$ -time and $(\exp(\tilde{O}(1/\epsilon^2)) \cdot N^2)$ -time, respectively). In the latter works, the problem is solved by actually constructing approximate max-cuts. Finding an approximate max-cut does not seem to follow from the mere existence of a tester for ρ -cut; yet, as mentioned above, our tester can be used to find such a cut in time linear in N (i.e., $(\tilde{O}(1/\epsilon^2) \cdot N + \exp(\tilde{O}(1/\epsilon^3)))$ -time).

One can turn the question around and ask whether approximation algorithms for dense instances can be transformed into corresponding testers as defined above. In several cases this is possible. For example, using some ideas of this work, the Max-Cut algorithm of de la Vega [1994] can be transformed into a tester of complexity comparable to ours. We do not know whether the same is true with respect to the algorithms in Arora et al. [1995]. Results on testing graph properties can be derived also from Alon et al. [1994].

1.2.4.3. RELATION TO “DUAL APPROXIMATION” (CF., [HOCHBAUM AND SHMOYS 1987; HOCHBAUM AND SHMOYS 1988]). To illustrate this relation, we consider the ρ -Clique Tester mentioned above. The traditional notion of approximating Max-Clique corresponds to distinguishing the case in which the max-clique has size at least ρN from, say, the case in which the max-clique has size at most $\rho N/2$. On the other hand, when we talk of testing “ ρ -Cliqueness”, the task is to distinguish the case in which an N -vertex graph has a clique of size ρN from the case in which it is ϵ -far from the class of N -vertex graphs having a clique of size ρN . This is equivalent to the “dual approximation” task of distinguishing the case in which an N -vertex graph has a clique of size ρN from the case in which any ρN subset of the vertices misses at least ϵN^2 edges. To demonstrate that these two tasks are vastly different we mention that whereas the former task is NP-Hard, for $\rho < 1/4$ (see Bellare et al. [1995b], Håstad [1996a], and Håstad [1996b]), the latter task can be solved in $\exp(O(1/\epsilon^2))$ -time, for any $\rho, \epsilon > 0$. We believe that there is no absolute sense in which one of these approximation tasks is more important than the other: Each of these tasks may be relevant in some applications and irrelevant in others.

As another illustration of the applicability to “dual approximation” problems, we discuss our results regarding testing k -Colorability. It is known that it is NP-Hard to distinguish 3-Colorable graphs from graphs in which every 3-partition of the vertex set violates at least a constant fraction of the edges [Petrak 1994]. In contrast, our k -Colorability Tester implies that solving the same promise problem is easy for *dense graphs*, where by dense graphs we mean N -vertex graphs with $\Omega(N^2)$ edges. This is the case since, for every $\epsilon > 0$, our tester can distinguish, in $\exp(k^2/\epsilon^3)$ -time, between k -Colorable N -vertex graphs and N -vertex graphs which remain non- k -Colorable even if one omits at most ϵN^2 of their edges.⁸

Another application of our 3-Colorability Tester uses the fact that, for every $\epsilon > 0$, in case the N -vertex graph is 3-colorable, the tester may retrieve in linear time a 3-partition which violates at most ϵN^2 edges. Thus, we may reduce the general problem of coloring 3-Colorable graphs with few edges to the same problem restricted to nondense graphs (i.e., N -vertex graphs with $O(N^2)$ edges).

⁸ As noted by Noga Alon, similar results, alas with much worse dependence on ϵ , can be obtained by using the results of Alon et al. [1994].

(The reduction produces a coloring for dense graphs with 3 times more colors than the number used by the coloring of the nondense graphs, but a factor of 3 seems small at the current state of art for this problem [Karger et al. 1994].) We remark that some known algorithms for this task seem to perform better when the maximum degree of vertices in the graph is smaller [Karger et al. 1994]. Furthermore, deciding k -Colorability even for N -vertex graphs of minimum degree at least $(k - 3)/(k - 2) \cdot N$ is NP-complete (cf., [Edwards 1986]). On the other hand, Edwards also gave a polynomial-time algorithm for k -coloring k -colorable N -vertex graphs of minimum degree at least αN , for any constant $\alpha > (k - 3)/(k - 2)$.

1.2.5. Our Techniques. Our algorithms share some underlying ideas. The first is the uniform selection of a small sample of vertices and the search for a *suitable* partition of this sample. In case of k -Colorability, certain k -Colorings of the subgraph induced by this sample will do, and are found by k -Coloring a slightly augmented graph. In case of the other algorithms, we exhaustively try all possible partitions. This is reminiscent of the *exhaustive sampling* of Arora et al. [1995], except that the partitions considered by us are always directly related to the combinatorial structure of the problem. We show how each possible partition of the sample induces a partition of the entire graph so that the following holds. If the tested graph has the property in question then, with high probability over the choice of the sample, there exists a partition of the sample which induces a partition of the entire graph so that the latter partition approximately satisfies the requirements established by the property in question. For example, in case the graph has a ρ -Cut there exists a 2-way-partition of the sample inducing a partition of the entire graph with at least $(\rho - \epsilon)N^2$ crossing edges. On the other hand, if the graph should be rejected by the test, then by definition no partition of the entire graph (and in particular none of the induced partitions) approximately obeys the requirements.

The next idea is to use an additional sample to approximate the quality of each such induced partition of the graph, and discover if at least one of these partitions approximately obeys the requirements of the property in question. An important point is that since the first sample is small (i.e., of size $\text{poly}(1/\epsilon)$), the total number of partitions it induces is only $\exp(\text{poly}(1/\epsilon))$. Thus, the additional sample must approximate only these many partitions (rather than all possible partitions of the entire graph) and it suffices that this sample be of size $\text{poly}(1/\epsilon)$.

The difference between the various algorithms is in the way in which partitions of the sample induce partitions of the entire graph. The simplest case is in testing Bipartiteness. For a partition (S_1, S_2) of the sample, all vertices in the graph that have a neighbor in S_1 are placed on one side, and the rest of the vertices are placed on the other side. In the other algorithms, the induced partition is less straightforward. For example, in case of ρ -Clique, a partition (S_1, S_2) of the sample S with $|S_1| \approx \rho|S|$, induces a candidate clique roughly as follows. Consider the set T of graph vertices each neighboring all of S_1 . Then the candidate clique consists of the ρN vertices with the highest degree in the subgraph induced by T . In the Bisection and General Partition testing algorithms, auxiliary guesses that are implemented by exhaustive search are used (to induce a partition on the entire graph).

A simple observation that is useful in our analyses is that we do not need the sample to approximate well all relevant quantities (e.g., the degree of all vertices in the graph). It suffices to approximate well most of these quantities. This observation may explain how we can manage with a sample of size independent of the size of the graph.

1.3. OTHER RELATED WORK

1.3.1. Property Testers Implicit in Previous Works. As mentioned above, results on testing graph properties can be derived from Alon et al. [1994]. That paper proves a constructive version of the Regularity Lemma of Szemerédi, and obtains from it a polynomial-time algorithm that given an N -vertex graph, $\epsilon > 0$ and $k \geq 3$, either finds a subgraph of size $f(\epsilon, k)$ which is not k -Colorable, or omits at most ϵN^2 edges and k -Colors the rest. Noga Alon has observed that the analysis can be modified to yield that almost all subgraphs of size $f(\epsilon, k)$ are not k -Colorable, which in turn implies a tester with query complexity $f(\epsilon, k)^2$ for k -Colorability. In comparison with our k -Colorability Tester, which takes a sample of $O(\epsilon^{-3}k^2 \log k)$ vertices, the k -Colorability tester derived from Alon et al. [1994] takes a much bigger sample—of size equaling a tower of $(k/\epsilon)^{20}$ exponents (i.e., $\log^* f(\epsilon, k) = (k/\epsilon)^{20}$).

1.3.2. Property Testing in the Context of Program Checking. There is an immediate analogy between program self-testing [Blum et al. 1993] and property-testing *with queries*. The difference is that in self-testing, a function f (represented by a program) is tested for being close to a fully specified function g , whereas in property-testing the test is whether f is close to any function in a function class $\mathcal{G}$. Interestingly, many self-testers [Blum et al. 1993; Rubinfeld and Sudan 1996] work by *first* testing that the program satisfies some properties that the function it is supposed to compute satisfies (and only then checking that the program satisfies certain constraints specific to the function). Rubinfeld and Sudan [1996] defined property testing, under the uniform distribution and using queries, and related it to their notion of Robust Characterization. Rubinfeld [1994] focuses on property testing as applied to properties which take the form of functional equations of various types.

1.3.3. Property Testing in the Context of Learning Theory. Departing from work in Statistics regarding the classification of distributions (e.g., [Hoeffding and Wolfowitz 1958; Cover 1973; Zeitouni and Kulkarni [1995]]), Ben-David [1992] and Kulkarni and Zeitouni [1993] considered the problem of classifying an unknown function into one of two classes of functions, given labeled examples. Ben-David studied this classification problem in the limit (of the number of examples), and Kulkarni and Zeitouni studied it in a PAC inspired model. For any fixed ϵ , the problem of testing the class $\mathcal{F}$ with distance parameter ϵ can be casted as such a classification problem (with $\mathcal{F}$ and the set of functions ϵ -away from $\mathcal{F}$ being the two classes). A different variant of the problem was considered by Yamanishi [1995].

1.3.4. Approximation in Dense Graphs. As stated previously, Arora et al. [1995] and de la Vega [1994] presented polynomial-time approximation schemes (PTAS) for dense instances of Max-Cut. The approach of Arora et al. uses Linear Programming and Randomized Rounding, and applies to other problems

which can be casted as “smooth” Integer Programs.⁹ The methods of de la Vega [1994] are purely combinatorial and apply also to similar graph partition problems. Following the approach of Alon et al. [1994], but using a relaxation of the regularity Lemma (and thus obtaining much improved running times), Frieze and Kanan [1996] devise PTAS for several graph partition problems such as Max-Cut and Bisection. We note that compared to all the above results, our respective graph partitioning algorithms have better running-times (not to mention that we obtain constant-time approximation schemes for approximating only the value of the partition). Like de la Vega, our methods use elementary combinatorial arguments related to the problem at hand. Still our methods suffice for dealing with the General Graph Partition Problem.

We note that Arora et al. [1995] showed that the “dense subgraph” problem, a generalization of ρ -Clique, has a PTAS for dense instances. Our General Graph Partition algorithm (with the appropriate setting of the parameters) improves on their result.

1.4. SUBSEQUENT WORK. As mentioned above, our representation of graphs by their adjacency predicate is most adequate for dense graphs. Another natural representation, most adequate for bounded-degree graphs was subsequently suggested in Goldreich and Ron [1997]: An N -vertex graph of degree bound d is represented there by the *incidence function*, $g: V \times [d] \mapsto V \cup \{0\}$, so that $g(u, i) = v$ if v is the i th vertex incident at u , and $g(u, i) = 0 \notin V$ if u has less than i neighbors.

As usual, the choice of representation has a fundamental impact on the potential algorithm. Here, the impact is even more dramatic since we seek algorithms that only inspect a relatively small fraction of the object (graph represented by a function). Furthermore, there is another fundamental impact of the choice of representation on the task of property testing. This has to do with our definition of distance, which is relative to the size of the domain of the function. In particular, distance ϵ in the adjacency predicate representation (adopted in this paper) means a symmetric difference of $2\epsilon \cdot N^2$ edges, whereas in the incident function representation (of Goldreich and Ron [1997]) this means a symmetric difference of $2\epsilon \cdot dN$ edges. (In both cases, the extra factor 2 is due to the redundant representation which is adopted for sake of simplicity.)

In contrast to our $\text{poly}(1/\epsilon)$ -query Bipartite tester (for the adjacency predicate representation), it was proven in Goldreich and Ron [1997] that testing Bipartiteness in the incident function representation requires $\Omega(\sqrt{N})$ queries. Interestingly, this bound is tight up to a polylogarithmic factor, as shown in Goldreich and Ron [1998] which presents a Bipartite tester for the incident function representation working in time $O(\text{poly}(\epsilon^{-1} \log N) \cdot \sqrt{N})$. We mention that Goldreich and Ron [1997] also presents $\text{poly}(1/\epsilon)$ -time algorithms for testing k -Connectivity, for $k \geq 1$, Planarity and other properties (all in the incident function representation).

In recent work, Kearns and Ron [1998] generalize our definition of property testing, and present testing algorithms that use only random examples for classes of functions that have been studied in the learning literature.

⁹ In Arora et al. [1996], the approach of Arora et al. [1995] is extended to other problems, such as Graph Isomorphism, using a new rounding procedure for the Assignment Problem.

We also mention the work of Ergun et al. [1998], in which they propose the study of *Spot Checker*, which in some contexts coincides with property testing.

1.5. ORGANIZATION. The paper is organized in two parts. The first part focuses on property testing in general: A definition is given and discussed in Section 2. In Section 3, we explore the relations between property testing and learning. General observations regarding property testing appear in Section 4.

The second part of the paper focuses on testing graph properties: The basic framework is presented in Section 5. In Section 6, we present the Bipartiteness tester, and its generalization to a k -Colorability Tester. In Section 7, we present our tester for ρ -Clique. In Section 8, we present our tester for ρ -CUT, and its generalization to Bisection. The general graph-partition problem is treated in Section 9. We conclude, in Section 10, with comments regarding extensions and limitations of the above algorithms and problems, as well as discuss other graph properties.

Appendix A contains a list of some recurring notation, and Appendix B recalls standard probabilistic inequalities which are extensively used.

Part I: General Property Testing

2. Definitions

Let $\mathcal{F} = \{\mathcal{F}_n\}$ be a parameterized class of functions, where the functions¹⁰ in $\mathcal{F}_n$ are defined over $\{0, 1\}^n$ and let $\mathcal{D} = \{\mathcal{D}_n\}$ be a corresponding class of distributions (i.e., $\mathcal{D}_n$ is a distribution on $\{0, 1\}^n$). We use $x \sim \mathcal{D}_n$ to denote that x is distributed according to the distribution $\mathcal{D}_n$. We say that a function f defined on $\{0, 1\}^n$ is ϵ -close to $\mathcal{F}_n$ with respect to $\mathcal{D}_n$ if there exists a function $g \in \mathcal{F}_n$ such that

$$\Pr_{x \sim \mathcal{D}_n}[f(x) \neq g(x)] \leq \epsilon. \quad (1)$$

Otherwise, f is ϵ -far from $\mathcal{F}_n$ (with respect to $\mathcal{D}_n$).

We shall consider several variants of testing algorithms, where the most basic one is defined as follows:

Definition 2.1 (Property Testing). Let $\mathcal{A}$ be an algorithm which receives as input a size parameter n , a distance parameter $0 < \epsilon < 1$, and a confidence parameter $0 < \delta < 1/2$. Fixing an arbitrary function f and distribution $\mathcal{D}_n$ over $\{0, 1\}^n$, the algorithm is also given access to a sequence of f -labeled examples, $(x_1, f(x_1)), (x_2, f(x_2)), \dots$, where each x_i is independently drawn from the distribution $\mathcal{D}_n$. We say that $\mathcal{A}$ is a *property testing algorithm* (or simply a *testing algorithm*) for the class of functions $\mathcal{F}$ if for every n , ϵ and δ and for every function f and distribution $\mathcal{D}_n$ over $\{0, 1\}^n$ the following holds

- if $f \in \mathcal{F}_n$, then with probability at least $1 - \delta$ (over the examples drawn from $\mathcal{D}_n$ and the possible coins tosses of $\mathcal{A}$), $\mathcal{A}$ *accepts* f (i.e., outputs 1);
- if f is ϵ -far from $\mathcal{F}_n$ (with respect to $\mathcal{D}_n$), then with probability at least $1 - \delta$, $\mathcal{A}$ *rejects* f (i.e., outputs 0).

¹⁰ The range of these functions may vary and for many of the results and discussions it suffices to consider Boolean functions.

The *sample complexity* of $\mathcal{A}$ is a function of n , ϵ , and δ bounding the number of labeled examples examined by $\mathcal{A}$ on input (n, ϵ, δ) .

Though it was not stated explicitly in the definition, we shall usually also be interested in bounding the running time of a property testing algorithm (as a function of the parameters n , δ , ϵ , and in some case of a complexity measure of the class $\mathcal{F}$). We consider the following variants of the above definition:

- (1) D_n may be a specific distribution that is known to the algorithm. In particular, we shall be interested in testing with respect to the uniform distribution.
- (2) D_n may be restricted to a known class of distributions (e.g., product distributions).
- (3) The algorithm may be given access to an *oracle* for the function f , which when queries on $x \in \{0, 1\}^n$, returns $f(x)$. In this case, we refer to the number of queries made by $\mathcal{A}$ (which is a function of n , ϵ , and δ), as the *query complexity* of $\mathcal{A}$.
- (4) In some cases, the algorithm might have the additional feature that whenever it outputs *fail* it also provides a *certificate* to the fact that $f \notin F$. Certificates are defined with respect to a *verification algorithm* that *accepts* a sequence of labeled examples whenever there exists $f \in \mathcal{F}_n$ which is consistent with the sequence. (We do not require that the algorithm reject each sequence which is not consistent with some $f \in \mathcal{F}_n$.) A certificate for $f \notin \mathcal{F}_n$ is an f -labeled sequence which is rejected by the verification algorithm.
- (5) The algorithm may have only *one-sided error*. Namely, in case $f \in \mathcal{F}$, the algorithm *always* accepts f .
- (6) The algorithm is given two distance parameters, ϵ_1 and ϵ_2 , and is required to pass with high probability every f which is ϵ_1 -close to F_n , and fail every f which is ϵ_2 -far from F_n .

3. On the Relation between Property Testing and PAC Learning

A *Probably Approximately Correct* (PAC) learning algorithm [Valiant 1984] works in the same framework as that described in Definition 2.1 except for the following (crucial) differences:

- (1) It is given a *promise* that the unknown function f (referred to as the *target function*) belongs to $\mathcal{F}$;
- (2) It is required to output (with probability at least $1 - \delta$) a *hypothesis* function h which is ϵ -close to f , where closeness is as defined in Eq. (1) (and ϵ is usually referred to as the *accuracy* parameter).

Note that the differences pointed out above effect the tasks in opposite directions. Namely, the absence of a promise makes testing potentially harder than learning, whereas deciding whether a function belongs to a class rather than finding the function may make testing easier.

In the learning literature, a distinction is made between *proper* (or *representation dependent*) learning and *nonproper* learning [Pitt and Valiant 1988]. In the former model, the hypothesis output by the learning algorithm is required to belong to the same function class as the target function f , that is, $h \in \mathcal{F}$, while in

the latter model, $h \in \mathcal{H}$, for some hypothesis class $\mathcal{H}$ such that $\mathcal{F} \subseteq \mathcal{H}$. We assume that a proper learning algorithm (for $\mathcal{F}$) either halts without output or outputs a function in $\mathcal{F}$, but it never outputs any function not in $\mathcal{F}$.¹¹ There are numerous variants of PAC learning (including learning with respect to specific distributions, and learning with access to an oracle for the target function f (which in the case of Boolean functions is referred to as a *membership oracle*)). Unless stated otherwise, whenever we refer in this section to PAC learning we mean the *distribution-free no-query nonproper model* described above. The same is true for references to property testing. In addition, apart from one example, we shall restrict our attention to classes of Boolean functions.

3.1. TESTING IS NOT HARDER THAN PROPER LEARNING

PROPOSITION 3.1.1. *If a function class $\mathcal{F}$ has a proper learning algorithm $\mathcal{A}$, then $\mathcal{F}$ has a property testing algorithm $\mathcal{A}'$ with sample complexity*

$$m_{\mathcal{A}'}(n, \epsilon, \delta) = m_{\mathcal{A}}\left(n, \frac{\epsilon}{2}, \frac{\delta}{2}\right) + O\left(\frac{\log(1/\delta)}{\epsilon}\right),$$

where $m_{\mathcal{A}}(\cdot, \cdot, \cdot)$ is the sample complexity of $\mathcal{A}$. Furthermore, the same relation holds between the running times of the two algorithms.

PROOF. In order to test if $f \in \mathcal{F}$ or is ϵ -far from any function in $\mathcal{F}$, we first run the learning algorithm $\mathcal{A}$ with confidence parameter $\delta/2$, and accuracy parameter $\epsilon/2$, using random examples labeled by f . If $\mathcal{A}$ does not output a hypothesis, then we reject f . If $\mathcal{A}$ outputs a hypothesis h (which must be in $\mathcal{F}$ since $\mathcal{A}$ is a proper learning algorithm), then we approximate the distance between h and f by drawing an additional sample of size $O(\epsilon^{-1} \log(1/\delta))$. If the approximated distance is less than $3\epsilon/4$ then we accept; otherwise, we reject.

In case $f \in \mathcal{F}$, with probability at least $1 - \delta/2$, $\mathcal{A}$'s output, h , is $\epsilon/2$ -close to f , and an additive Chernoff bound (see Appendix B), tells us that with probability at least $1 - \delta/2$ over the additional sample, we shall not reject it. In case f is ϵ -far from $\mathcal{F}$, any hypothesis $h \in \mathcal{F}$ is at least ϵ -far from f , and with probability at least $1 - \delta/2$ over the additional sample, f is rejected. $\square$

In particular, the above proposition implies that if for every n , $\mathcal{F}_n$ has polynomial (in n) VC-dimension [Vapnik and Chervonenkis 1971; Blumer et al. 1989],¹² then $\mathcal{F}$ has a tester whose sample complexity is polynomial in n , $1/\epsilon$, and $\log(1/\delta)$. The reason is that classes with polynomial VC-dimension can be properly learned from a sample of the above size [Blumer et al. 1989]. However, the running time of such a proper learning algorithm, and hence of the resulting testing algorithm might be exponential in n .

¹¹ We remark that in case the functions in $\mathcal{F}$ have an easy to recognize representation, one can easily guarantee that the algorithm never outputs a function not in $\mathcal{F}$, by simply checking the hypothesis' representation. Standard classes considered in works on proper learning (e.g., Decision-Trees) typically have this feature.

¹² The Vapnik Chervonenkis (VC) dimension of a class $\mathcal{F}_n$ is defined to be the size d of the largest set $X \subseteq \{0, 1\}^n$ for which the following holds. For each (of the 2^d) partitions (X_0, X_1) of X there exists a function $f \in \mathcal{F}_n$ such that for every $x \in X_0$, $f(x) = 0$, and for every $x \in X_1$, $f(x) = 1$. A set X that has this feature is said to be *shattered* by $\mathcal{F}_n$.

COROLLARY 3.1.2. *Every class that is learnable with constant confidence using a sample of size $\text{poly}(n/\epsilon)$ (and thus has a $\text{poly}(n)$ VC dimension [Blumer et al. 1989]), is testable with a $\text{poly}(n/\epsilon) \cdot \log(1/\delta)$ sample (in at most exponential time).*

3.2. TESTING MAY BE HARDER THAN LEARNING. In contrast to Proposition 3.1.1 and to Corollary 3.1.2, we show that there are classes that are efficiently learnable (though not by a proper learning algorithm) but are not efficiently testable. This is proven by observing that many hardness results for proper learning (cf. [Pitt and Valiant 1988; Blum and Rivest 1989; Pitt and Warmuth 1993]) actually establish the hardness of testing (for the same classes). Furthermore, we believe that it is more natural to view these hardness results as referring to testing and derive the hardness for proper learning via Proposition 3.1.1. Thus, the separation between efficient learning and efficient proper learning translates to a separation between efficient learning and efficient testing.

PROPOSITION 3.2.1. *If $\mathcal{NP} \not\subseteq \mathcal{BPP}$, then there exist function classes that are not $\text{poly}(n/\epsilon)$ -time testable but are $\text{poly}(n/\epsilon)$ -time (nonproperly) learnable.*

PROOF. The proposition follows from the fact that many of the representation dependent hardness results¹³ have roughly the following form. An NP-complete problem is reduced to the following decision problem: Given a set S of labeled examples, does there exist a function in $\mathcal{F}$ that is consistent with S ? A learning algorithm is forced to find a consistent function if one exists by letting the support of the distribution D (which is allowed to be arbitrary) lie solely on S , and setting ϵ to be smaller than $1/|S|$. Actually, since the consistency problem is that of deciding if there exists a consistent function and not necessarily of finding such a function, it follows that the corresponding testing problem (using the same D and ϵ) is hard as well. Details follow.

Let $\mathcal{F}$ be a fixed class of functions and suppose that the following decision problem is NP-complete

input. A sequence $(x_1, \sigma_1), \dots, (x_t, \sigma_t)$, where $x_i \in \{0, 1\}^n$, $\sigma_i \in \{0, 1\}$, and $t = \text{poly}(n)$.

question. Is there a function $f \in \mathcal{F}_n$ so that $f(x_i) = \sigma_i$, for all $i \in [t] \stackrel{\text{def}}{=} \{1, \dots, t\}$.

Assuming that there exists a $\text{poly}(n/\epsilon)$ -time property testing algorithm, denoted $\mathcal{A}$, for the class $\mathcal{F}$, we construct a polynomial-time decision procedure for the above problem (contradicting the assumption that $\mathcal{NP} \not\subseteq \mathcal{BPP}$). We invoke $\mathcal{A}$ with parameters n , ϵ and δ , where $\epsilon < 1/t$ (say $\epsilon = 1/(2t)$), and say $\delta = 1/3$. Suppose that $\mathcal{A}$ requires $m \stackrel{\text{def}}{=} m_{\mathcal{A}}(n, \epsilon, \delta)$ samples. We uniformly select m indices, denoted $i_1, \dots, i_m$, (possibly with repetitions) out of $[t]$ and feed $\mathcal{A}$ with the labeled sample $(x_{i_1}, \sigma_{i_1}), \dots, (x_{i_m}, \sigma_{i_m})$. We decide according to $\mathcal{A}$'s output.

We analyze the performance of our algorithm by relating it to the performance of the property testing algorithm on the uniform distribution over the set $\{x_i : i \in [t]\}$. Suppose first that there exists $f \in \mathcal{F}_n$ so that $f(x_i) = \sigma_i$, for all $i \in [t]$.

¹³ See, for example, Gold [1998], Angluin [1978], Pitt and Valiant [1988], Blum and Rivest [1989], and Pitt and Warmuth [1993].

In this case, we provide $\mathcal{A}$ with a random sample labeled by a function in $\mathcal{F}_n$ and thus with probability at least $1 - \delta$ the test must accept. Thus, our decision procedure accepts yes-instances with probability at least $1 - \delta$.

Suppose now that there exists no function $f \in \mathcal{F}_n$ such that $f(x_i) = \sigma_i$, for all $i \in [t]$. This implies that the function f defined by $f(x_i) \stackrel{\text{def}}{=} \sigma_i$, for all $i \in [t]$ (and $f(x) = 0$ for $x \notin \{x_i : i \in [t]\}$), is at distance at least $1/t > \epsilon$ from $\mathcal{F}$ (with respect to the uniform distribution over $\{x_i : i \in [t]\}$). Since we provide $\mathcal{A}$ with a random sample labeled by this f , the test must reject with probability at least $1 - \delta$. Hence, our decision procedure rejects no-instances with probability at least $1 - \delta$.

This establishes, in particular, that testing the class of k -Term DNF (where k is a constant) is NP-Hard (see Pitt and Valiant [1988]). On the other hand, k -Term DNF (for constant k) is efficiently learnable (using the hypothesis class of k -CNF) [Valiant 1984; Pitt and Valiant 1988]. $\square$

We stress that whereas Proposition 3.1.1 generalizes to learning and testing under specific distributions, and to learning and testing with queries, the proof of Proposition 3.2.1 uses the premise that the testing (or proper learning) algorithm works for *any* distribution and does *not* make queries.

3.3. TESTING MAY BE EASIER THAN LEARNING. We start by presenting a function class that is easy to test but cannot be learned with polynomial sample complexity, regardless of the running-time.

PROPOSITION 3.3.1. *There exist function classes $\mathcal{F}$ such that:*

- $\mathcal{F}$ has a property testing algorithm whose sample complexity and running time are $O(\epsilon^{-1} \cdot \log(1/\delta))$ (i.e., independent of n);
- Any learning algorithm for $\mathcal{F}$ must have sample complexity exponential in n .

PROOF. It is possible to come up with quite a few examples of functions classes for which the above holds. We give one example below. For each n let $\mathcal{F}_n$ include all functions f over $\{0, 1\}^n$, such that for every $y \in \{0, 1\}^{n-1}$, $f(1y) = 1$ (and if the first bit of the input is 0 then no restriction is made).

Given $m = O(\epsilon^{-1} \log(1/\delta))$ examples, labeled by an unknown f and drawn according to an arbitrary distribution D_n , the testing algorithm will simply verify that for all examples x whose first bit is 1, $f(x) = 1$. If $f \in \mathcal{F}_n$, it will always accept it, and if f is ϵ -far from $\mathcal{F}_n$ (with respect to D_n) then the probability that it does not observe even a single example of the form $(1y, 0)$ (and as a consequence, accepts f) is bounded by $(1 - \epsilon)^m < \delta$. On the other hand, the VC-dimension of $\mathcal{F}_n$ is 2^{n-1} (since the set $\{0y : y \in \{0, 1\}^{n-1}\}$ is shattered by $\mathcal{F}_n$). By Blumer et al. [1989], learning this class requires a sample of size $\Omega(2^n)$. $\square$

The impossibility of learning the function class in Proposition 3.3.1 is due to its exponential VC-dimension, (i.e., it is a pure information theoretic consideration). We now turn to function classes of exponential (rather than double exponential) size. Such classes are always learnable with a polynomial sample; the question is whether they are learnable in polynomial-time. We present a function class that is easy to test but cannot be learned in polynomial-time (even under the uniform distribution), provided certain trapdoor one-way permutations

exist (e.g., factoring is intractable). We start by defining this assumption, which in some sense is weaker than the standard one (cf., Goldreich [1995, Sec. 2.4]). (We comment that the standard term “trapdoor permutation” is somewhat misleading since what is being defined is a collection of permutations.)

Definition 3.3.2. (Weak Trapdoor Permutations with Dense Domains). Let $\{p_\alpha: D_\alpha \mapsto D_\alpha\}_{\alpha \in \{0,1\}^*}$ be a family of permutations. We say that this family has *dense domains* if for some positive polynomial $q(\cdot)$ and all α 's, both $D_\alpha \subseteq \{0,1\}^{|\alpha|}$ and $|D_\alpha| \geq 2^{|\alpha|}/q(|\alpha|)$ hold. The family is *one-way* if it satisfies

- (1) (*easy to evaluate*). There exists a polynomial-time algorithm which given α and x outputs $p_\alpha(x)$.
- (2) (*hard to invert*). No probabilistic polynomial-time algorithm can, given uniformly chosen α and x , output $p_\alpha^{-1}(x)$ with, say, success probability at least $1/n$ (where the probability is taken uniformly over the coin tosses of the algorithm and all possible choices of $\alpha \in \{0,1\}^n$ and $x \in D_\alpha$).

The family is said to have *weak trapdoors* if for each $\alpha \in \{0,1\}^*$ (equivalently, for each p_α) there exists a $\text{poly}(|\alpha|)$ -size circuit which inverts p_α .

The weak trapdoor condition is a relaxation of the standard condition which requires also that one can efficiently generate $(\alpha, \text{trapdoor})$ -pairs. The dense domain condition is indeed nonstandard, but does hold for all popular candidates (e.g., RSA and Rabin functions). See discussion in Canetti et al. [1996]. Another nonstandard aspect of the definition is associating a permutation with each string, whereas the standard definition associates permutations only with a subset of all strings. Again, we note that in case of the popular candidates, permutations are associated with a polynomial fraction of all strings of certain length, and so we can modify these constructions so that a permutation is associated with each string.¹⁴

PROPOSITION 3.3.3. *If there exist weak trapdoor one-way permutations with dense domains, then there exists a family of functions that can be tested in $\text{poly}(n/\epsilon)$ -time but cannot be learned in $\text{poly}(n/\epsilon)$ -time, even with respect to the uniform distribution. Furthermore, the functions can be computed by $\text{poly}(n)$ -size circuits.*

PROOF. By Canetti et al. [1996], any collection of trapdoor one-way permutations with dense domains can be converted into a collection of trapdoor one-way permutations where the domain of each p_α is $\{0,1\}^{|\alpha|}$. Thus, we use such a collection

$$\{p_\alpha: \{0,1\}^{|\alpha|} \mapsto \{0,1\}^{|\alpha|}\}_{\alpha \in \{0,1\}^*}$$

We consider the function class $\mathcal{OW} = \{\mathcal{OW}_n\}$, where $\mathcal{OW}_n$ consists of the multivalued functions f_α , so that $f_\alpha(x) \stackrel{\text{def}}{=} (\alpha, p_\alpha^{-1}(x))$ for every $\alpha, x \in \{0,1\}^n$.

¹⁴ One way of doing so proceeds in two steps, and is analogous to the Canetti et al. [1996] transformation mentioned subsequently. First, one introduces dummy permutations for each string that is not associated with a permutation in the original construction. This may weaken the one-way property but still some “one-wayness” remains. Next, one amplifies the one-wayness by taking a “direct product”—that is, for each $\alpha_1, \dots, \alpha_t$ and $x_1, \dots, x_t$, we define a new function indexed by the α -sequence that consists of the concatenation of the values of each function associated with an α_i evaluated at the corresponding x_i .

Using the weak trapdoor condition, we know that the functions in $\mathcal{W}_n$ can be computed by $\text{poly}(n)$ -size circuits.

To test if $f \in \mathcal{W}$, we merely examine sufficiently many f -labeled examples. Specifically, $m \stackrel{\text{def}}{=} O(\epsilon^{-1} \log(1/\delta))$ examples will do. For each labeled example, $(x, (\alpha, y))$, if $p_\alpha(y) \neq x$, then we reject f . In addition, we also reject if we see two examples, $(x_1, (\alpha_1, y_1))$ and $(x_2, (\alpha_2, y_2))$, so that $\alpha_1 \neq \alpha_2$. Otherwise, we accept f .

We show that this test works for any distribution on the examples and so the class $\mathcal{W}$ is efficiently testable (in a distribution-free sense). Clearly, the test always accepts $f_\alpha \in \mathcal{W}_n$. Assuming that f is ϵ -far from $\mathcal{W}_n$, with respect to some distribution D , we wish to show that f is rejected by the test with high probability. We consider two cases.

Case 1. Suppose there exists an α such that the probability that the first element of $f(x)$ equals α is at least $1 - \epsilon/2$, where the probability is taken over $x \sim D$. Since f is ϵ -far from $\mathcal{W}_n$, we have in particular

$$\Pr_{x \sim D_n}[f(x) = f_\alpha(x)] < 1 - \epsilon$$

and so

$$\begin{aligned} q &\stackrel{\text{def}}{=} \Pr_{x \sim D_n}[x \neq p_\alpha(y) \text{ and } f(x) = (\alpha, y) \text{ for some } y] \\ &= \Pr_{x \sim D_n}[f(x) = (\alpha, y) \text{ for some } y] - \Pr_{x \sim D_n}[f(x) = f_\alpha(x)] \\ &> \left(1 - \frac{\epsilon}{2}\right) - (1 - \epsilon) = \frac{\epsilon}{2}. \end{aligned}$$

Hence, the test accepts with probability at most $(1 - q)^m < \delta$.

Case 2. Suppose that for every α the probability that the first element of $f(x)$ equals α is at most $1 - (\epsilon/2)$. Then, the probability that all m examples have a label starting with the same α is at most

$$\sum_{\alpha} \Pr_{x \sim D_n}[\exists y \text{ such that } f(x) = (\alpha, y)]^m,$$

which is bounded above by $(1 - (\epsilon/2))^{m-1} < \delta$.

Thus, in both cases, the test rejects ϵ -far functions with sufficiently high probability. We now turn to show that it is infeasible to learn the class $\mathcal{W}$ under the uniform distribution. This is done by using any efficient learning algorithm, $\mathcal{A}$, in order to construct an algorithm that contradicts the one-way condition. The inverting algorithm operates as follows, on inputs α and x , $p_\alpha(x)$. First it uniformly generates an f_α -labeled sample $\{(x_i, (\alpha, y_i))\}$ for $\mathcal{A}$. This is done by uniformly selecting y_i and setting $x_i = p_\alpha(y_i)$. (Here we use the efficient evaluation property of the collection.) Note that indeed $f_\alpha(x_i) = (\alpha, p_\alpha^{-1}(x_i)) = (\alpha, y_i)$, and that x_i is uniformly distributed. When the learning stage of $\mathcal{A}$ is over, the inverter supplies $\mathcal{A}$ with x , asking for its label. With probability at least $(1 - \delta) \cdot (1 - \epsilon) > 1 - \delta - \epsilon$, taken uniformly over all possible $x \in \{0, 1\}^{|\alpha|}$ and the internal coin tosses of both the inverter and $\mathcal{A}$, algorithm $\mathcal{A}$ returns the correct label; that is, $(\alpha, p_\alpha^{-1}(x))$. Thus, an efficient learning algorithm is transformed

into an efficient inverting algorithm for the family, in contradiction to the one-wayness condition.¹⁵ $\square$

The class presented in Proposition 3.3.3 consists of multi-valued functions. We leave it as an open problem whether a similar result holds for a class of Boolean functions.

3.4 VARIANTS

3.4.1. Learning and Testing with Queries (under the Uniform Distribution). Let the class of parity functions, $\mathcal{PAR} = \{\mathcal{PAR}_n\}$, where $\mathcal{PAR}_n \stackrel{\text{def}}{=} \{f_S: S \subseteq \{1, \dots, n\} \text{ and } f_S: \{0, 1\}^n \mapsto \{0, 1\} \text{ so that } f_S(x) = \sum_{i \in S} x_i \bmod 2\}$. Work on linearity testing,¹⁶ culminating in the result of Bellare et al. [1995a], implies that there exists a testing algorithm for the class $\mathcal{PAR}$, under the uniform distribution, whose query complexity is $O(\epsilon^{-1} \log(1/\delta))$. The running-time is a factor n bigger, merely needed to write down the queries. On the other hand, any learning algorithm for this class must use at least n queries (or examples). The reason is that any query (or example) gives rise to a single linear constraint on the coefficients of the linear function, and with less than n such constraints the function is not uniquely defined. Furthermore, every two linear functions disagree with probability $1/2$ on a uniformly chosen input.

An example of a $\text{poly}(n/\epsilon)$ -time testable (with queries) class which is not learnable with $\text{poly}(n/\epsilon)$ queries is the class of multi-variate (i.e., n -variate) polynomials. Specifically, let $\mathcal{POLY} = \{\mathcal{POLY}_n\}$, where $\mathcal{POLY}_n$ consists of n -variate polynomials of total degree n over the field $\text{GF}(q)$, where q is the first prime in the interval $[n^4, 2n^4]$. Work on low-degree testing [Babai et al. 1991a; 1991b; Gemmell et al. 1991], culminating in the result of Rubinfeld and Sudan [1996], implies that there exists a testing algorithm for the class $\mathcal{POLY}$, under the uniform distribution, whose query complexity is $O(\min\{n^3, n/\epsilon\} \cdot \log(1/\delta))$. The running-time is a factor $O(n \log n)$ bigger, merely needed to write down the queries and do some simple algebra. It is not hard to show that one cannot possibly learn $\mathcal{POLY}_n$, under the uniform distribution, using only $\text{poly}(n)$ queries. Again, the reason is that any query (or example) gives rise to a single linear constraint on the coefficients of the polynomial. Since there are exponentially (in n) many coefficients, this leaves the polynomial not uniquely defined. Finally, one invokes Schwarz's Lemma [Schwartz 1980], by which two such degree n polynomials can agree on at most $n/q < 1/n^3$ fraction of the domain.

3.4.2. Agnostic Learning and Testing. In a variant of PAC learning, called *Agnostic PAC learning* [Kearns et al. 1992], there is no promise concerning the target function f . Instead, the learner is required to output a hypothesis h from a certain hypothesis class $\mathcal{H}$, such that h is ϵ -close to the function in $\mathcal{H}$ which is closest to f . The absence of a promise makes agnostic learning closer in spirit to property testing than basic PAC learning. However, we were not able to translate this similarity in spirit to anything stronger than the relation between testing and proper learning. Namely, since agnostic learning with respect to a hypothesis

¹⁵ Actually, the inverter is stronger than what is required to contradict one-wayness: It can invert any p_α on all but a small fraction of the range elements.

¹⁶ See, for example, Blum et al. [1993], Babai et al. [1991b], Feige et al. [1991], Bellare et al. [1993], and Bellare and Sudan [1994].

class $\mathcal{H}$ implies proper learning of the class $\mathcal{H}$, it also implies property testing of $\mathcal{H}$.

3.4.3. Learning and Testing Distributions. A *distribution learning* algorithm for a class of distributions, $\mathcal{D} = \{\mathcal{D}_n\}$, receives (in addition to the parameters n , ϵ and δ) an (unlabeled) sample of strings in $\{0, 1\}^n$, distributed according to an unknown distribution $D \in \mathcal{D}_n$. The algorithm is required to output a distribution D' (either in form of a machine that generates strings according to D' , or in form of a machine that on input $x \in \{0, 1\}^n$ outputs $D'(x)$), such that with probability at least $1 - \delta$, the variation distance between D and D' is at most ϵ . (For further details, see Kearns et al. [1994].) In contrast, a *distribution testing* algorithm, upon receiving a sample of strings in $\{0, 1\}^n$ drawn according to an unknown distribution D , is required to accept D , with probability at least $1 - \delta$, if $D \in \mathcal{D}_n$, and to reject (with probability $\geq 1 - \delta$) if D is ϵ -far from $\mathcal{D}_n$ (with respect to the variation distance).

The context of learning and testing distributions offers another demonstration to the importance of a promise (i.e., the fact that the learning algorithm is required to work only when the target belongs to the class, whereas the testing algorithm needs to work for all targets that are either in the class or far away from it).

PROPOSITION 3.1.5.1. *There exist distribution classes that are efficiently learnable (in both senses mentioned above) but cannot be tested with a subexponential sample (regardless of the running-time).*

PROOF. Consider the class of distributions $\mathcal{D} = \{\mathcal{D}_n\}$ consisting of all distributions, D_n^p , which are generated by n independent tosses of a coin with bias p . Clearly, this class can be efficiently learned (by merely approximating the bias p of the target distribution). However, a tester cannot distinguish the case in which a sample of subexponential size is taken from the uniform distribution $D_n^{1/2}$ (and thus should be accepted), and the case in which such a sample is taken from a ‘typically bad’ distribution B_n^S which is uniform over $S \subset \{0, 1\}^n$, where $|S| = 2^{n-1}$. Formally, we consider the behavior of the test when given a sample from $D_n^{1/2}$ versus its behavior when given a sample from B_n^S , where S is uniformly chosen among all subsets of size 2^{n-1} . $\square$

We note that the above proof holds for any distribution class that contains the uniform distribution and is far from distributions such as the B_n^S s.

4. General Observations

4.1. PROPERTY TESTING MAY BE VERY HARD

PROPOSITION 4.1.1. *There exists a function class $\mathcal{F} = \{\mathcal{F}_n\}$ for which any testing algorithm must inspect the value of the function at a constant fraction of the inputs (i.e., on $\Omega(2^n)$ inputs). This holds even for testing with respect to the uniform distributions, for any constant distance parameter $\epsilon < 1/2$ and confidence parameter $\delta < 1/2$, and even when allowing the algorithm to make queries and use unlimited computing time.*

PROOF. Suppose for simplicity that $\epsilon = 1/4$ and $\delta = 1/5$. The proof easily generalizes to general constant ϵ , $\delta < 1/2$ by appropriately modifying the size of

$\mathcal{F}_n$ (which is defined below). We will use the Probabilistic Method to demonstrate the existence of a function class, $\mathcal{F} = \{\mathcal{F}_n\}$, satisfying the claim, so that $\mathcal{F}_n$ consists of $2^{(1/10) \cdot 2^n}$ Boolean functions operating on $\{0, 1\}^n$. We'll show that there exists a class $\mathcal{F}$ so that a uniformly selected function is both far from it and indistinguishable from it when observing $O(2^n)$ values.

First we show that, with high probability, a uniformly selected function $g: \{0, 1\}^n \mapsto \{0, 1\}$ is ϵ -far from any set, $\mathcal{F}_n$, of $2^{(1/10) \cdot 2^n}$ functions. Let $N \stackrel{\text{def}}{=} 2^n$ and U_N be uniformly distributed on $\{0, 1\}^N$. Then,

$\Pr_g[g \text{ is } \epsilon\text{-close to } \mathcal{F}_n]$

$$\begin{aligned} &\leq \Pr_g[\exists f \in \mathcal{F}_n \text{ such that } g(x) \neq f(x) \text{ for less than } \epsilon N \text{ } x\text{'s}] \\ &\leq |\mathcal{F}_n| \cdot \Pr[U_N \text{ has at most } \epsilon N \text{ } 1\text{'s}] \\ &\leq 2^{N/10} \cdot 2 \exp(-N/8) \\ &= \exp(-\Omega(N)). \end{aligned}$$

Thus, with overwhelmingly high probability (over the choices of g), the function g is ϵ -far from the class $\mathcal{F}_n$. We now consider any fixed sequence, S , of $T \stackrel{\text{def}}{=} N/20$ inputs and compare the values assigned to them by (a random) g versus the values assigned to them by a uniformly chosen function in $\mathcal{F}_n$. Clearly, in the first case the values are uniformly distributed. Let $\delta_S(\mathcal{F}_n)$ denote the statistical difference between the uniform distribution and the distribution of function-values (on the inputs in S) induced by a uniformly selected function in $\mathcal{F}_n$. That is,

$$\delta_S(\mathcal{F}_n) \stackrel{\text{def}}{=} \frac{1}{2} \cdot \sum_{\alpha \in \{0,1\}^T} |\Pr_{f \in \mathcal{F}_n}[f(S) = \alpha] - 2^{-T}|,$$

where for a sequence $S = x_1, \dots, x_T$, $f(S) = f(x_1) \cdots f(x_T)$. We consider the probability, taken over all possible choices of $\mathcal{F}_n$ (consisting of $2^{(1/10) \cdot 2^n}$ functions), that $\delta_S(\mathcal{F}_n) > 1/2$. By using a multiplicative Chernoff bound (with multiplicative factors $1/2$ and $3/2$ —see Appendix B), we get

$$\begin{aligned} &\Pr_{\mathcal{F}_n} \left[\delta_S(\mathcal{F}_n) > \frac{1}{2} \right] \\ &\leq \Pr_{\mathcal{F}_n} [\exists \alpha \in \{0, 1\}^T \text{ such that } |\Pr_{f \in \mathcal{F}_n}(f(S) = \alpha) - 2^{-T}| > 2^{-(T+1)}] \\ &\leq 2^T \cdot 2 \exp \left(-\frac{1}{3} \cdot \left(\frac{1}{2} \right)^2 \cdot 2^{-T} |\mathcal{F}_n| \right) \\ &= 2^{N/20} \cdot 2 \exp \left(-\frac{1}{12} \cdot 2^{N/20} \right) \\ &= \exp(-2^{\Omega(N)}) \end{aligned}$$

Summing the probabilities over all possible $\binom{N}{T} < 2^N$ sequences we conclude that with overwhelmingly high probability, over the choice of $\mathcal{F}_n$, all δ_g 's are bounded above by $1/2$. Let us fix such an $\mathcal{F}_n$. Consequently, the difference between the acceptance probability of a truly random g and the acceptance probability of a uniformly selected $f \in \mathcal{F}_n$ is at most $1/2$. This does not allow one both to accept every $f \in \mathcal{F}_n$ with probability at least 0.8 and accept a random g with probability at most 0.21 (and so to reject ϵ -far functions with probability at least 0.8 —the extra 0.01 over-compensates for the case that a random g is ϵ -close to $\mathcal{F}_n$). $\square$

4.2. THE ALGEBRA OF PROPERTY TESTING. Suppose that two function classes are testable within certain complexity. What can we infer about their INTERSECTION, UNION, and COMPLEMENT? Unfortunately, in general, we can only say that their union is testable within comparable complexity. That is,

PROPOSITION 4.2.1. *Let $\mathcal{F}' = \{\mathcal{F}'_n\}$ and $\mathcal{F}'' = \{\mathcal{F}''_n\}$ be function classes testable within complexities $c'(n, \epsilon, \delta)$ and $c''(n, \epsilon, \delta)$, respectively. Then, the class $\mathcal{F} = \{\mathcal{F}_n\}$, where $\mathcal{F}_n = \mathcal{F}'_n \cup \mathcal{F}''_n$, is testable within complexity $c(n, \epsilon, \delta) = c'(n, \epsilon, (\delta/2)) + c''(n, \epsilon, (\delta/2))$.*

PROOF. The testing algorithm for $\mathcal{F}$ consists of testing for membership in both $\mathcal{F}'$ and $\mathcal{F}''$ and accepting if and only if at least one of the tests has accepted. The validity of this test relies on the fact that if f is ϵ -far from $\mathcal{F} = \mathcal{F}' \cup \mathcal{F}''$, then it is ϵ -far from both $\mathcal{F}'$ and $\mathcal{F}''$. $\square$

The fact that a claim analogous to the one used in the above proof does not hold for intersection is the reason that an analogous tester does not work for the intersection class. That is, it may be the case that f is far from $\mathcal{F} = \mathcal{F}' \cap \mathcal{F}''$ and yet it is very close to both $\mathcal{F}'$ and $\mathcal{F}''$. Thus, a function (close to both $\mathcal{F}'$ and $\mathcal{F}''$), may pass both the corresponding property tests, but still may be far from $\mathcal{F}$. In particular

PROPOSITION 4.2.2. *There exist function classes $\mathcal{F}' = \{\mathcal{F}'_n\}$ and $\mathcal{F}'' = \{\mathcal{F}''_n\}$ such that both are trivially testable under the uniform distribution (i.e., by an oblivious algorithm that always accepts provided $\epsilon > 2^{-n}$), whereas the class $\mathcal{F} = \{\mathcal{F}_n\}$, where $\mathcal{F}_n = \mathcal{F}'_n \cap \mathcal{F}''_n$ is not testable under the uniform distribution with query complexity $o(2^n)$ even for constant $\epsilon, \delta < 1/2$.*

PROOF. Let $\mathcal{F} = \{\mathcal{F}_n\}$ be as guaranteed in Proposition 4.1.1, and let $\mathcal{F}'_n \stackrel{\text{def}}{=} \mathcal{F}_n \cup \{f : f(0^n) = 0\}$, $\mathcal{F}''_n \stackrel{\text{def}}{=} \mathcal{F}_n \cup \{f : f(0^n) = 1\}$. Clearly both $\mathcal{F}'$ and $\mathcal{F}''$ are testable as claimed since every function is 2^{-n} -close to both $\mathcal{F}'$ and $\mathcal{F}''$. On the other hand, $\mathcal{F}'_n \cap \mathcal{F}''_n = \mathcal{F}_n$, and so the negative result of Proposition 4.1.1 applies. $\square$

Finally, we observe that property testing is not preserved under complementation. That is,

PROPOSITION 4.2.3. *There exists a function class $\mathcal{G} = \{\mathcal{G}_n\}$ such that for every*

$$\epsilon \geq \frac{1}{2^n - 1}$$

the algorithm that accepts every function is a tester for $\mathcal{G}$, while the class of functions not in $\mathcal{G}$ is not testable in subexponential complexity.

PROOF. Consider the function class $\mathcal{F}$ used in the proof of Proposition 4.1.1. As shown there, this class is not testable in subexponential complexity. Furthermore, as we show subsequently, $\mathcal{F}$ can be chosen so that for every n and for every function $f \in \mathcal{F}_n$, there exists at most one function $f' \in \mathcal{F}_n$, such that f and f' differ on exactly one input. In other words, there exist at least $2^n - 1$ functions not in $\mathcal{F}_n$ that differ from f on exactly one input. Let $\mathcal{G} = \{\mathcal{G}_n\}$ consist of all functions not in $\mathcal{F}$. Based on our additional claim (concerning differences between functions in $\mathcal{F}_n$), as long as $\epsilon \geq 1/(2^n - 1)$, the trivial algorithm which accepts all functions constitutes a tester for $\mathcal{G}$. This is true since for every $f \notin \mathcal{G}_n$ (i.e., $f \in \mathcal{F}_n$), and for every distribution, there exists some function $g \in \mathcal{G}_n$ that is at distance less than $1/(2^n - 1)$ from f . Namely, the total probability mass of the x 's for which there exists $g \in \mathcal{G}_n$ such that g and f differ only on x is at most 1, and there are at least $2^n - 1$ such x 's.

To prove the claim regarding differences between functions in $\mathcal{F}$, we again use the probabilistic method. Observe first that the probability that a particular function f belongs to $\mathcal{F}_n$, is $2^{N/10}/2^N = 2^{-0.9 \cdot N}$, where $N \stackrel{\text{def}}{=} 2^n$. Thus, for any fixed function f , the probability that there exist two functions f_1 and f_2 in $\mathcal{F}_n$ such that both f_1 and f_2 differ from f on exactly one input is bounded by $N^2 \cdot (2^{-0.9 \cdot N})^2 < 2^{-1.7 \cdot N}$. Finally the probability that this event occurs for some f is at most $2^N \cdot 2^{-1.7 \cdot N} = 2^{-0.7 \cdot N}$, which is extremely small. By adding the negligible probability (bounded in Proposition 4.1.1) that the uniformly chosen $\mathcal{F}_n$ is not hard to test, Proposition 4.2.3 follows. $\square$

PART II: Testing Graph Properties

5. Testing Graph Properties—Preliminaries

In the following sections, we concentrate on testing graph properties using queries and with respect to the uniform distribution. In Section 10.1, we discuss some extensions beyond this basic model. We start by defining the basic model and giving an overview of our results.

5.1. GENERAL GRAPH NOTATION. We consider undirected, simple graphs (no multiple edges or self-loops). For a simple graph G , we denote by $V(G)$ its vertex set and assume, without loss of generality, that $V(G) = \{1, \dots, |V(G)|\}$. Graphs are represented by their (symmetric) adjacency matrix. Thus, graphs are associated with the Boolean function corresponding to this matrix (i.e., the value of a pair $(u, v) \in V(G) \times V(G)$ indicates whether $(u, v) \in E(G)$). This brings us to associated undirected graphs with directed graphs where each edge in the undirected graph is associated with a pair of anti-parallel edges. Specifically, for a graph G , we denote by $E(G)$ the set of ordered pairs which correspond to edges in G (i.e., $(u, v) \in E(G)$ iff there is an edge between u and v in G). In the sequel, whenever we say ‘edge’ we mean a directed edge, and the degree of a vertex is the total number of edges incident to it (i.e., the sum of its in-degree and out-degree). For two (not necessarily disjoint) sets of vertices, X_1 and X_2 , we

let

$$E(X_1, X_2) \stackrel{\text{def}}{=} \{(u, v) \in E(G) : u \in X_1, v \in X_2 \text{ or } u \in X_2, v \in X_1\}.$$

The distance between two N -vertex graphs, G_1 and G_2 , is defined as the fraction of entries $(u, v) \in [N]^2$ ($[N] \stackrel{\text{def}}{=} \{1, \dots, N\}$), among all N^2 entries, which are in the symmetric difference of $E(G_1)$ and $E(G_2)$. Namely,

$$\text{dist}(G_1, G_2) \stackrel{\text{def}}{=} \frac{|(E(G_1) \setminus E(G_2)) \cup (E(G_2) \setminus E(G_1))|}{N^2}.$$

This notation is extended naturally to a set, $\mathcal{C}$, of N -vertex graphs; that is,

$$\text{dist}(G, \mathcal{C}) \stackrel{\text{def}}{=} \min_{G' \in \mathcal{C}} \{\text{dist}(G, G')\}$$

Another notation used extensively in subsequent sections is the set of neighbors of a vertex v , that is, $\Gamma(v) \stackrel{\text{def}}{=} \{u : (v, u) \in E(G)\}$. This notation is extended to sets of vertices in the natural manner; i.e., $\Gamma(S) \stackrel{\text{def}}{=} \bigcup_{v \in S} \Gamma(v)$.

5.2. OUR ALGORITHMS AND THEIR ANALYSES. We present testers for Bipartiteness, k -Colorability (for $k \geq 3$), ρ -Clique, ρ -Cut (and ρ - k -Cut), ρ -Bisection and the General Graph Partition property. The latter generalizes all the former ones, but yields worse complexity bounds for the special cases. Also, the testers for ρ -Cut and ρ - k -Cut, which actually work by approximating the size of the maximum cut, generalize Bipartiteness and k -Colorability, respectively. However, the former have two-sided error probability and higher complexity, whereas the latter can be directly tested with one-sided error probability.

For all properties, we also present linear (in N) time algorithms that, given a graph that has the property, find a partition that is approximately-good with respect to that property (e.g., in the case of ρ -Clique, the algorithm finds a set of size ρN that misses few edges to being a clique). While these algorithms are clearly beneficial when the corresponding approximation problems are known to be hard, they are also meaningful in the case of Bipartiteness (where an exact efficient algorithm is known). This is true since in order to partition a dense bipartite graph, one has to work in time quadratic in N (i.e., linear in the number of edges), while we offer an approximation algorithm that works in time linear in N . The complexities of our algorithms are summarized in Figure 1.

5.2.1. Common Themes. As mentioned in the introduction, our algorithms and their analyses share a few themes. In particular, all algorithms uniformly select a set of vertices and perform edge-queries only on pairs of vertices in the sample. In all cases the *natural* algorithm for the property, which checks whether the property (approximately) holds on the subgraph induced by the sample of vertices, is indeed a testing algorithm for the property. However, only in the case of Bipartiteness and k -Colorability do we analyze directly the natural algorithm. For the other properties, we present and analyze more complex algorithms. However, the correctness of the natural algorithms follows from the correctness of the algorithms that we present, and the resulting complexities of the former are only slightly higher than those of the latter.

PROPERTY	TESTING		PARTITIONING Time (on top of testing)
	Query	Time	
Bipartiteness	$O\left(\frac{\Delta^2}{\epsilon^3}\right)$	$O\left(\frac{\Delta^2}{\epsilon^3}\right)$	$O\left(\frac{\Delta}{\epsilon}\right) \cdot N$
k -Colorability	$O\left(\frac{k^4 \cdot \Delta_k^2}{\epsilon^6}\right)$	$\exp\left(O\left(\frac{k^2 \cdot \Delta_k^2}{\epsilon^3}\right)\right)$	$O\left(\frac{\Delta_k}{\epsilon}\right) \cdot N$
ρ -Clique	$O\left(\frac{\Delta^2 \cdot \rho^2}{\epsilon^6}\right)$	$\exp\left(O\left(\frac{\Delta \cdot \rho}{\epsilon^2}\right)\right)$	$O\left(\frac{\Delta^2}{\epsilon^2}\right) \cdot N$
ρ -Cut	$O\left(\frac{\Delta^2}{\epsilon^4}\right)$	$\exp\left(O\left(\frac{\Delta}{\epsilon^3}\right)\right)$	$O\left(\frac{\Delta}{\epsilon^2}\right) \cdot N$
ρ - k -Cut	$O\left(\frac{\Delta^2}{\epsilon^4}\right)$	$\exp\left(O\left(\frac{\Delta^2}{\epsilon^3}\right)\right)$	$O\left(\frac{\Delta_k}{\epsilon}\right) \cdot N$
ρ -Bisection	$O\left(\frac{\Delta^2}{\epsilon^8}\right)$	$\exp\left(O\left(\frac{\Delta}{\epsilon^3}\right)\right)$	$O\left(\frac{\Delta}{\epsilon^2}\right) \cdot N$
General k -Partition	$\Delta^2 \cdot \left(\frac{O(k^2)}{\epsilon}\right)^{2k+8}$	$\exp\left(\Delta \cdot \left(\frac{O(k^2)}{\epsilon}\right)^{k+1}\right)$	$O\left(\frac{\Delta_k}{\epsilon^2}\right) \cdot N$

FIG. 1. Summary of Results. Here Δ and Δ_k denote logarithmic factors. Specifically, $\Delta \stackrel{\text{def}}{=} \log(1/(\epsilon\delta))$ and $\Delta_k \stackrel{\text{def}}{=} \log(k/(\epsilon\delta))$. For simplicity we use Δ_k also in cases where the dependency is slightly better (e.g., in k -Colorability). We also note that in some of the cases and for certain values of δ , one can obtain slightly better complexities (in terms of the dependence on δ). This is done by considering the complexity for constant δ , and then amplifying the confidence by repeating the execution of the algorithm $\log(1/\delta)$ times.

For all our testers, the sample of vertices is viewed as consisting of two parts—one part, denoted U , has the role of inducing partitions on $V(G)$ (i.e., all graph vertices), and the other part, denoted S , evaluates these partitions. In the Bipartiteness and k -Colorability testers this view of the sample is only taken in the analysis. In the other algorithms we explicitly use these two parts of the sample differently. In particular, the latter algorithms consider all possible partitions of U . For each such partition of U , the second part of the sample, S , is partitioned according to a property-specific rule that uses the neighborhood relations between vertices in S and vertices in U .¹⁷

In all cases we want U to be *representative* of certain *good* partitions of $V(G)$, which exist whenever G has the property that is being tested. The notion of a good partition depends on the property being tested, and the sense in which U is representative varies accordingly. For example, when testing ρ -Clique, a good partition of a graph G is $(C, V(G) \setminus C)$, where C is a clique of size ρN . In this case U is representative if for *almost all* vertices v in the graph, if v neighbors all vertices in $U \cap C$, then it neighbors almost all vertices in C . Note that we do not require that the above hold for all vertices v , and this feature is common to all our definitions of a representative sample U . Consequently, U need be only of size $\text{poly}(\epsilon^{-1} \log(1/\delta))$, and so independent of N . This is especially important since most of our algorithms run in time exponential in $|U|$.

5.2.2. Specific Themes. Though the algorithms and their analyses have much in common, there are several themes that are only shared by some of the algorithms. Consider first the Bipartite and the k -Colorability testers. As noted above, for both properties we directly analyze the natural algorithms, which uniformly select a small sample of vertices and check whether the induced

¹⁷ In the ρ -Clique tester, we slightly deviate from the above formula by using an additional sample of vertices that aids in partitioning S .

subgraph is Bipartite (respectively, k -Colorable). These two algorithms have one-sided error as they always accept graphs that have the property. To show that, with high probability, the algorithms reject graphs that are far from having the property, we prove the counterpositive. Namely, that if a graph is accepted with probability greater than δ , then it is close to having the property. We do so by showing that the partition of the *sampled* vertices found by the algorithm when accepting the graph can be used to define a good partition of all graph vertices (where this definition is constructive.)

For all other properties, we first describe an algorithm that actually partitions $V(G)$. This graph-partitioning algorithm produces several (i.e., $\exp(\text{poly}(\log(1/\delta)/\epsilon))$) partitions of $V(G)$ and outputs the best one. The corresponding tester runs the graph-partitioning algorithm on a small sample S and evaluates the resulting partitions of S . Thus, S serves as a sample that approximates the quality of the graph-partitions (determined by the corresponding graph-partitioning algorithm). We will focus in this introductory discussion on the graph-partitioning algorithms.

In the case of ρ -Clique, the graph-partitioning algorithm tries to find an approximate ρ -clique (i.e., a set of vertices of size ρN that is close to being a clique). To this end, it selects a small sample of vertices U , and for each subset U' of U having size $\rho/2 \cdot |U|$, it considers all vertices that neighbor every vertex in U' . Assume first that G in fact has a clique C of size ρN . The idea is that in such a case, with high probability over the choice of U , the intersection of U and C will be of size at least $\rho/2 \cdot |U|$. Therefore, for some subset U' of U , the set of vertices, T , that neighbor every vertex in U' contains the clique C . However, T might contain many other vertices. Nonetheless, we can show that if we order the vertices in T roughly according to their degree in the subgraph induced by T , then the first ρN vertices in this order will be close to being a clique. The above implies that if G has a clique of size ρN then with high probability the algorithm finds an approximate clique. On the other hand, if every subset of size ρN in G is far from having a clique then no such approximate-clique exists.

The graph-partitioning algorithms for ρ -Cut, ρ -Bisection and the General Partition property work in $\ell = O(1/\epsilon)$ stages, and consider a fixed (standard) equal-partition $V^1, \dots, V^\ell$ of $V(G)$.¹⁸ In the i th stage, the set V^i is partitioned. The partition of V^i is determined by a partition of U^i (which is part of the sample U), and possibly other varying parameters. For example, in the case of ρ -Cut (where the partitioning algorithm actually finds a close-to-*maximum* cut) the vertices in V^i are partitioned as follows: The vertices that have more neighbors on side 1 of the partition of U^i are put on side 2 of the partition of V^i , and vice-versa. The fact that the algorithms partition only a set of size $O(\epsilon N)$ in each stage (using a different sample U^i) is essential to the proofs of correctness of the algorithms.

5.3. TWO TECHNICAL CONVENTIONS. As noted above, we represent N -vertex graphs by their $N \times N$ adjacency matrix. This representation is redundant and means that each edge appears twice. In some places, this convention simplifies

¹⁸ In particular, this partition is defined according to lexicographical order (i.e., where each vertex is represented by a string of length $\log_2 N$). We stress that this partition is arbitrary and has nothing to do with the desired partition determined by the property being tested.

the analysis, but in others it results in the need to artificially double certain natural quantities. Still we believe we made the right choice.

Our algorithms are presented as uniformly selecting sets of vertices of certain sizes (without repetitions). However, the analysis refers to uniform and independent selection of a number of vertices (with possible repetitions). Formally-inclined readers are thus encouraged to consider the algorithms as selecting multisets (rather than sets) of the prescribed size. Other readers may just ignore this point.

6. Testing Bipartiteness and Vertex-Colorability

In Section 6.1, we describe an algorithm for testing the class, $\mathcal{B}$, of bipartite graphs. This is a special case of testing k -Colorability, considered in the Section 6.2. We choose to present the case of $k = 2$ separately because it is both simpler to describe, and it served as a good prelude to the general case. Moreover, the algorithm presented for $k = 2$ has lower complexity (in terms of its dependence on the distance parameter, ϵ) than the one described in the Section 6.2.

6.1. TESTING BIPARTITENESS. We start by describing a testing algorithm whose query complexity is $O(\log^2(1/(\epsilon\delta))/\epsilon^4)$. We later point out how this algorithm can be slightly modified so that its query complexity decreases to $O(\log^2(1/(\epsilon\delta))/\epsilon^3)$.

Bipartite Testing Algorithm

- (1) Choose uniformly a set, denoted X , of $O(\log(1/(\epsilon\delta))/\epsilon^2)$ vertices.
- (2) For every pair of vertices $v_1, v_2 \in X$, query if $(v_1, v_2) \in E(G)$. Let G_X be the induced subgraph.
- (3) If G_X is a bipartite graph, then output *accept*; otherwise, output *reject*.

Before stating the main theorem of this section, we introduce the following definitions. Recall that N is the number of vertices of G .

Definition 6.1.1. (Violating Edges and Good Partitions). We say that an edge $(u, v) \in E(G)$ is a *violating edge* with respect to a partition (V_1, V_2) of $V(G)$ if either $u, v \in V_1$ or $u, v \in V_2$. If a partition (V_1, V_2) has at most ϵN^2 violating edges then we say that it is ϵ -good. Otherwise, it is ϵ -bad. A partition that has no violating edges is called *perfect*.

Thus, if G is bipartite, then there exists a perfect partition of $V(G)$, and if G is ϵ -far from bipartite then *every* partition of $V(G)$ is ϵ -bad.

THEOREM 6.1.2. *The Bipartite Testing Algorithm is a property testing algorithm for the class of bipartite graphs whose query complexity and running time are $O(\log^2(1/(\epsilon\delta))/\epsilon^4)$. Furthermore, if the tested graph G is bipartite, then it is accepted with probability 1, and, with probability at least $1 - \delta$ (over the choice of the sampled vertices), it is possible to construct an ϵ -good partition of $V(G)$ in (additional) time $O(\log(1/(\epsilon\delta))/\epsilon \cdot N)$,*

PROOF. It is clear that if G is bipartite then any subgraph of G is bipartite and hence G will always be accepted. Since it is possible to determine if G_X is bipartite by simply performing a breadth-first-search (BFS) on G_X , the bound on the running time of the testing algorithm directly follows. Note that if G_X is

bipartite, then the BFS provides us with a perfect partition of X , while if it is not bipartite, then it gives a certificate that G is not bipartate. This certificate is in form of a cycle of odd length in G_X (which is also a cycle in G). Thus, the heart of this proof is to show that if G is ϵ -far from bipartite then the test will reject it with probability at least $1 - \delta$. To this end, we prove the counterpositive of the previous statement: For any graph G , if the *Bipartite Testing Algorithm* accepts G with probability greater than δ , then $V(G)$ must have an ϵ -good partition.

We view the set of sampled vertices X as a union of two disjoint sets U and S , where

$$t \stackrel{\text{def}}{=} |U| = O\left(\frac{\log(1/(\delta\epsilon))}{\epsilon}\right), \text{ and } m \stackrel{\text{def}}{=} |S| = O\left(\frac{t + \log(1/\delta)}{\epsilon}\right).$$

The role of U , or more precisely of a given partition (U_1, U_2) of U , is to define a partition of *all* of $V(G)$. In particular, if the test accepts G , then we know that X has a perfect partition, and we will be interested in the partition of U induced by this perfect partition of X . The role of S is to *test* the partitions of $V(G)$ defined by the partition of U in the following sense. If a certain partition (V_1, V_2) of $V(G)$, defined by a partition (U_1, U_2) of U , is ϵ -bad, then with very high probability there is no partition (S_1, S_2) of S such that $(U_1 \cup S_1, U_2 \cup S_2)$ is a perfect partition of X . We next make the above notions more precise.

Given a partition (U_1, U_2) of U we define the following partition (V_1, V_2) of $V(G)$: let $V_1 \stackrel{\text{def}}{=} \Gamma(U_2)$, and $V_2 \stackrel{\text{def}}{=} V(G) \setminus \Gamma(U_2)$. That is, V_1 is the set of neighbors of U_2 , and V_2 contains all neighbors of U_1 (which are not neighbors of U_2), as well as the rest of the vertices—namely those which do not have a neighbor in U . Note that the partition of U is not relevant to the placement of vertices that have no neighbor in U . Thus, we first ensure that most vertices in $V(G)$ (or actually most “influential” vertices in $V(G)$) have a neighbor in U .

Definition 6.1.3 (Influential Vertices and Covering Sets). We say that a vertex $v \in V(G)$ is *influential* if it has degree at least $(\epsilon/3)N$. Recall that a degree of a vertex is the sum of its in-degree and its out-degree (which is twice its degree in the undirected representation of the graph). We call U a *covering* set for $V(G)$ if all but at most $(\epsilon/6)N$ of the influential vertices in $V(G)$ have a neighbor in U (here each neighbor from which there is one outgoing edge, and one ingoing edge, is counted once).

Note that in the above definition, we did not require of U that *every* influential vertex have a neighbor in U , but rather that this hold for almost all influential vertices. This slackness (which appears in various forms in our other algorithms as well) is what allows our algorithm to use a sample size that is independent of the size of the graph.

CLAIM 6.1.4. *With probability at least $1 - \delta/2$, a uniformly chosen set U of size $t = O(\log(1/(\delta\epsilon))/\epsilon)$ is a covering set for $V(G)$.*

PROOF. Let $t = (6 \cdot \log(12/(\delta\epsilon))/\epsilon)$. For a given influential vertex v , the probability that v does not have any neighbor in a uniformly chosen set U of size

t is at most

$$\left(1 - \frac{\epsilon}{6}\right)^t \leq \exp\left(-t \cdot \frac{\epsilon}{6}\right) = \frac{\epsilon \cdot \delta}{12}. \quad (2)$$

Hence, the expected number of influential vertices that do not have a neighbor in a random set U is $(\delta \cdot \epsilon)/12 \cdot N$, and by Markov's inequality (see Appendix B), the probability that there are more than $(\epsilon/6)N$ such vertices is less than $\delta/2$. $\square$

Given a covering set U and a partition of U , we can concentrate on the violating edges between vertices in $\Gamma(U_i)$, for $i = 1, 2$. The reason being that the total number of edges incident to vertices not in $\Gamma(U)$ is small. This motivates the following definition, where a useful partition of U implies that G has a good partition (see Lemma 6.1.8).

Definition 6.1.5 (useful partitions). Let $U \subset V(G)$. A partition (U_1, U_2) of U is called ϵ -useful (or just useful) if

$$|\{(v, v') \in E(G) : \exists i \in \{1, 2\} \text{ such that } v, v' \in \Gamma(U_i)\}| < \frac{\epsilon}{3} \cdot N^2. \quad (3)$$

Otherwise, it is ϵ -unuseful.

In other words, a partition of U is unuseful if there are too many violating edges among the neighbors either of U_1 or of U_2 in the corresponding partition defined on $V(G)$. As the following claim shows, if (U_1, U_2) is an unuseful partition, then with high probability we shall see evidence of its unusefulness in the sample S . The evidence is in form of an edge $(v, v') \in S \times S$ between neighbors of vertices in, say, U_1 . Let $u \in U_1$ (respectively, $u' \in U_1$) be a neighbor of v (respectively, v'). In case $u = u'$, there is a triangle in G_X (which means that the test would reject). In case $u \neq u'$, due to the edge (v, v') , in any perfect partition of $X = U \cup S$, u and u' must belong to different sides of the partition (and so (U_1, U_2) cannot be used in the partition of $X = U \cup S$).

CLAIM 6.1.6. Let U be a set of size t and let (U_1, U_2) be a fixed ϵ -unuseful partition of U . Then, for a uniformly chosen set $S \subset V(G)$ of cardinality, $m = O(t + \log(1/\delta)/\epsilon)$,

$$\Pr_S[\forall v, v' \in S, i \in \{1, 2\} : (v, v') \notin E(\Gamma(U_i), \Gamma(U_i))] < \delta \cdot 2^{-(t+1)}.$$

The claim says that, if (U_1, U_2) is an ϵ -unuseful partition of U , then with very high probability there exists no partition of S so that the combined partition of $X = U \cup S$ (respecting (U_1, U_2)) is a perfect partition of the subgraph induced by X .

PROOF. If (U_1, U_2) is an ϵ -unuseful partition, then by Eq. (3):

$$\Pr_{v,v'}[\exists i \in \{1, 2\} \text{ such that } (v, v') \in E(\Gamma(U_i), \Gamma(U_i))] \geq \frac{\epsilon}{3}. \quad (4)$$

Since S can be chosen by drawing $m/2$ independent random pairs of vertices (v, v') ,

$$\begin{aligned}
 & \Pr_S[\forall v, v' \in S, i \in \{1, 2\} : (v, v') \notin E(\Gamma(U_i), \Gamma(U_i))] \\
 & \leq \left(1 - \frac{\epsilon}{3}\right)^{m/2} \\
 & = \exp(-\Omega(t + \log(1/\delta))) \\
 & < \delta \cdot 2^{-(t+1)} \quad \square
 \end{aligned}$$

Since there are 2^t possible partitions of U , using Claim 6.1.6, we have

COROLLARY 6.1.7. *For every set U of size t , if all partitions of U are ϵ -unuseful, then with probability at least $1 - \delta/2$ there is no perfect partition of $X = U \cup S$, where S is chosen uniformly. (In such a case G_X is found to be nonbipartite, and the test rejects G .)*

On the other hand, if U has a useful partition, denoted (U_1, U_2) , and U is a covering set for $V(G)$, then this partition induces a good partition of the entire graph. More precisely,

LEMMA 6.1.8. *For every graph G , if there exists a covering set U of $V(G)$, which has an ϵ -useful partition (U_1, U_2) , then G is ϵ -close to bipartite. In particular, the following partition (V_1, V_2) of $V(G)$ is ϵ -good: $V_1 \stackrel{\text{def}}{=} \Gamma(U_2)$, $V_2 \stackrel{\text{def}}{=} V(G) \setminus \Gamma(U_2)$.*

PROOF. Let us count the number of violating edges with respect to the partition (V_1, V_2) :

- Edges incident to noninfluential vertices: There are at most N such vertices and by definition each has at most $(\epsilon/3)N$ incident edges, giving a total of $(\epsilon/3)N^2$.
- Edges incident to influential vertices which do not have neighbors in U : Since U is a covering set, there are at most $(\epsilon/6)N$ such vertices and each has at most $2N$ incident edges, totaling to $(\epsilon/3)N^2$.
- Violating edges which are incident to neighbors of U . We consider two cases:
 - Edges of the form $(v, v') \in E(V_1, V_1)$. Since $V_1 = \Gamma(U_2)$, these edge have both end-points in $\Gamma(U_2)$.
 - Edges of the form $(v, v') \in E(V_2, V_2)$. By definition of V_2 , both v and v' are not in $\Gamma(U_2)$. However, since $v, v' \in \Gamma(U)$, it follows that these edges have both end-points in $\Gamma(U_1)$.

By the ϵ -usefulness of (U_1, U_2) , there are at most $(\epsilon/3)N^2$ such vertices.

Thus, we have a total of at most ϵN^2 violating edges, as required. $\square$

Combining Lemma 6.1.8 with Claim 6.1.4 and Corollary 6.1.7, we complete the proof of Theorem 6.1.2, as follows: If G is accepted with probability greater than δ , then, by Claim 6.1.4, the probability that G is accepted *and* U is a covering set is greater than $\delta/2$. Thus, there exists a covering set $U \in V(G)$, such that if U is chosen, then G is accepted with probability greater than $\delta/2$ (where here the probability is taken only over the choice of S). But in such a case it follows from

Corollary 6.1.7 that (the covering set) U must have a useful partition, and we can apply Lemma 6.1.8 to show that G must be ϵ -close to bipartite.

Finally, let G be a bipartite graph, and let (V_1, V_2) be a perfect partition of $V(G)$. Then, for every covering set U of $V(G)$ (where such a set is chosen with probability at least $1 - \delta/2$), there exists a partition (U_1, U_2) of U , such that $U_1 \subseteq V_1$ and $U_2 \subseteq V_2$. Thus, necessarily, (U_1, U_2) is a useful partition of U which by Lemma 6.1.8 defines an ϵ -good partition of $V(G)$. Furthermore, given such a partition of a covering set U , for *every* set S there exists a perfect partition of $U \cup S$, of the form $(U_1 \cup S_1, U_2 \cup S_2)$. On the other hand, by Claim 6.1.6, for any set U , with probability at least $1 - \delta/2$ over the choice of S , there will be no perfect partition of $U \cup S$ that induces an unuseful partition of U . Therefore, with probability at least $1 - \delta$ over the choice of U and S , the testing algorithm (using BFS) will find a perfect partition of $U \cup S$ that induces a useful partition of U , which can then be used to construct an ϵ -good partition of $V(G)$ (as defined in Lemma 6.1.8) in time $O(|U| \cdot N) = O(\epsilon^{-1} \log(1/(\epsilon\delta))) \cdot N$.

This completes the proof of Theorem 6.1.2. $\square$

6.1.1. Improving the Query Complexity. We can save a factor of $1/\epsilon$ in the query complexity and in the running time of the testing algorithm. This is done simply by observing that we do not need to perform queries for all pairs of vertices in S . Instead, we can choose S to be a uniformly distributed random sample of $m/2$ pairs of vertices. We then need only to query which of these $m/2 = O(\log(1/(\epsilon\delta))/\epsilon^2)$ pairs are edges, as well as query all $m \cdot t = O(\log^2(1/(\epsilon\delta))/\epsilon^3)$ pairs (u, v) where $u \in U$ and $v \in S$. Note that the proof of Theorem 6.1.2 does not refer to any edges between vertices in S , except for the $m/2$ pairs mentioned above (which are used for establishing Claim 6.1.6).

6.1.2. Impossibility of Testing Without Queries. A natural question that may arise is whether queries are really necessary for testing bipartiteness, or perhaps it might be possible to test this property from a random labeled sample (of pairs of vertices) alone. We show that queries are in fact necessary in the sense that any testing algorithm that uses only a random sample (of pairs of vertices) must have very large sample complexity. More precisely:

PROPOSITION 6.1.2.1. *Any property testing algorithm for the class of bipartite graphs that observes only a random labeled sample, must have sample complexity $\Omega(\sqrt{N})$.*

PROOF. Consider the following two classes of graphs: $\mathcal{G}^1$ is the class of all complete bipartite graphs G in which both sides are of equal cardinality. That is, $V(G) = V_1 \cup V_2$, $|V_1| = |V_2| = N/2$, and $E(G) = \{(v_1, v_2) : v_1 \in V_i, v_2 \in V_j, i \neq j\}$. $\mathcal{G}^2$ is the class of graphs that consists of two disjoint cliques of size $N/2$. That is, $V(G) = V_1 \cup V_2$, $|V_1| = |V_2| = N/2$, and $E(G) = \{(v, v') : v, v' \in V_1 \text{ or } v, v' \in V_2\}$. Clearly, all graphs in $\mathcal{G}^2$ are $1/4$ -far from bipartite. Note that all graphs in both classes have the same edge density, since every vertex has degree N . Intuitively, we want to show that if the edge-labeled sample is not large enough then a hypothetical property testing algorithm cannot distinguish between random samples labeled by graphs in $\mathcal{G}^1$ and random sample labeled by graphs in $\mathcal{G}^2$.

For simplicity, let us fix δ to be $1/4$. Then, by definition, a property testing algorithm for the class of bipartite graphs should accept each $G \in \mathcal{G}^1$ with

probability at least $3/4$, and should accept each $G \in \mathcal{G}^2$ with probability less than $1/4$. Therefore, the difference in acceptance probability between an arbitrary $G \in \mathcal{G}^1$ and an arbitrary $G \in \mathcal{G}^2$ must be greater than $1/2$. Since the above should be true for any pair of graphs taken from the two classes, it should hold for a random pair of graphs chosen from the two classes. Suppose we first draw an unlabeled random sample of m pairs of vertices, and then label it by a graph G chosen randomly either from the class $\mathcal{G}^1$ or from the class $\mathcal{G}^2$. Assume first that the sample is such that no vertex appears in more than one pair in the sample. Then, regardless of whether G was chosen uniformly in $\mathcal{G}^1$ or in $\mathcal{G}^2$, each of the 2^m possible labeling of the sample has equal probability. If the sample does include two pairs that share a vertex, then we cannot make such a claim. Let us say in this case that the sample is *informative*. However, the probability that a random sample of size m is informative is at most $\binom{m}{2} \cdot 2/N < m^2/N$. By the argument made above on noninformative samples, the difference between the acceptance probability of a random graph in $\mathcal{G}^1$ and the acceptance probability of a random graph in $\mathcal{G}^2$ is at most the probability that a random sample is informative. But in order that this probability be greater than $1/2$, the size of the random sample must be greater than $\sqrt{N/2}$. $\square$

6.2. TESTING k -COLORABILITY ($k > 2$). In this section, we present an algorithm for testing the k -Colorability property for any given k . Namely, we are interested in determining if the vertices of a graph G can be colored by k colors so that no two adjacent vertices are colored by the same color, or if any k -partition of the graph has at least ϵN^2 violating edges (i.e., edges between pairs of vertices that belong to the same side of the partition).

The test itself is analogous to the bipartite test described in the previous section: We sample from the vertices of the graph, query all pairs of vertices in the sample to find which are edges in G , and check if the induced subgraph is k -Colorable. The query complexity of the algorithm is polynomial in $1/\epsilon$, $\log(1/\delta)$ and k . In lack of efficient algorithms for k -Colorability, for $k \geq 3$, we use the obvious exponential-time algorithm on the induced subgraph (which is typically small). Note that the number of queries made is larger than in the (improved) Bipartite Tester (i.e., by a factor of $\tilde{O}(k^4/\epsilon^3)$).

k -Colorability Testing Algorithm

- (1) Choose uniformly set of $O(k^2(\log(k/\delta)/\epsilon^3))$ vertices, denoted X .
- (2) For every pair of vertices $v_1, v_2 \in X$, query if $(v_1, v_2) \in E(G)$. Let G_X be the induced subgraph.
- (3) If G_X is k -Colorable, then output *accept*; otherwise, output *reject*.

Similarly to the bipartite case, we define violating edges and good k -partitions.¹⁹

Definition 6.2.1 (Violating Edges and Good k -Partitions). We say that an edge $(u, v) \in E(G)$ is a *violating edge* with respect to a k -partition $\pi: V(G) \rightarrow [k]$ if $\pi(u) = \pi(v)$. We shall say that a k -partition is ϵ -good if it has at most ϵN^2 violating edges (otherwise, it is ϵ -bad). The partition is *perfect* if it has no violating edges.

¹⁹ k -partitions are associated with mappings of the vertex set into the canonical k -element set $[k]$. The partition associated with $\pi: V(G) \rightarrow [k]$ is $(V_1 \stackrel{\text{def}}{=} \pi^{-1}(1), \dots, V_k \stackrel{\text{def}}{=} \pi^{-1}(k))$. We shall use the mapping notation π , and the explicit partition notation $(V_1, \dots, V_k)$, interchangeably.

THEOREM 6.2.2. *The k -Colorability Testing Algorithm is a property testing algorithm for the class of k -Colorable graphs whose query complexity and running time are*

$$O\left(\frac{k^4 \cdot \log^2(k/\delta)}{\epsilon^6}\right) \quad \text{and} \quad \exp\left(O\left(\frac{k^2 \cdot \log^2(k/\delta)}{\epsilon^3}\right)\right),$$

respectively. If the tested graph G is k -Colorable, then it is accepted with probability 1, and with probability at least $1 - \delta$, (over the choice of the sampled vertices), it is possible to construct an ϵ -good k -partition of $V(G)$ in (additional) time $O(\log(k/\delta)/\epsilon^2) \cdot N$.

PROOF. If G is k -Colorable then every subgraph of G is k -Colorable, and hence G will always be accepted. As in the bipartite case, the crux of the proof is to show that every G which is ϵ -far from the class of k -Colorable graphs, denoted $\mathcal{G}_k$, is rejected with probability at least $1 - \delta$. Again, we establish this claim by proving its counterpositive, namely, that every G that is accepted with probability greater than δ must be ϵ -close to $\mathcal{G}_k$. This is done by giving a (constructive) proof of the existence of an ϵ -good k -partition of $V(G)$. Hence, in case $G \in \mathcal{G}_k$, we also get an efficient probabilistic procedure for finding an ϵ -good k -partition of $V(G)$. Note that if the test rejects G , then we have a certificate that $G \notin \mathcal{G}_k$, in form of the (small) subgraph induced by X which is not k -Colorable.

We view the set of sampled vertices X as a union of two disjoint sets U and S , where U is a union of ℓ (disjoint) sets $U^1, \dots, U^\ell$, each of size m . The size of S is m as well, where $m = O(\ell \cdot \epsilon^{-1} \log(k/\delta))$ and $\ell = \lceil 4k/\epsilon \rceil$. The roles of U and S are analogous to their roles in the bipartite case. The set U (or rather a k -partition of U) is used to define a k -partition of $V(G)$. The set S ensures that with high probability, the k -partition of U which is induced by the perfect k -partition of $X = U \cup S$, defines an ϵ -good partition of $V(G)$.

In order to define a k -partition of $V(G)$ given a k -partition of U , we first introduce the notion of a *clustering* of the vertices in $V(G)$ with respect to this partition of U . More precisely, we define the clustering based on the k -partition of a subset $U' \subset U$ (which is determined later), where this partition, denoted $(U'_1, \dots, U'_k)$, is the one induced by the k -partition of U . The clustering is defined so that vertices in the same cluster have neighbors in the same components of the partition of U' . For every $A \subseteq [k]$, the A -cluster, denoted C_A , contains all vertices in $V(G)$ that have neighbors in U'_i for every $i \in A$ (and do not have neighbors in the other U'_i 's). The clusters impose restrictions on possible extensions of the partition of U' to partitions $(V_1, \dots, V_k)$ of $V(G)$, which do not have violating edges incident to vertices in U' . Namely, vertices in C_A should NOT be placed in any V_i such that $i \in A$. As a special case, $C_\emptyset$ is the set of vertices that do not have any neighbors in U' (and hence can be put in any component of the partition). In the other extreme, $C_{[k]}$ is the set of vertices that in any extension of the partition of U' will cause violations. For each i , the vertices in $C_{[k] \setminus \{i\}}$ are *forced* to be put in V_i , and thus are easy to handle. In the bipartite case we focused on the clusters $C_{\{1\}}$ and $C_{\{2\}}$, where vertices in $C_{\{i\}}$ were forced to the side opposite to i . (The cluster $C_\emptyset$ was explicitly shown to be unimportant and the cluster $C_{\{1,2\}}$ was dealt with implicitly.) In the case of

k -Coloring the situation is more complex. In particular, the clusters C_A where $|A| < k - 1$ do not force a placement of vertices.

Definition 6.2.3 (Clusters). Let U' be a set of vertices, and let π' be a perfect k -partition of U' . Define $U'_i \stackrel{\text{def}}{=} \{v \in U' : \pi'(v) = i\}$. For each subset $A \subseteq [k]$ we define the A -cluster with respect to π' as follows:

$$C_A \stackrel{\text{def}}{=} \left(\bigcap_{i \in A} \Gamma(U'_i) \right) - \left(\bigcup_{i \notin A} \Gamma(U'_i) \right). \quad (5)$$

The relevance of the above clusters becomes clear given the following definitions of extending and consistent partitions.

Definition 6.2.4 (Consistent Extensions). Let U' and π' be as above. We say that a k -partition π of $V(G)$ extends a k -partition π' of U' if $\pi(u) = \pi'(u)$ for every $u \in U'$. An extended partition π is consistent with π' if $\pi(v) \neq \pi'(u)$ for every $u \in U'$ and $v \in \Gamma(u) \setminus C_{[k]}$, where $C_{[k]}$ is the $[k]$ -cluster with respect to π' .

Thus, each vertex v in the cluster C_A (with respect to π' defined on U') is forced to satisfy $\pi(v) \in \bar{A} \stackrel{\text{def}}{=} [k] \setminus A$, for every k -partition π which extends π' in a consistent manner. There are no restrictions regarding vertices in $C_\emptyset$ and vertices in $C_{[k]}$ (the latter is guaranteed artificially in the definition and the consequences will have to be treated separately). For $v \in C_{[k] - \{i\}}$, the consistency condition forces $\pi(v) = i$, but unlike the bipartite case we cannot ignore the A -clusters with $|A| < k - 1$.

We now focus on the main problem of the analysis. Given a k -partition of U , what is a good way to define a k -partition of $V(G)$? Our main idea is to claim that with high probability the set U contains a subset U' so that the clusters with respect to the induced k -partition of U' determine whatever needs to be determined. That is, if these clusters allow to place some vertex in a certain component of the partition, then doing so does not introduce too many violating edges. The first step in implementing this idea is the notion of a *restricting* vertex. A vertex $v \in C_A$ is *restricting* if for every $i \in \bar{A}$, adding v to U'_i (and thus to U') will cause many of its neighbors to move to a cluster corresponding to a bigger subset. That is, v 's neighbors in the B -cluster (with respect to $(U'_1, \dots, U'_k)$) move to the $(B \cup \{i\})$ -cluster (with respect to $(U'_1, \dots, U'_i \cup \{v\}, \dots, U'_k)$). More precisely:

Definition 6.2.5 (Restricting Vertex). A pair (v, i) , where $v \in C_A$, $A \neq [k]$ and $i \in \bar{A}$ is said to be *restricting* with respect to a k -partition π' of U' if v has at least $(\epsilon/4)N$ neighbors in $\bigcup_{B: i \notin B} C_B$. Otherwise, (v, i) is *nonrestricting*. A vertex $v \in C_A$, $A \neq [k]$, is *restricting* with respect to π' if for every $i \in \bar{A}$, the pair (v, i) is restricting. Otherwise, v is *nonrestricting*. As always, the clusters are with respect to π' .

Given a perfect k -partition of U , we construct $U' \subset U$ in steps starting with the empty set. At step j , we add to U' a vertex $u \in U^j$ (recall that $U = U^1 \cup \dots \cup U^\ell$), which is a restricting vertex with respect to the k -partition of the current set U' . If no such vertex exists, the procedure terminates. When the procedure terminates (and as we shall see it must terminate after at most ℓ steps), we will be able to define, based on the k -partition of the final U' , an

ϵ -good k -partition of $V(G)$. The procedure defined below is viewed at this point as a mental experiment. Namely, it is provided in order to show that with high probability there exists a subset U' of U with certain desired properties (which we later exploit). We later discuss how to implement this procedure when we are actually interested in choosing U' for the purposes of partitioning all of $V(G)$ efficiently.

Restriction Procedure (Construction of U')

Input: a perfect k -partition of $U = U^1, \dots, U^\ell$.

- (1) $U' \leftarrow \emptyset$.
- (2) For $j = 1, 2, \dots$ do the following. Consider the current set U' and its partition π' (induced by the perfect k -partition of U).
 - If there are less than $\epsilon/8 \cdot N$ restricting vertices with respect to π' , then halt and output U' .
 - If there are at least $\epsilon/8 \cdot N$ restricting vertices but there is no restricting vertex in U^j , then halt and output an error symbol.
 - Otherwise (there is a restricting vertex in U^j), add the first (by any fixed order) restricting vertex to U' (and proceed to the next iteration).

CLAIM 6.2.6. *For every U and every perfect k -partition of U , after at most $\ell = \lceil 4k/\epsilon \rceil$ iterations, the Restriction Procedure halts and outputs either U' or error.*

PROOF. If the procedure has not halted, then in each iteration a restricting vertex is added to U' causing at least $\epsilon/4 \cdot N$ vertices in $V(G)$ to move to a cluster corresponding to a bigger subset. Since each vertex can be moved at most k times (before it belongs to $C_{[k]}$), the maximal number of iterations before the procedure halts is $4k/\epsilon$. $\square$

Before we show how U' can be used to define a k -partition π of $V(G)$, we need to ensure that with high probability, the restriction procedure in fact outputs a set U' and not an error symbol. To this end, we first extend the notion of *covering set* to the context of k -Coloring. Though the notion here may seem somewhat remote from the one used in the Bipartite case, it can be shown that the two are related.

Definition 6.2.7 (Covering Sets—for k -Coloring). We say that U is a *covering set* for $V(G)$, if for every perfect k -partition of U , the Restriction Procedure, given this partition as input, halts with an output $U' \subset U$ (rather than an error symbol).

In other words, U is such that for every perfect k -partition of U and for each of the at most ℓ iterations of the procedure, if there exist at least $(\epsilon/8) \cdot N$ restricting vertices with respect to the current partition of U' , then U^j will include at least one such restricting vertex.

LEMMA 6.2.8. *With probability at least $1 - \delta/2$, a uniformly chosen set U of size $\ell \cdot m = O(k^2 \log(k/\delta)/\epsilon^3)$ is a covering set.*

PROOF. Let us first consider a single iteration of the Restriction Procedure. If there are at least $(\epsilon/8) \cdot N$ restricting vertices with respect to the partition π' of the current U' , then the probability that in a uniformly chosen sample of size m ($= |U^j|$) there will be no restricting vertex with respect to π' , is at most $(1 - (\epsilon/8))^m$. By our choice of $m = O(\ell \cdot \epsilon^{-1} \log(k/\delta))$, the latter is bounded by

$(\delta/2)k^{-\ell}$. Thus, the lemma reduces to proving that, for every j , the number of possible pairs (U', π') that we need to consider for the j th iteration is at most k^{j-1} .

We shall prove the above claim inductively. Let the set U' at the start of iteration j (before adding a new restricting vertex to it) be denoted by $U'(j-1)$, and let its partition be denoted by π'_{j-1} . For the base case, $j=1$, the set $U'(0)$ is empty and the claim trivially holds. Assuming the claim holds for $j-1$, we now prove it for j . In the j th iteration, for each of the possible pairs $(U'(j-1), \pi'_{j-1})$, such that there exist at least $(\epsilon/8) \cdot N$ restricting vertices with respect to π'_{j-1} , the vertex $u_j \in U^j$ which is the first restricting vertex in U^j , is uniquely defined.²⁰ Hence, for each such pair $(U'(j-1), \pi'_{j-1})$, there is a single possible extension $U'(j)$ of $U'(j-1)$, namely, $U'(j) = U'(j-1) \cup \{u_j\}$. The new partition, π'_j , which extends π'_{j-1} can be one of at most k possibilities (depending only on $\pi'_j(u_j)$). $\square$

Definition 6.2.9 (Closed Partitions). Let U' be a set and π' a k -partition of it. We call (U', π') *closed* if there are less than $\epsilon/8 \cdot N$ restricting vertices with respect to π' .

Clearly, if the Restriction Procedure outputs a set U' , then this set together with its (induced) partition are closed. If (U', π') is closed, then most of the vertices in $V(G)$ are nonrestricting. Recall that a nonrestricting vertex v , belonging to a cluster C_A , $A \neq [k]$, has the following property. There exists at least one index $i \in \bar{A}$, such that (v, i) is nonrestricting. It follows from Definition 6.2.5 that for every consistent extension of π' to π which satisfies $\pi(v) = i$ there are at most $(\epsilon/2)N$ violating edges incident to v .²¹ However, even if v is nonrestricting, there might be indices $i \in \bar{A}$ such that (v, i) is restricting, and hence there may exist a consistent extension of π' to π that satisfies $\pi(v) = i$ in which there are more than $(\epsilon/2)N$ violating edges incident to v . Therefore, we need to define for each vertex its set of *forbidden* indices that will not allow to have $\pi(v) = i$ for a restricting pair (v, i) .

Definition 6.2.10 (Forbidden Sets). Let (U', π') be closed and consider the clusters with respect to π' . For each $v \in V(G) \setminus U'$ we define the *forbidden set* of v , denoted F_v , as the smallest set satisfying

— $F_v \supseteq A$, where $v \in C_A$.

— For every $i \in \bar{A}$, if v has at least $\epsilon/4 \cdot N$ neighbors in the clusters C_B for which $i \notin B$, then i is in F_v . That is, if (v, i) is restricting, then $i \in F_v$.

For $u \in U'$, we artificially define $F_u = [k] \setminus \{\pi'(u)\}$.

LEMMA 6.2.11. *Let (U', π') be an arbitrary closed pair and F_v 's be as in Definition 6.2.10. Then:*

$$(1) |\{v : (v \notin C_{[k]}) \wedge (F_v = [k])\}| \leq (\epsilon/8)N.$$

²⁰ It may be the case that no such restricting vertex exists in U^j , but the probability for this event has been bounded above.

²¹ First note that by definition of a consistent extension, no vertex in cluster C_B , where $i \in B$, can have π -value i . Thus, all violated edges incident to v are incident to vertices in clusters C_B so that $i \notin B$. Using the definition of a restricting pair (v, i) , we are done.

- (2) Let π be any k -partition of $V(G) \setminus \{v : F_v = [k]\}$ such that $\pi(v) \notin F_v$ for every v . Then, the number of edges $(v, v') \in E(G)$ for which $\pi(v) = \pi(v')$ is at most $\epsilon/2 \cdot N^2$.

The lemma can be thought of as saying that any k -partition which respects the forbidden sets is good (i.e., does not have many violating edges). However, the partition applies only to vertices for which the forbidden set is not $[k]$. The first item tells us that there cannot be many such vertices which do not belong to the cluster $C_{[k]}$. We deal with vertices in $C_{[k]}$ at a later stage.

PROOF. The first item follows from the closeness of (U', π') . Namely, if $F_v = [k]$ and $v \notin C_{[k]}$, then, by the second item of Definition 6.2.10, it follows that for every $i \in \bar{A}$, vertex v has at least $(\epsilon/4) \cdot N$ neighbors in clusters C_B such that $i \notin B$. But in this case, it is a restricting vertex with respect to π' . By Definition 6.2.9, as applied to (U', π') , there are at most $(\epsilon/8) \cdot N$ such vertices.

For the second item, consider a vertex v such that $\pi(v) = i$ and so $v \in C_A$ where $i \notin F_v \supseteq A$. All edges (v, u) and (u, v) such that $u \in C_B$ and $i \in B$ cannot be violating edges since $i \in F_u$ (by the first item in Definition 6.2.10). As for edges (v, u) and (u, v) where $u \in C_B$ and $i \notin B$, and vertex v can have at most $(\epsilon/2) \cdot N$ such edges (according to the second item in Definition 6.2.10). The total of violating edges is hence at most $(\epsilon/2) \cdot N^2$. $\square$

We next show that with high probability over the choice of S , the k -partition π' of U' (induced by the k -partition of $U \cup S$) is such that $C_{[k]}$ is small. This implies that the vertices in $C_{[k]}$ (which were left out of the partition in the previous lemma) can be placed in any component of the partition without contributing too many violating edges (which are incident to them).

Definition 6.2.12 (Useful k -partitions). We say that a pair (U', π') is ϵ -useful if $|C_{[k]}| < (\epsilon/8)N$. Otherwise it is ϵ -unuseful.

The next claim follows from our choice of m and the above definition.

CLAIM 6.2.13. Let U' be a fixed set of size ℓ and π' be a fixed k -partition of U' so that (U', π') is ϵ -unuseful. Let S be a uniformly chosen set of size m . Then, with probability at least $1 - (\delta/2)k^{-\ell}$, there exists no perfect k -partition of $U' \cup S$ which extends π' .

PROOF. By definition, for $C_{[k]}$ defined based on (U', π') , we have $|C_{[k]}| \geq (\epsilon/8) \cdot N$, and so a uniformly chosen S contains a vertex in $C_{[k]}$ with the claimed probability, in which case any extension of π' to S is nonperfect. $\square$

By the same argument applied in the proof of Lemma 6.2.8, we have that the number of possible closed pairs (U', π') determined by all possible k -partitions of U is at most k^ℓ . Therefore, we get the following corollary to the above claim:

COROLLARY 6.2.14. If all closed pairs (U', π') that are determined by all possible k -partitions of U are unuseful, then with probability at least $1 - \delta/2$ over the choice of S , there is no perfect k -partition of $X = U \cup S$.

We can now wrap up the main part of the proof of Theorem 6.2.2. If G is accepted with probability greater than δ , then by Lemma 6.2.8, the probability that it is accepted and U is a covering set is greater than $\delta/2$. In particular, there must exist at least one covering set U , such that if U is chosen then G is accepted

with probability greater than $\delta/2$ (with respect to the choice of S). That is, (with probability greater than $\delta/2$) there exists a perfect partition of $U \cup S$. But in such a case (by applying Corollary 6.2.14), there must be a useful closed pair (U', π') (where $U' \subset U$). If we now partition $V(G)$ as described in Lemma 6.2.11, where vertices with forbidden set $[k]$ are placed arbitrarily, then from the two items of Lemma 6.2.11 and the usefulness of (U', π') it follows that there are at most ϵN^2 violating edges with respect to this partition. This completes the main part of the proof and the rest refers to the efficient procedure for finding ϵ -good partitions.

Similar to the bipartite case, if $G \in \mathcal{G}_k$, then with probability at least $1 - \delta$ (over the choice of U and S), the k -Coloring of G_X (recall that $X = U \cup S$) is such that the induced (perfect) coloring of U determines a useful pair (U', π') which can be used to determine a partition of $V(G)$. Details follow.

EFFICIENT CONSTRUCTION OF AN ϵ -GOOD k -PARTITION OF $V(G)$. For the efficient implementation, we assume the testing algorithm is run with distance parameter $\epsilon/2$ and confidence parameter $\delta/2$. The main point we need to address is the question of efficiently implementing the Restricting Procedure (i.e., constructing U' given U and a perfect k -partition of U), and the definition of forbidding sets. We first observe, that in the Restricting Procedure we do not actually need to determine (in each iteration) if there are more or less than $(\epsilon/8) \cdot N$ restricting vertices. Since we know that with high probability U^j contains a restricting vertex if many such vertices exist, we need only scan U^j in search for such a vertex. Note that no harm is done when despite the fact that there are too few restricting vertices in iteration j nevertheless U^j contains one. This is true since the bound on the number of iterations performed by the Restriction Algorithm, is unrelated to the actual number of restricting vertices in each iteration.

In order to recognize restricting vertices, we do the following. We uniformly select ℓ sets of vertices, $W^1, \dots, W^\ell$, each of size $\Theta(k \log(k/(\epsilon\delta))/\epsilon^2)$. Consider the j th iteration of the restricting procedure, where we search for a restricting vertex in U^j (with respect to the current U' and π'). Using W^j , we approximate for each $u \in U^j$ and each $i \in \bar{A}$ (where $u \in C_A$) how many neighbors u has in B -clusters such that $i \notin B$. To do so we simply perform queries on all pairs of vertices in $U' \times W^j$, and all pairs in $U^j \times W^j$. For each $u \in U^j$, let $R(u, i)$ be the set of neighbors of u in W_j that belong to B -clusters such that $i \notin B$. If for some $u \in U^j$, $|R(u, i)|/|W^j| \geq 3\epsilon/8$ for every $i \in \bar{A}$, then we conclude that u is a restricting vertex, and we select the first such vertex in U^j to add to U' . By our choice of the size of each W^j and an additive Chernoff bound, with probability at least $1 - \delta/4$, for every j , if U^j contains a restricting vertex then it will be detected. Furthermore, any vertex in U^j that is considered as restricting in fact has at least $(\epsilon/8) \cdot N$ neighbors (in $V(G)$) that belong to B -clusters such that $i \notin B$, for every $i \in \bar{A}$. Note that the first such vertex is uniquely determined by the set W^j , and that the time for implementing the restriction procedure is negligible with respect to the running time of the testing algorithm.

Once U' is constructed, we cluster all vertices in $V(G)$ according to U' and π' . To do so we simply perform all queries between vertices in $V(G)$ and U' . As for implementing the definition of forbidden sets, here each vertex $v \in V(G)$ samples its neighbors to determine F_v . Namely, for each vertex v we select a

sample of $\Theta(\log(k/(\epsilon\delta))/\epsilon^2)$ vertices, and approximate, for each $i \in \bar{A}$ (where $v \in C_A$), the number of neighbors that v has in B -clusters such that $i \notin B$. The expected number of vertices v for which these approximations differ significantly from the expected value is $O(\epsilon\delta \cdot N)$, and hence with probability at least $1 - \delta/4$, there are $(\epsilon/4) \cdot N$ vertices for which these approximations are in fact far from correct. Adding the contribution of these vertices to the number of violating edges accounted by Lemma 6.2.11 and Claim 6.2.13, we get an ϵ -good partition. The running time is hence governed by the implementation of the forbidden sets, and is $O(\log(k/(\epsilon\delta))/\epsilon^2) \cdot N$.

This completes the proof of Theorem 6.2.2. $\square$

7. Testing MAX-CLIQUE

Let $\omega(G)$ denote the size of the largest clique in graph G , and $\mathcal{C}_\rho \stackrel{\text{def}}{=} \{G : \omega(G) \geq \rho \cdot |V(G)|\}$ be the set of graphs having cliques of density at least ρ . Recall that $N \stackrel{\text{def}}{=} |V(G)|$. The main result of this section is:

THEOREM 7.1. *There exists a property testing algorithm, $\mathcal{A}$, for the class $\mathcal{C}_\rho$ whose query complexity and running time are*

$$O\left(\frac{\log^2(1/(\epsilon\delta))\rho^2}{\epsilon^6}\right) \text{ and } \exp\left(O\left(\frac{\log(1/(\epsilon\delta))\rho}{\epsilon^2}\right)\right),$$

respectively. In particular, $\mathcal{A}$ uniformly selects $O(\log(1/(\epsilon\delta))\rho/\epsilon^4)$ vertices in G and queries the oracle only on the existence of edges between (some of) these vertices. In case $G \in \mathcal{C}_\rho$, one can also retrieve in time $O(\log(1/(\epsilon\delta))\rho/\epsilon^2) \cdot N$ a set of $\rho \cdot N$ vertices in G , which is almost a clique (in the sense that it lacks at most $\epsilon \cdot N^2$ edges to being a clique).

Theorem 7.1 is proven by presenting a seemingly unnatural algorithm/tester (see below). However, as a corollary, we observe that the “natural” algorithm, which uniformly selects $\text{poly}(\epsilon^{-1} \log(1/\delta))$ many vertices and accepts iff they induce a subgraph with a clique of density $\rho - \epsilon/2$, is a valid $\mathcal{C}_\rho$ -tester as well.

COROLLARY 7.2. *Let $q(\cdot, \cdot)$ be the query complexity of algorithm $\mathcal{A}$ guaranteed by Theorem 7.1 (i.e., $q(\epsilon, \delta) = \text{poly}(\epsilon^{-1} \log(1/\delta))$), and let $\rho, \epsilon, \delta > 0$. Let R be a uniformly selected set of $m \stackrel{\text{def}}{=} 20 \cdot q(\epsilon/2, \delta/5)$ vertices in $V(G)$, and G_R be the subgraph (of G) induced by R . Then,*

—if $G \in \mathcal{C}_\rho$, then $\Pr_R[\omega(G_R) \leq (\rho - \epsilon/2) \cdot m] < \delta$.

—if $\text{dist}(G, \mathcal{C}_\rho) > \epsilon$, then $\Pr_R[\omega(G_R) \geq (\rho - \epsilon/2) \cdot m] < \delta$.

7.1 PROOF OF COROLLARY. Our presentation presupposes that $\mathcal{A}$ is given oracle access to a graph whose vertices may be an arbitrary subset of $[V(G)]$. If one insists that $\mathcal{A}$ only tests graphs with $|V(G)|$ vertices then another auxiliary trick is needed. Instead of providing $\mathcal{A}$ with oracle access to G_R (as done below), we provide it with oracle access to a graph in which each vertex of G_R is duplicated $|V(G)|/|R|$ times and edges are duplicated in the natural manner.

With the above technicality being settled, let us present the underlying ideas of the proof. The first item follows easily from the fact that $q(\epsilon, \delta) = \Omega(\epsilon^{-2} \log(1/\delta))$ (as such a sample is likely to hit enough vertices of the clique). The

issue is the second item. The basic idea is that if G_R has a clique of size $\rho' \stackrel{\text{def}}{=} \rho - (\epsilon/2)$ and we were to run $\mathcal{A}$ on it (with density parameter ρ'), then $\mathcal{A}$ would accept with high probability. On the other hand, when $\mathcal{A}$ is invoked on G (with density parameter ρ' and distance parameter $\epsilon/2$), algorithm $\mathcal{A}$ rejects with high probability. Loosely speaking, these two facts stand in contradiction since the samples that $\mathcal{A}$ sees in the two cases are statistically close. (Here we use the hypothesis that $\mathcal{A}$ queries the graph only on edges between vertices in a uniformly selected sample.) Actually, the above would be true if we were to set $m = O(q(\epsilon/2, \delta/5)^2)$ (as done in preliminary versions of this paper), as in both cases $\mathcal{A}$ would be likely to see a uniformly distributed set of $q(\epsilon/2, \delta/5)$ distinct vertices. This, however, does not happen with $m = O(q(\epsilon/2, \delta/5))$, since when sampling G_R algorithm $\mathcal{A}$ is likely to see several occurrences of some vertices (whereas this is unlikely when sampling G). Thus, we start by presenting an interface between the graph and algorithm $\mathcal{A}$. Loosely speaking, the interface obtains a sample 10 times bigger than what is required by $\mathcal{A}$, and passes to $\mathcal{A}$ a uniformly chosen subset of the required size consisting of distinct vertices.

INTERFACE TO $\mathcal{A}$. We use the fact that $q(\epsilon/2, \delta/5) > 2 \cdot \log(5/\delta)$ (otherwise, we would have needed to increase the size of m). Let $q \stackrel{\text{def}}{=} q(\epsilon/2, \delta/5)$, and suppose that $q < N/2$ (as, otherwise, we can just scan the entire graph). Given a sample of $10q$ vertices (with possible repetitions), the interface passes on the first q distinct vertices, and aborts if such elements do not exist. We adopt the convention by which whenever the interface aborts, the (modified) algorithm accepts.

Comment 7.1.1. Actually, we need to further modify the above interface so that collisions may occur with small probability, as would be the case when sampling independently q elements from $V(G)$. This is easy to do (e.g., one always passes the first vertex, and for $i = 2, \dots, q$ and $j = 1, \dots, i - 1$, one repeats the j th element in the i th location with probability $1/N$ and otherwise passes the i th element). In the rest of the proof, we ignore this modification of the interface.

CLAIM 7.1.2. *Suppose that $q \leq m/2$, and that S is a set of m distinct elements. Then, the probability that the interface aborts when given a sample of $10q$ elements uniformly and independently selected in S is smaller than $\delta/5$.*

PROOF. It suffices to upper bound the probability that a random multiset of $10q$ elements selected in a set of size $2q$ contains less than q distinct elements. Let ζ_i represent the event that either the $i - 1$ first multiset elements contain at least q distinct elements or the i th element in the multiset is different from each of the previous ones. Clearly, $\Pr[\zeta_i = 1 | \sum_{j < i} \zeta_j < q] \geq 1/2$, whereas $\Pr[\zeta_i = 1 | \sum_{j < i} \zeta_j \geq q] = 1$. Note that $\Pr[\sum_{i=1}^{10q} \zeta_i < q]$ represents the probability of aborting.

$$\begin{aligned} \Pr \left[\sum_{i=1}^{10q} \zeta_i < q \right] &< \sum_{s=1}^{q-1} \binom{10q}{s} \cdot 0.5^{10q-s} \\ &< 2^{q-1} \cdot \exp(-2 \cdot (0.5 - 0.1)^2 \cdot 10q) \\ &< 2^{-q/2}, \end{aligned}$$

which is at most $\delta/5$, as required. $\square$

Claim 7.1.2. will be applied both with respect to the set $V(G)$, and with respect to the subset R as in the statement of the corollary. The next claim refers to the latter set R , and is quite obvious.

CLAIM 7.1.3. *Suppose that $q \leq m/2$, and suppose that the interface is given a uniformly selected set $R \subset V(G)$ of size m . Then conditioned on not having aborted, the interface passes on a set of q vertices that is uniformly distributed among all such subsets of $V(G)$.*

We now turn to the actual claims of the corollary. As stated above, the first item is obvious (given that $q(\epsilon/2, \delta/5) = \Omega(\epsilon^{-2} \log(1/\delta))$). Our aim is to prove the second item. Suppose that $\text{dist}(G, \mathcal{C}_\rho) > \epsilon$, and consider what happens if we were to run $\mathcal{A}$, through the interface, on G_R with density parameter $\rho' = \rho - \epsilon/2$, distance parameter $\epsilon' = \epsilon/2$ and confidence parameter, $\delta' = \delta/5$. In such a case, algorithm $\mathcal{A}$ requires a sample of size $q = q(\epsilon', \delta')$, which is supplied by taking a random sample of $10q$ vertices of R and passing them through the interface to $\mathcal{A}$. By Claim 7.1.2, the interface aborts with probability less than $\delta/5$. By Claim 7.1.3, for a uniformly selected R , conditioned on the abort event not happening, algorithm $\mathcal{A}$ obtains a sample which is identically distributed to the sample it obtains when testing G under the same set of parameters, denoted $\pi' \stackrel{\text{def}}{=} (\rho', \epsilon', \delta')$. Denoting by $\mathcal{B}$ the composition of the interface with $\mathcal{A}$, we conclude that

$$|\Pr_R[\mathcal{B}^{G_R}(\pi') = 1] - \Pr[\mathcal{B}^G(\pi') = 1]| < \frac{\delta}{5}. \quad (6)$$

Observe that $\text{dist}(G, \mathcal{C}_{\rho'}) > \epsilon/2$ (as otherwise there exists $G' \in \mathcal{C}_{\rho'}$ so that $\text{dist}(G, G') \leq \epsilon/2$, whereas $\text{dist}(G', \mathcal{C}_\rho) \leq \rho\epsilon/2$ for all $G' \in \mathcal{C}_{\rho'} = \mathcal{C}_{\rho-(\epsilon/2)}$). Using this fact, Eq. (6) and the fact that the interface increases the success probability of $\mathcal{A}$ by at most $\delta/5$, we have

$$\begin{aligned} \Pr_R[\mathcal{B}^{G_R}(\pi') = 1] &< \Pr[\mathcal{B}^G(\pi') = 1] + \frac{\delta}{5} \\ &< \Pr[\mathcal{A}^G(\pi') = 1] + 2 \cdot \frac{\delta}{5} \\ &< \frac{\delta}{5} + \frac{2\delta}{5} \end{aligned}$$

On the other hand, using the fact that $\mathcal{B}$'s accepting probability is lower bounded by that of $\mathcal{A}$, we have,

$$\begin{aligned} \Pr_R[\mathcal{B}^{G_R}(\pi') = 1] &\geq \Pr_R[\omega(G_R) > \rho' \cdot m] \cdot \min_{G': \omega(G') > \rho' \cdot m} \{\Pr[\mathcal{B}^{G'}(\pi') = 1]\} \\ &> \Pr_R[\omega(G_R) > \rho' \cdot m] \cdot \left(1 - \frac{\delta}{5}\right). \end{aligned}$$

Approximate-Clique Finding Algorithm

Let $t = \Theta\left(\frac{\log(1/(\epsilon\delta)) \cdot \rho}{\epsilon^2}\right)$, and $r = \Theta\left(\frac{\log(1/(\epsilon\delta)) \cdot \rho}{\epsilon^2}\right)$.

- (1) Uniformly select two independent samples, U, W of sizes t and r , respectively.
- (2) For each $U' \subset U$ of cardinality $\frac{\rho}{2} \cdot t$, perform the following steps:
 - (a) Let $T(U') \subseteq V(G)$ be the subset of vertices in G which neighbor all vertices in U' (i.e., $T(U') \stackrel{\text{def}}{=} \{v \in V : \Gamma(v) \supseteq U'\} = V \cap (\bigcap_{u \in U'} \Gamma(u))$.
If $|T(U')| < \rho N$ then continue with the next subset U' .
 - (b) Let $W(U') \subseteq W$ be the subset of vertices in W which neighbor all vertices in U' (i.e., $W(U') \stackrel{\text{def}}{=} \{v \in W : \Gamma(v) \supseteq U'\} = W \cap (\bigcap_{u \in U'} \Gamma(u))$.
 - (c) For each $v \in T(U')$, compute $\hat{d}(v) \stackrel{\text{def}}{=} 2 \cdot |\Gamma(v) \cap W(U')|$. Let $C(U') \subseteq T(U')$ be the set of ρN vertices of the highest $\hat{d}(\cdot)$ value in $T(U')$. (Ties are broken by lexicographic order).
- (3) Among all sets $C(U')$, let $\tilde{C}$ be the one for which $\text{dist}(C(U'), C_N^\rho)$ is minimizes.
Output $\tilde{C}$.

FIG. 2. Approximate-Clique Finding Algorithm.

We conclude that

$$\begin{aligned} \Pr_R[\omega(G_R) > \rho' \cdot m] &< \frac{\Pr_R[\mathcal{B}^{\text{Gr}}(\pi') = 1]}{1 - (\delta/5)} \\ &< \frac{2\delta}{5} + \frac{3\delta}{5} \end{aligned}$$

as required. This completes the proof of Corollary 7.2. $\square$

7.2. THE APPROXIMATE-CLIQUE FINDER. We start by presenting and analyzing an algorithm that, given a graph G that has a clique of size ρN , finds a set of vertices of size ρN that is close to being a clique. Namely, for any given ϵ and δ , with probability at least $1 - \delta$, the algorithm finds a set of vertices of size ρN , such that there are at most ϵN^2 pairs of vertices in this set that are not connected by an edge. Our testing algorithm builds on this algorithm, which we call the *Approximate-Clique Finding* algorithm. The Approximate-Clique Finding algorithm (described in Figure 2) runs in time proportional to N^2 . We later show how an approximate clique can be found in time linear in N as asserted in Theorem 7.1. Throughout the analysis, we assume that $\epsilon < \rho^2$ (as otherwise $\text{dist}(G, \mathcal{C}_\rho) \leq \rho^2 \leq \epsilon$ for every graph G).

Notation. Let C_N^ρ denote the class of N -vertex graphs consisting of a clique of size ρN and $(1 - \rho) \cdot N$ isolated vertices. In the sequel, we denote by $\text{dist}(G', C_N^\rho)$ the relative distance (as a fraction of N^2) between a graph G' and C_N^ρ . In case G' contains less than N vertices we augment it by $N - |V(G')|$ isolated vertices. In all cases $|V(G')| \leq N$. With slight abuse of notation, for a set $X \subseteq V(G)$, we let $\text{dist}(X, C_N^\rho)$ denote the relative distance between the subgraph of G induces by X and C_N^ρ .

7.2.1. A Mental Experiment. To motivate our algorithm, we start with the following mental experiment. Assume that G has a clique C of size ρN . Suppose

we had an oracle that would tell us for every given vertex v whether v is a neighbor of every vertex in C . By querying the oracle on each vertex in the graph, we could determine the set of vertices, denoted $T(C)$, that neighbor every vertex in C . Note that since C is a clique, C is a subset of $T(C)$. Unfortunately, there might be many other vertices in $T(C)$ that do not belong to C . However, assume that we order the vertices in $T(C)$ according to their degree in the subgraph induced by $T(C)$ (breaking ties arbitrarily), and let C' be first ρN vertices according to this order. Then, we claim that C' is a clique (though it might be different from C). To see this, observe that by definition of $T(C)$, each vertex in C neighbors every vertex $T(C)$ (except itself). Thus, each vertex in C has degree $2(|T(C)| - 1)$ in the subgraph induced by $T(C)$, which is the maximum possible. (Recall that according to our convention, the degree of a vertex is the sum of its in-degree and its out-degree, which is twice its degree in the undirected representation of the graph.) Since $|C| = \rho N$, every vertex in C' must have degree $2(|T(C)| - 1)$ as well (because the vertices in C are all candidates for the set C' whose size is ρN as well). In other words, every vertex in C' neighbors every (other) vertex in $T(C)$, and in particular it neighbors every other vertex in C' , making C' a clique.

Suppose next that instead of having an oracle to C , we were given a uniformly chosen set U' in C of sufficient size (i.e., of size $\Theta(\epsilon^{-2} \cdot \log(1/(\epsilon\delta)))$). Let $T(U')$ be the set of vertices that neighbor every vertex in U' . Then, with high probability over the choice of U' , almost every vertex in $T(U')$ neighbors almost all vertices in C . Similarly to the above, we could order the vertices in $T(U')$ according to their degree in the graph induced by $T(U')$, and take the first ρN vertices. It can be shown (and we prove something slightly stronger later) that the resulting set is close to being a clique.

7.2.2. The Actual Algorithm. Since a uniformly chosen set in C is not provided to the algorithm, it instead “guesses” such a set. More precisely, it uniformly selects a set U from all graph vertices, and it considers all its subsets U' of size $\rho/2|U|$. Since with high probability $|U \cap C| \geq \rho/2|U|$, there exists a subset U' contained in C . Furthermore, with high probability, for this U' , almost all vertices in $T(U')$ (the set of vertices that neighbor every vertex in U'), neighbor almost every vertex in C . The next idea is to *approximate* the degree of each $v \in T(U')$ in the subgraph induced by $T(U')$, instead of computing it exactly. While this is not necessary for the efficiency of the Approximate-Clique Finding algorithm (as it runs in time quadratic in N anyhow), it will be useful to apply this modification here so as to simplify the analysis of the tester (which is presented later).

To approximate the degree of vertices in $T(U')$, we uniformly select an additional set, W , and let $W(U') \subseteq W$ contain all vertices in W that neighbor every vertex in U' . Thus, $W(U')$ is effectively sampling from $T(U')$. We now order the vertices in $T(U')$ according to the number of neighbors they have in $W(U')$, and take the first ρN vertices according to this order (if $|T(U')| < \rho N$ then necessarily $U' \not\subseteq C$ and we don't need to consider this U'). Thus for every subset U' we obtain a set of ρN vertices, and we output the one that misses the fewest edges to being a clique. (We note that U and W together play the role that the subset U of the sampled vertices plays in our other algorithms. Namely, U and W together are used to determined partitions of $V(G)$.)

The resulting Approximate-Clique Finding algorithm is given in Figure 2. In order to establish the correctness of the algorithm, we first show that with high probability, the set U has certain desired properties.

As before, let C be a clique of size ρN in G . For $\epsilon_1 \in [0, 1]$, we say that a set $U' \subset C \subset V(G)$ is ϵ_1 -clique-representative with respect to C , if for all but $\epsilon_1 N$ of the vertices, $v \in V(G)$,

$$\text{if } |\Gamma(v) \cap C| < (\rho - \epsilon_1)N \text{ then } \Gamma(v) \cap U' \neq U'. \quad (7)$$

That is, for most vertices v for which $\Gamma(v) \cap U' = U'$, it holds that $|\Gamma(v) \cap C| \geq (\rho - \epsilon_1)N$. Note that for every $v \in C$ the above condition holds for all $\epsilon_1 \geq 0$ (since $\Gamma(v) \supseteq C \supset U'$).

CLAIM 7.2.2.1. *Let $t = \Omega(\log(1/(\epsilon_1 \delta))/\epsilon_1 \rho)$, where $\epsilon_1 \in [0, 1]$. Let U be a uniformly chosen set of t vertices in G . Then, with probability at least $1 - \delta/2$, the set U contains an ϵ_1 -clique-representative subset of size $(\rho/2)t$.*

PROOF. Using a multiplicative Chernoff Bound (see Appendix B) we obtain that $|U \cap C| \geq \frac{1}{2}\rho t$ with probability at least $1 - \exp(-\Omega(\rho t)) > 1 - \delta/4$. Let us now consider a uniformly selected subset $U' \subset C$ of cardinality $t' \stackrel{\text{def}}{=} (\rho/2)t$. Clearly, for each $v \in C$, Eq. (7) holds. For each $v \in V(G) \setminus C$

$$\Pr_{U'}[\text{Eq. (7) does not hold for } v] = (1 - \epsilon_1)^{t'} < \frac{\epsilon_1 \delta}{4}, \quad (8)$$

where the last inequality is due to $t' = (\rho/2)t$ and the hypothesis regarding t . Thus, the expected number of vertices which violate Eq. (7) is bounded by $\delta/4 \cdot \epsilon_1 N$. Applying Markov's Inequality (see Appendix B) we conclude that with probability at least $1 - \delta/4$ there are at most $\epsilon_1 N$ vertices that violate Eq. (7). The lemma follows. $\square$

We next show that for an ϵ_1 -clique-representative U' , if we take a subset of $T(U')$ of size ρN which approximately contains the vertices that have highest degree in the subgraph induced by $T(U')$, then we obtain an approximate clique.

LEMMA 7.2.2.2. *Let $\epsilon_2 < \epsilon/\rho$ and $\epsilon_1 = \epsilon_2^2$. Let U' be ϵ_1 -clique-representative (with respect to C) and $T = T(U') \stackrel{\text{def}}{=} \bigcap_{u \in U'} \Gamma(u)$. Let α be such that αN is the degree of the $(\rho - \epsilon_2) \cdot N$ th vertex of highest degree in T . (We stress that we consider degrees in the subgraph of G induced by T .) Let $\tilde{C} \subseteq T$ be a set of size $\rho \cdot N$ that contains at least $(\rho - 3\epsilon_2) \cdot N$ vertices of degree at least $(\alpha - 2\epsilon_2) \cdot N$ (in the subgraph induced by T). Then, C satisfies $\text{dist}(\tilde{C}, C_N^\rho) \leq 8\epsilon_2 \rho$.*

The somewhat complicated formulation of the lemma is meant to fit its application in the analysis of the Approximate-Clique Finding Algorithm. Specifically, it addresses the fact that the algorithm uses approximations to the degrees of vertices in T , and that these approximations are slightly inaccurate for most vertices, and very inaccurate for few vertices.

PROOF. Clearly $T \supseteq C$ (as U' is a subset of the clique C). Let $H \stackrel{\text{def}}{=} \{v \in V(G) \setminus C : |\Gamma(v) \cap C| \geq (\rho - \epsilon_1) \cdot N\}$ be the set of vertices outside of C having many neighbors in C . Let $R \stackrel{\text{def}}{=} T \setminus (C \cup H)$ (i.e., the rest of T). Since U' is ϵ_1 -clique-representative, it follows that $|R| < \epsilon_1 N$ (since a vertex not in $C \cup H$

may enter T only if it is adjacent to all vertices in U' whereas it is adjacent to less than $(\rho - \epsilon_1)N$ vertices in C). Let the degree of vertex v in the subgraph induced by T be denoted by $\deg_T(v)$. (Recall that according to our convention, the degree of a vertex is the sum of its in-degree and its out-degree, which is twice its degree in the undirected representation of the graph.) We have,

$$\begin{aligned} \sum_{v \in C} \deg_T(v) &\geq 2 \cdot |\{(u, v) : u, v \in C\}| + 2 \cdot |\{(u, v) : u \in C, v \in H\}| \\ &\geq 2 \cdot (|C|^2 + |H| \cdot (\rho - \epsilon_1) \cdot N) \\ &> 2\rho N \cdot \left(\rho N + \left| H \right| - \frac{\epsilon_1}{\rho} N \right) \end{aligned} \quad (9)$$

Thus, the average value of $\deg_T(v)$ for $v \in C$ is at least $2 \cdot (\rho N + |H| - (\epsilon_1/\rho)N)$. On the other hand, the maximum value of $\deg_T(v)$ for $v \in C$ is bounded above by $2|T| = 2(|C| + |H| + |R|) \leq 2(\rho N + |H| + \epsilon_1 N)$. Therefore, the difference between the maximum value and the average value of $\deg_T(v)$ for $v \in C$ is $2 \cdot (\epsilon_1 + \epsilon_1/\rho) \cdot N$. By a simple counting argument (which is essentially a variant of Markov's Inequality) we have that for every k ,

$$\left| \left\{ v \in C : \deg_T(v) < 2 \left(\rho N + \left| H \right| - \frac{\epsilon_1}{\rho} N \right) - k(2\epsilon_1 \cdot (1 + \rho^{-1}) \cdot N) \right\} \right| < \frac{|C|}{k}.$$

Setting $k = \rho/\epsilon_2$, we get $|C|/k = \epsilon_2 N$, and $k \cdot 2\epsilon_1(1 + 1/\rho) \leq 4\epsilon_2$. Thus, at least $(\rho - \epsilon_2) \cdot N$ vertices in C have degree (in the subgraph induced by T) of at least $2\rho N + 2|H| - (2\epsilon_1/\rho) - 4\epsilon_2 N$. Since $(2\epsilon_1/\rho) = (2\epsilon_2^2/\rho)$, $\epsilon_2 < \epsilon/\rho$, and $\epsilon < \rho^2$, we have that α , as defined in the lemma, satisfies

$$\alpha \geq 2\rho + \frac{2|H|}{N} - 6\epsilon_2 = 2\rho + \frac{2(|T| - (|R| + |C|))}{N} - 6\epsilon_2 \geq \frac{2|T|}{N} - 8\epsilon_2. \quad (10)$$

By the lemma's hypothesis, we have $|\tilde{C}| = \rho \cdot N$. Also, denoting by $Z \subseteq \tilde{C}$ the set of vertices v for which $\deg_T(v) \geq (\alpha - 2\epsilon_2) \cdot N$, we have by the lemma's hypothesis $|Z| \geq (\rho - 3\epsilon_2) \cdot N$. By Eq. (10), we also have, for each $v \in Z$,

$$\begin{aligned} \deg_{\tilde{C}}(v) &\geq \deg_T(v) - 2 \cdot |T \setminus \tilde{C}| \\ &\geq [(2|T| - 8\epsilon_2 N) - 2\epsilon_2 N] - 2[|T| - \rho \cdot N] \\ &= 2(\rho - 5\epsilon_2) \cdot N. \end{aligned}$$

Summing up the degrees (in $\tilde{C}$) of all vertices in $\tilde{C}$, we obtain

$$\begin{aligned} \sum_{v \in \tilde{C}} \deg_{\tilde{C}}(v) &\geq \sum_{v \in Z} \deg_{\tilde{C}}(v) \\ &\geq 2 \cdot |Z| \cdot ((\rho - 5\epsilon_2) \cdot N) \\ &\geq 2 \cdot (\rho - 3\epsilon_2) \cdot (\rho - 5\epsilon_2) \cdot N^2 \end{aligned}$$

$$\begin{aligned}
&> 2 \cdot (\rho^2 - 8\epsilon_2\rho) \cdot N^2 \\
&= 2|\tilde{C}|^2 - 16\epsilon_2\rho \cdot N^2
\end{aligned}$$

It follows that $\tilde{C}$ is $8\epsilon_2\rho$ -close to being a clique and so the subgraph induced by it satisfies the claim of the lemma. $\square$

For a fixed U' (of cardinality $\rho/2 \cdot t$), let $T = T(U') \stackrel{\text{def}}{=} \bigcap_{u \in U'} \Gamma(u)$. We first prove that $\deg_{W(U')}(\cdot) (= \hat{d}(\cdot))$ provides a good estimate of $\deg_T(\cdot)$ (where recall that $W'(U) = W \cap T$). For $\epsilon_2 \in [0, 1]$, we say that a set W is ϵ_2 -representative (with respect to T), if for all but $\epsilon_2 N$ of the vertices, $v \in V(G)$,

$$\left| \frac{\deg_W \cap T(v)}{r} - \frac{\deg_T(v)}{N} \right| < \epsilon_2, \quad (11)$$

where for set of vertices Q , $\deg_Q(v)$ is the degree of v in the subgraph of G induced by Q .

CLAIM 7.2.2.3. *Let $r = \Omega(\log(1/\epsilon_2\delta)/\epsilon_2^2\rho)$, where $\epsilon_2 \in [0, 1]$. Suppose that $|T| \geq \rho N$ and that W is a uniformly selected subset of r vertices in $V(G)$. Then, with probability at least $1 - (\delta/4)$, the set W is ϵ_2 -representative of T .*

PROOF. By applying a multiplicative Chernoff Bound, we get

$$\Pr_W \left[\left| W \cap T \right| < \frac{\rho}{2} r \right] = \exp(-\Omega(\rho r)) < \frac{\delta}{8}.$$

We now consider a uniformly chosen $W' \subset T$ of size $r' \geq (\rho/2)r$. By applying a Chernoff Bound, we get for any fixed $v \in V(G)$,

$$\Pr_{W'} \left[\left| \frac{\deg_{W'}(v)}{r} - \frac{\deg_T(v)}{N} \right| \geq \epsilon_2 \right] = \exp(-\Omega(\epsilon_2^2 \rho r)) < \frac{\epsilon_2 \delta}{8}.$$

Applying Markov's Inequality, the claim follows. $\square$

As a corollary to Lemma 7.2.2.2, we have,

COROLLARY 7.2.2.4. *Let C be a ρN -Clique in G , let $\epsilon_2 < \epsilon/\rho$ and $\epsilon_1 = \epsilon_2^2$. Suppose that $U' \subset C$ is ϵ_1 -clique-representative with respect to C , and that W is ϵ_2 -representative with respect to $T(U')$. Then, the set of ρN vertices of highest $\hat{d}(\cdot)$ value (in $T(U') \supseteq C$) is $8\epsilon_2\rho$ -close to being a ρN -Clique. Recall that $\hat{d}(v) \stackrel{\text{def}}{=} 2 \cdot |\Gamma(v) \cap W(U')| = d_{W(U')}(v)$, and that $W(U') = W \cap T(U')$.*

PROOF (OF COROLLARY 7.2.2.4). Let $\tilde{C}$ be the set of ρN vertices with highest $\hat{d}(\cdot)$ value in $T = T(U')$ (where ties are broken arbitrarily). Let α be as defined in Lemma 7.2.2.2, (i.e., $\deg_T(v) \geq \alpha N$ for at least $(\rho - \epsilon_2)N$ vertices in T). By the hypothesis that W is ϵ_2 -representative of T , it follows that at least $(\rho - 2\epsilon_2)N$ of the vertices v of T satisfy

$$\frac{1}{r} \hat{d}(v) \geq \frac{1}{N} \deg_T(v) - \epsilon_2 \geq \alpha - \epsilon_2.$$

Clique-Degree Tester

Let $t = \Theta\left(\frac{\log(1/(\epsilon\delta)) \cdot \rho}{\epsilon^2}\right)$, $r = \Theta\left(\frac{\log(1/(\epsilon\delta)) \cdot \rho}{\epsilon^2}\right)$, and $m = \Theta\left(\frac{t + \log(1/\delta)}{\epsilon^2}\right)$.

- (1) Uniformly select three independent samples, U , W and $S = \{s_1, \dots, s_m\}$, of sizes t , r and m , respectively.
- (2) For each $U' \subset U$ of cardinality $\frac{\rho}{2} \cdot t$, perform the following steps:
 - (a) Let $S(U') \subseteq S$ be the subset of vertices which neighbor all vertices in U' (i.e., $S(U') \stackrel{\text{def}}{=} \{v \in S : \Gamma(v) \supseteq U'\} = S \cap (\bigcap_{u \in U'} \Gamma(u))$).
 - (b) Let $W(U') \subseteq W$ be the subset of vertices which neighbor all vertices in U' (i.e., $W(U') \stackrel{\text{def}}{=} \{v \in W : \Gamma(v) \supseteq U'\} = W \cap (\bigcap_{u \in U'} \Gamma(u))$).
 - (c) For each $v \in S(U')$, compute $\hat{d}(v) \stackrel{\text{def}}{=} 2 \cdot |\Gamma(v) \cap W(U')|$. Let $\hat{C}(U') \subseteq S(U')$ be the set of $\lfloor \rho m \rfloor$ vertices of the highest $\hat{d}(\cdot)$ value in $S(U')$. (Ties are broken by lexicographic order, and in case $|S(U')| < \lfloor \rho m \rfloor$ we let $\hat{C}(U') \stackrel{\text{def}}{=} S(U')$.)
 - (d) If the following two conditions hold then accept and halt.

Condition (a): $|\hat{C}(U')| \geq (\rho - \epsilon/80)m$.

Condition (b): $|\{(s_{2k-1}, s_{2k}) \in (\hat{C}(U') \times \hat{C}(U')) \setminus E(G)\}| \leq \frac{2\epsilon}{3} \cdot \frac{m}{2}$.
- (3) If none of these iterations made the algorithm accept G then it halts and rejects G .

FIG. 3. Clique-Degree Tester.

Since $\tilde{C}$ contains vertices with highest $\hat{d}(\cdot)$ value, it must contain at least $(\rho - 2\epsilon_2)N$ vertices of $\hat{d}(\cdot)$ value at least $(\alpha - \epsilon_2)r$. Using the hypothesis regarding W , we conclude that $\tilde{C}$ contains at least $(\rho - 3\epsilon_2)N$ vertices of $\deg_T(\cdot)$ value at least $(\alpha - 2\epsilon_2)N$. Using the hypothesis that U' is ϵ_1 -clique-representative with respect to C , we may now invoke Lemma 7.2.2.2 and the corollary follows. $\square$

The correctness of the Approximate-Clique Finding Algorithm follows from Claims 7.2.2.1 and 7.2.2.3, and Corollary 7.2.2.4, where $\epsilon_2 = \epsilon/(8\rho)$. Note that W need be ϵ_2 -representative only with respect to $T(U')$ where U' is ϵ_1 -clique-representative with respect to C .

7.3. THE CLIQUE-DEGREE TESTER. Given the correctness of the Approximate-Clique Finding algorithm, if we could sample from each $C(U')$ to tests whether it is close to being a clique, then we would obtain a testing algorithm for $\mathcal{C}_\rho$. Namely, if $G \in \mathcal{C}_\rho$, then from the previous subsection we know that with high probability over the choice of U and W , one of the sets $C(U')$ is close to being a clique, and if G is far from $\mathcal{C}_\rho$, then every set of size ρN , and in particular every $C(U')$, is far from being a clique. What we would like to do is test each $C(U')$ without actually first constructing it, that is without ordering all (the $\Omega(\rho N)$) vertices in $T(U')$. To this end, we uniformly choose an additional set of vertices, S , and essentially run the Approximate-Clique Finding algorithm on S . Namely, for every subset U' of U (of size $\rho/2|U|$) we let $S(U')$ be the subset of vertices in S that neighbor every vertex in U' (so that $S(U') = S \cap T(U')$). We then order the vertices in $S(U')$ according to the number of neighbors they have in $W(U')$. Finally we check whether for some U' the first $\rho|S|$ vertices in $S(U')$ (according to the above order) are close to being a clique. The resulting testing algorithm, called the *Clique-Degree Tester*, is described in Figure 3.

The correctness of the Clique-Degree Tester follows by two observations: (1) with high probability there exists an iteration where the sets U and W are as required in Corollary 7.2.2.4, and (2) the set S is a “good” sample of $T(U')$. We start by formulating the second observation in the lemma below. We note that when the lemma is applied, X is set to be $T(U')$, and the order on X is as determined by $\hat{d}(\cdot)$ (which is computed based on the number of neighbors each vertex has in $W(U')$ —see Figure 3).

LEMMA 7.3.1. *Let X be any subset of V , and assume $|X| \geq (\rho - \epsilon/40)N$. Consider a fixed ordering $x_1, \dots, x_{|X|}$ of the vertices in X , and let X' be the first $\min(\rho N, |X|)$ vertices in X according to the above ordering. Let $S = \{s_1, \dots, s_m\}$ be a uniformly selected set in V of size $m = \Omega(t + \log(1/\delta)/\epsilon^2)$, and let $S' \subseteq S$ be the first $\min(\lfloor \rho m \rfloor, |S \cap X|)$ vertices in $S \cap X$ according to the order defined on X . Then*

$$\Pr_S \left[\left| \frac{|\{(s_{2k-1}, s_{2k}) \in E(S', S')\}|}{m/2} - \frac{|E(X', X')|}{N^2} \right| > \frac{\epsilon}{3} \right] < \frac{\delta}{8} \cdot 2^{-t}.$$

Before proving Lemma 7.2.1, we state a simple claim that follows directly from an additive Chernoff bound.

CLAIM 7.3.2. *Let $S = \{s_1, \dots, s_m\}$ be a uniformly selected set of vertices of size $m \geq t + \log(32/\delta)/2\epsilon_3^2$, where $\epsilon_3 \in [0, 1]$. Then, for any fixed set of vertices X ,*

$$\Pr_S \left[\left| \frac{|S \cap X|}{m} - \frac{|X|}{N} \right| > \epsilon_3 \right] < \frac{\delta}{16} \cdot 2^{-t}.$$

PROOF OF LEMMA 7.3.1. Let $\epsilon_3 \stackrel{\text{def}}{=} \epsilon/40$, and let X'' be the first $(\rho - \epsilon_3)N$ vertices in X . (Throughout the proof, whenever we refer to a number of vertices such that the number is not an integer, we mean the floor of that number.) By Claim 7.3.2,

$$\Pr_S \left[\left| \frac{|S \cap X''|}{m} - (\rho - \epsilon_3) \right| > \epsilon_3 \right] < \frac{\delta}{16} \cdot 2^{-t}.$$

Thus, assume from now on that S is such that:

- (1) $|S \cap X''| \leq \rho m$, from which it follows (by definition of S') that $S' \cap X'' = S \cap X''$;
- (2) $|S \cap X''| \geq (\rho - 2\epsilon_3)m$, from which it follows that is $|S' \cap X''| \geq (\rho - 2\epsilon_3)m$.

Next observe that

$$\begin{aligned} \{(s_{2k-1}, s_{2k}) \in E(S', S')\} &= \{(s_{2k-1}, s_{2k}) \in E(S' \cap X'', S' \cap X'')\} \\ &\quad \cup \{(s_{2k-1}, s_{2k}) \in E(S' \setminus X'', S')\}. \end{aligned} \quad (12)$$

(Recall that for two sets of vertices A and B , we let $E(A, B)$ denote the set of edges with one end point in A and the other in B). By Item (1) above, and a

Chernoff bound to get that

$$\begin{aligned}
 \Pr_S \left[\left| \frac{|\{(s_{2k-1}, s_{2k}) \in E(S' \cap X'', S' \cap X'')\}|}{m/2} - \frac{|E(X'', X'')|}{N^2} \right| > \frac{\epsilon}{9} \right] \\
 = \Pr_S \left[\left| \frac{|\{(s_{2k-1}, s_{2k}) \in E(S \cap X'', S \cap X'')\}|}{m/2} - \frac{|E(X'', X'')|}{N^2} \right| > \frac{\epsilon}{9} \right] \\
 < \frac{\delta}{16} \cdot 2^{-t}. \quad (13)
 \end{aligned}$$

By definition of X' and X'' , we have that

$$\frac{|E(X', X')|}{N^2} - \frac{|E(X'', X'')|}{N^2} \leq 2\epsilon_3\rho < \frac{\epsilon}{9}. \quad (14)$$

By Item (2) above, we know that $|S' \setminus X''| \geq 2\epsilon_3 m$, and so

$$\begin{aligned}
 & \frac{|\{(s_{2k-1}, s_{2k}) \in E(S' \setminus X'', S')\}|}{m/2} \\
 & \leq \frac{|\{s_k \in S' \setminus X''\}|}{m/2} \leq 4\epsilon_3 < \frac{\epsilon}{9}. \quad (15)
 \end{aligned}$$

Summing up the probabilities of errors and substituting the bounds of Eqs. (14)–(15) in Eq. (13), the lemma follows. $\square$

COROLLARY 7.3.3. *Let $\mathcal{A}$ be the algorithm of Figure 3.*

- (1) *If $G \in \mathcal{C}_\rho$ then $\Pr[\mathcal{A}(G) = \text{accept}] > 1 - \delta$.*
- (2) *If $\text{dist}(G, \mathcal{C}_\rho) > \epsilon$, then $\Pr[\mathcal{A}(G) = \text{accept}] < \delta$.*

The main part of Theorem 7.1 follows from Corollary 7.3.3. The construction of an approximate clique in time linear in N is discussed following the proof of the corollary.

PROOF (OF COROLLARY 7.3.3). Let $\epsilon_2 = \epsilon/(24\rho)$ and let $\epsilon_1 = \epsilon_2^2$. In proving Part (1), we let C be an arbitrary ρN -Clique in G . By Claim 7.2.2.1 (and our choice of ϵ_1 and t), with probability at least $1 - (\delta/2)$, the set U contains an ϵ_1 -clique-representative (with respect to C) subset of size $(\rho/2)t$. Let us denote this subset by U' and recall that $U' \subset C$. We now consider the execution of Steps (1)–(4) with this U' . By Claim 7.2.2.3 (and our choice of ϵ_2 and r), with probability at least $1 - (\delta/4)$, the set W is ϵ_2 -representative of $T(U')$. Let $\tilde{C}$ denote the set of ρN vertices of highest $\hat{d}(\cdot)$ value in $T(U')$. By Corollary 7.2.2.4 and $8\epsilon_2\rho = \epsilon/3$, the set $\tilde{C}$ is $\epsilon/3$ -close to being a ρN -Clique. Applying Claim 7.3.2 to $T(U')$ with $\epsilon_3 = \epsilon/80$, Condition (a) of Step (4) holds with probability greater than $1 - (\delta/8)$. Applying Lemma 7.3.1 (with $X = T(U')$ and the order on X determined by $\hat{d}(\cdot)$), Condition (b) of Step (4) also holds with probability greater than $1 - (\delta/8)$ (as the fraction of missing edges is at most $(\epsilon/3)$ more than the

($\epsilon/3$) loss of $\tilde{C}$). Summing up the error probabilities, Part (1) of this corollary follows.

We now turn to prove Part (2). For any fixed choice of U' and W , we consider the set, denoted $C(U')$, of $\min(\rho N, |T(U')|)$ vertices of highest $\hat{d}(\cdot)$ value in $T = T(U')$. If $|C(U')| < (\rho - \epsilon/40)N$, then necessarily $T = C(U')$. Applying Claim 7.3.2 to T , Condition (a) of Step (4) is violated with probability greater than $1 - 2^{-t} \cdot \delta$. Otherwise, we apply Lemma 7.3.1 (again (with $X = T(U')$ and the order on X determined by $\hat{d}(\cdot)$). Since $\text{dist}(C(U'), C_N^{\rho}) \geq \epsilon$, with probability greater than $1 - 2^{-t} \cdot \delta$, Condition (b) of Step (4) is violated. We conclude that for every possible choice of U' and W , Step (4) accepts with probability bounded by $2^{-t} \cdot \delta$. Recalling that U and W are selected at random and less than 2^t possible $U' \subset U$ are tried, Part (2) follows. $\square$

A LINEAR (IN N) TIME ALGORITHM FOR FINDING AN APPROXIMATE CLIQUE. Given a graph G having a clique of size ρN , we can find an approximate clique (with high probability) as follows. We first run the Clique-Degree Tester (which with high probability accepts G), and record the sets U' and $W(U')$ which gave rise to the set $\hat{C}(U') \subseteq S(U')$ that was close to being a clique. We then determine the set $T(U')$ of vertices that neighbor every vertex in U' and order them according to $\hat{d}(\cdot)$ (i.e., according to the number of neighbors they have in $W(U')$). Finally, we take the first ρN vertices according to this order. Thus, the running time of this algorithm is the running time of the Clique-Degree Tester plus

$$O(|U'| \cdot N + |W| \cdot N) = O(\epsilon^{-2} \rho \log(1/(\epsilon\delta)) \cdot N)$$

as desired. The correctness of this algorithm follows by combining the arguments used above in the proof of Corollary 7.3.3. This completes the proof of Theorem 7.1. $\square$

8. Cut Problems

In this section, we present algorithms for testing ρ -Cut and ρ -Bisection. In both cases, we start by describing an algorithm that actually finds a partition that approximately maximizes the number of edges crossing the cut. In the case of the Bisection algorithm, the two sides of the partition are of equal size. We then show how these algorithms can be modified so as to *approximate* the size of the maximum cut. These modifications directly yield testing algorithms for the respective properties. Finally, we show how to achieve improved partitioning algorithms by first running the approximation algorithms. In the Bisection Subsection, we also describe how the algorithms can be easily modified to deal with the case in which one seeks a bisection that *minimizes* the number of edges crossing the cut.

8.1. TESTING MAX-CUT. For a given partition (V_1, V_2) of $V(G)$, let $\mu(V_1, V_2)$ denote the edge density of the cut defined by (V_1, V_2) . Namely,

$$\mu(V_1, V_2) \stackrel{\text{def}}{=} \frac{|\{(v, v') \in E(G) : \text{for } j \neq j', v \in V_j \text{ \& } v' \in V_{j'}\}|}{|V(G)|^2}.$$

Let $\mu(G)$ denote the edge density of the largest cut in G . Namely, it is the maximum of $\mu(V_1, V_2)$ taken over all partitions (V_1, V_2) of $V(G)$. The main results of this subsection are summarized below.

THEOREM 8.1

- (1) *There exists an algorithm that on input ϵ and δ and oracle access to a graph G , with probability at least $1 - \delta$, outputs a value $\hat{\mu}$ such that $|\hat{\mu} - \mu(G)| \leq \epsilon$. The algorithm has query complexity and running time*

$$O\left(\frac{\log^2(1/(\epsilon\delta))}{\epsilon^7}\right) \quad \text{and} \quad \exp\left(O\left(\frac{\log(1/(\epsilon\delta))}{\epsilon^3}\right)\right),$$

respectively.

- (2) *There exists an algorithm that on input ϵ and δ , and oracle access to G , runs in time*

$$\exp\left(O\left(\frac{\log(1/(\epsilon\delta))}{\epsilon^3}\right)\right) + O\left(\frac{\log(1/(\epsilon\delta))}{\epsilon^2}\right) \cdot N$$

and with probability at least $1 - \delta$ outputs a partition (V_1, V_2) of $V(G)$ such that $\mu(V_1, V_2) \geq \mu(G) - \epsilon$.

For any given ρ , let $\mathcal{MC}_\rho \stackrel{\text{def}}{=} \{G : \mu(G) \geq \rho\}$ be the class of graphs with cuts of density ρ . Note that for $\rho > 1/2$, the class $\mathcal{MC}_\rho$ is empty, and so the problem of testing ρ -Cut for $\rho > 1/2$ is trivial. For $\rho < 1/2$, we get the following as a corollary to Item (1) of Theorem 8.1.1.

COROLLARY 8.1.2. *For every constant $0 \leq \rho < 1/2$, there exists a property testing algorithm for the class $\mathcal{MC}_\rho$ whose query complexity and running time are*

$$O\left(\frac{\log^2(1/(\epsilon\delta))}{\epsilon^7}\right) \quad \text{and} \quad \exp\left(O\left(\frac{\log(1/(\epsilon\delta))}{\epsilon^3}\right)\right),$$

respectively.

Although very appealing, an approximation of $\mu(G)$ does not *directly* translate to a tester for the class $\mathcal{MC}_\rho$, for any ρ . The following proof provides the slightly more subtle connection. In particular, the complexity of the property testing algorithm asserted in Corollary 8.1.2 increases as a function of the inverse of $(1/2) - \rho$ (which is assumed to be a constant). When ρ is arbitrarily close to $1/2$, then we need to run a variant of the Bisection testing algorithm described in Subsection 8.2.4.

PROOF. Let $\gamma \stackrel{\text{def}}{=} \sqrt{1 - 2\rho}$. The testing algorithm runs the approximation algorithm referred to in Item (1) of Theorem 8.1, with approximation parameter $\epsilon' \stackrel{\text{def}}{=} (\gamma \cdot \epsilon)/3$ and confidence parameter δ (where ϵ and δ are the distance parameter and confidence parameter, respectively, of the testing algorithm). It accepts G if and only if $\hat{\mu} \geq \rho - \epsilon'$.

If $\mu(G) \geq \rho$, then by Item (1) of Theorem 8.1.1, G is accepted with probability $1 - \delta$, as required. Conversely, if the graph is accepted with probability greater

than δ , then $\mu(G) \geq \rho - 2\epsilon'$. We claim that this implies that G is ϵ -close to some graph G' in the class $\mathcal{MC}_\rho$. Details follow.

Let (V_1, V_2) be a partition of $V(G)$ such that $\mu(V_1, V_2) \geq \rho - 2\epsilon'$. Then necessarily, $2|V_1| \cdot |V_2| \geq (\rho - 2\epsilon')N^2$. If $2|V_1| \cdot |V_2| \geq \rho N^2$, then to obtain G' we can simply add edges between vertices in V_1 and vertices in V_2 until $\mu(V_1, V_2) = \rho$. In this case, $\text{dist}(G, G') \leq 2\epsilon' < \epsilon$. Otherwise, we cannot obtain G' by simply adding edges between V_1 and V_2 , as the total possible number of edges between V_1 and V_2 is less than ρN^2 . Instead, we first move a sufficient number of vertices from the larger set among V_1 and V_2 to the smaller set so as to “make room” for the needed number of added edges. Assume without loss of generality that $|V_1| < |V_2|$, and consider a partition (V'_1, V'_2) such that $V'_1 \supset V_1$, and $|V'_1|$ has the minimum value such that $2|V'_1| \cdot |V_2| \geq \rho N^2$. It is not hard to verify (by solving two quadratic equations) that $|V'_1| - |V_1| \leq \epsilon/\gamma \cdot N$, and so

$$\mu(V'_1, V'_2) \geq \mu(V_1, V_2) - \frac{\epsilon'}{\gamma} \geq \rho - \epsilon.$$

We can now proceed as in the first case by adding edges between V'_1 and V'_2 until we obtain a cut of the desired density. $\square$

A more natural property tester follows as in the case of ρ -Clique:

COROLLARY 7.1.3. *Let ρ be any non-negative constant smaller than $1/2$. Let $m = \text{poly}(\epsilon^{-1} \log(1/\delta))$ and let R be a uniformly selected set of m vertices in $V(G)$. Let G_R be the subgraph (of G) induced by R . Then,*

- if $G \in \mathcal{MC}_\rho$, then $\Pr_R[\mu(G_R) > \rho - (\epsilon/2)] > 1 - \delta$.
- if $\text{dist}(G, \mathcal{MC}_\rho) > \epsilon$, then $\Pr_R[\mu(G_R) \leq \rho - (\epsilon/2)] > 1 - \delta$.

Our algorithms can be easily generalized to approximate and test Max- k -way-Cut for $k > 2$ (see Section 8.1.5). Furthermore, since maximizing the density of cut edges effectively minimizes the density of edges inside the different components of the partition, the approximation algorithm for Max-Cut (and Max- k -way-Cut) can be used to test Bipartiteness (and respectively k -Colorability) as well. However, as opposed to our Bipartite (respectively, k -Colorability) testing algorithm, here we achieve a two-sided error (rather one-side error). That is, even if the input graph is bipartite (respectively, k -Colorable) it might be rejected (here) with probability greater than 0. Furthermore, for constant k , the sample complexity and running time of the algorithms presented here are also worse than those specifically intended to test bipartiteness and k -Colorability.

8.1.1. Organization. We start by presenting a quadratic-time partitioning algorithm, which given a graph G constructs a cut (i.e., a partition the vertices of the graph into two disjoint sets) of edge density at least $\mu(G) - (3/4)\epsilon$. This algorithm runs in time $\exp(\text{poly}(\log(1/\delta)/\epsilon)) \cdot N^2$ and is the basis for the approximation algorithm of Item (1) of Theorem 8.1.1. The algorithm claimed in Item (2) follows by combining the two algorithms. The extension to k -way cuts is presented in Section 8.1.5.

8.1.2. A Preliminary Graph Partitioning Algorithm. Let $\ell = \lceil 4/\epsilon \rceil$, and let $(V^1, \dots, V^\ell)$ be a fixed partition of $V(G)$ into ℓ sets of (roughly) equal size (say,

according to the order of the vertices in $V(G)$). In the *Graph Partitioning Algorithm* given below, we describe how to construct a partition (V_1, V_2) of $V(G)$ in ℓ iterations, where in the i th iteration we construct a partition (V_1^i, V_2^i) of V^i . The algorithm is essentially based on the following observation.

Let (H_1, H_2) be any fixed partition of $V(G)$. (In particular, we may consider a partition that defines a maximum cut). Let $v \in H_1$ and assume that v has at least as many neighbors in H_1 as it has in H_2 (i.e., $|\Gamma(v) \cap H_1| \geq |\Gamma(v) \cap H_2|$). Then by moving v from H_1 to H_2 we cannot decrease the edge density of the cut (and we might even increase it). Namely,

$$\mu(H_1 \setminus \{v\}, H_2 \cup \{v\}) \geq \mu(H_1, H_2).$$

Furthermore,

$$\mu(H_1 \setminus \{v\}, H_2 \cup \{v\}) - \mu(H_1, H_2) = \frac{2 \cdot (|\Gamma(v) \cap H_1| - |\Gamma(v) \cap H_2|)}{N^2}. \quad (16)$$

8.1.2.1. THE PLAN. Taking this observation one step further, we next show how we can define a new partition of V based on some (little) information concerning (H_1, H_2) , where we move $\Theta(\epsilon N)$ vertices between the components of the partition. Although we cannot ensure that the size of the cut does not decrease, as was the case when moving a single vertex, we can show that the size of the cut decreases by $O(\epsilon^2 N^2)$. (Note that, in the worst case, by moving $\Theta(\epsilon N)$ vertices, the size of a cut may decrease by $\Theta(\epsilon N^2)$.) We then show how such a process could be used by a graph partitioning algorithm given such information (which can be viewed as access to certain oracles). The oracle-aided algorithm will work in $O(1/\epsilon)$ stages. It will be viewed as starting from a partition corresponding to a maximum cut and moving $O(\epsilon N)$ vertices in each stage. The total decrease in the size of the cut will hence be bounded by $O((1/\epsilon) \cdot \epsilon^2 N) = O(\epsilon N)$. Finally, we show how the process can be implemented approximately (without any additional information).

8.1.2.2. AN “IDEAL” PROCEDURE. Let X be a subset of $V(G)$ of size N/ℓ , (where we assume for simplicity that ℓ divides N), let $W \stackrel{\text{def}}{=} V \setminus X$, and let (W_1, W_2) be the partition of W induced by (H_1, H_2) . That is, $W_1 \stackrel{\text{def}}{=} H_1 \cap W$, and $W_2 \stackrel{\text{def}}{=} H_2 \cap W$. Recall that (H_1, H_2) is some fixed partition of $V(G)$, and that in particular, we may consider a partition that defines a maximum cut. Assume we knew for every vertex $x \in X$ how many neighbors it has on each side of the partition (W_1, W_2) . In such a case, define X_{UB} to be the set of *unbalanced* vertices in X with respect to (W_1, W_2) . That is, X_{UB} is the set of vertices that have significantly (say $(1/8)\epsilon N$) more neighbors on one side of the partition than it has in the other. Analogously, define $X_{\text{B}} = X \setminus X_{\text{UB}}$ to be the set of *balanced* vertices with respect to (W_1, W_2) .

Assume we partition X into (X_1, X_2) as follows: Vertices in X_{UB} which have more neighbors in W_1 , are put in X_2 ; vertices in X_{UB} that have more neighbors in W_2 , are put in X_1 ; and vertices in X_{B} are placed arbitrarily. Based on this partition of X , we define a new partition of V : $(H'_1, H'_2) = (W_1 \cup X_1, W_2 \cup X_2)$, which differs from (H_1, H_2) only in the placement of vertices in X . Then the difference between $\mu(H'_1, H'_2)$ and $\mu(H_1, H_2)$, is only due to the change in the

Graph Partitioning Algorithm (for Max-Cut)

- (1) Choose $\ell = \lceil \frac{4}{\epsilon} \rceil$ sets $U^1, \dots, U^\ell$ each of size $t = \Theta\left(\frac{1}{\epsilon^2} \log \frac{1}{\epsilon \delta}\right)$, where U^i is chosen uniformly in $V \setminus V^i$. Let $\bar{U} = \langle U^1, \dots, U^\ell \rangle$.
- (2) For each sequence of partitions $\pi(\bar{U}) = \langle (U_1^1, U_2^1), \dots, (U_1^\ell, U_2^\ell) \rangle$ (where for each i , (U_1^i, U_2^i) is a partition of U^i) do
 - (a) For $i = 1 \dots \ell$, partition V^i into two disjoint sets V_1^i and V_2^i as follows:
 - For each $v \in V^i$,
 - i. If $|\Gamma(v) \cap U_1^i| \geq |\Gamma(v) \cap U_2^i|$ then put v in V_2^i .
 - ii. Otherwise put v in V_1^i .
 - (b) Let $V_1^{\pi(\bar{U})} = \bigcup_{i=1}^\ell V_1^i$, and let $V_2^{\pi(\bar{U})} = \bigcup_{i=1}^\ell V_2^i$.
- (3) Among all partitions $(V_1^{\pi(\bar{U})}, V_2^{\pi(\bar{U})})$, created in Step (2), let $(\tilde{V}_1^{\pi(\bar{U})}, \tilde{V}_2^{\pi(\bar{U})})$ be the one which defines the largest cut, and output it.

FIG. 4. Graph Partitioning Algorithm for Max-Cut.

number of edges between vertices in X and vertices in W , and between pairs of vertices in X . By definition of X_{UB} , and the way it was partitioned, the number of cut edges between vertices in X_{UB} and vertices in W could not have decreased. By definition of X_B , the arbitrary placement of these vertices decreased the number of cut edges between X_B and W by at most $|X_B| \cdot 2 \cdot (1/8)\epsilon N \leq (\epsilon/4\ell)N^2$, and the number of cut edges between pairs of vertices in X decreased by at most $|X|^2 = (1/\ell^2)N^2 \leq (\epsilon/4\ell)N^2$.

Now, let X be V^1 (i.e., the first N/ℓ vertices in lexicographical order), let (H_1, H_2) define a maximum cut, and let the partition resulting from the process defined above be denoted by (H_1^1, H_2^1) . Assume we continue iteratively, where in stage i we perform the above partitioning process for V^i , given the partition (H_1^{i-1}, H_2^{i-1}) determined in stage $i - 1$. That is, in stage i , we assume we know which vertices in V^i are unbalanced with respect to the partition of $V \setminus V^i$ induced by (H_1^{i-1}, H_2^{i-1}) . Then we can apply the same argument used above to each pair of consecutive partitions (H_1^i, H_2^i) and (H_1^{i+1}, H_2^{i+1}) , and get that $\mu(H_1^i, H_2^i)$ is smaller than $\mu(H_1, H_2) = \mu(G)$ by no more than an additive factor of $\ell \cdot \epsilon/2\ell = \epsilon/2$.

8.1.2.3. THE ACTUAL ALGORITHM. The graph partitioning algorithm, depicted in Figure 4, approximately implements the iterative procedure described above, starting from a partition (H_1^0, H_2^0) , which defines a maximum cut. Clearly, we do not have a clue as to what (H_1^0, H_2^0) is, and hence, in particular, we have no direct way of determining for a given vertex v in V^1 whether it is balanced or unbalanced with respect to the partition (W_1^0, W_2^0) of $W^0 \stackrel{\text{def}}{=} V \setminus V^1$ induced by this partition. However, we can approximate the number of neighbors v has on each side of (W_1^0, W_2^0) by sampling. Namely, if we uniformly choose a set of vertices U^1 of size $t = \text{poly}(\log(1/\delta)/\epsilon)$ in W^0 , then (as we later prove formally), with high probability over the choice of U^1 there exists a partition (U_1^1, U_2^1) of U^1 , which is *representative* with respect to (W_1^0, W_2^0) and V^1 in the following sense. For all but a small fraction of vertices v in V^1 , the number of neighbors v has in U_1^1 (U_2^1), relative to the size of U^1 , is approximately the same as the number of neighbors v has in W_1^0 (W_2^0), relative to the size of $V(G)$. Clearly such an approximation suffices since what is important when deciding where to put the vertices in V^1 is

to determine where to put the unbalanced vertices. If U^1 has a representative partition, then we say that U^1 is *good*. Since we do not know which of the 2^t partitions of U^1 is the representative one (assuming one exists), we simply try them all.

The choice of U^1 together with each of its partitions determines a partition of V^1 . While we must consider all partitions (U_1^1, U_2^1) of U^1 , we are only interested in the (hopefully representative) partition for which $U_1^1 \subset W_1^0$ and $U_2^1 \subset W_2^0$. Denote this partition by (U_1^1, U_2^1) . Let (V_1^1, V_2^1) be the partition of V^1 that is determined by this partition of U^1 , and let (H_1^1, H_2^1) be the resulting partition of $V(G)$. Namely, (H_1^1, H_2^1) is the same as (H_1^0, H_2^0) *except* for the placement of the vertices in V^1 , which is as in (V_1^1, V_2^1) . If in fact (U_1^1, U_2^1) is representative (with respect to (W_1^0, W_2^0) and V^1), then $\mu(H_1^1, H_2^1)$ is not much smaller than $\mu(H_1^0, H_2^0) = \mu(G)$. We continue in the same manner, where in stage i we randomly pick a set U^i , and for each of its partitions we determine a partition of V^i . Therefore, we are actually constructing $(2^t)^\ell = 2^{\ell \cdot t}$ possible partitions of $V(G)$, one for each partition of all the U^i 's. However, in order to show that at least one of these partitions defines a cut that is not much smaller than the maximum cut, we only need to ensure that for each i , with high probability, U^i is good with respect to (W_1^{i-1}, W_2^{i-1}) , where the latter partition is determined by the choice of $U^1, \dots, U^{i-1}$, and their representative partitions, $(U_1^1, U_2^1), \dots, (U_1^{i-1}, U_2^{i-1})$. The actual code is depicted in Figure 4. In the following lemma, we formalize the intuition given previously as to why the partitioning algorithm works.

LEMMA 8.1.2.3.1. *Let (H_1, H_2) be a fixed partition of $V(G)$. Then with probability at least $1 - \delta/2$ over the choice of $\bar{U} = \langle U^1 \cdots U^\ell \rangle$, there exists a sequence of partitions $\pi(\bar{U})$, such that $\mu(V_1^{\pi(\bar{U})}, V_2^{\pi(\bar{U})}) \geq \mu(H_1, H_2) - (3/4)\epsilon$.*

PROOF. For a given sequence of partitions $\pi(\bar{U})$, we consider the following $\ell + 1$ *hybrid* partitions. The hybrid (H_1^0, H_2^0) is simply (H_1, H_2) . The i th hybrid partition, (H_1^i, H_2^i) , has the vertices in $V^{i+1}, \dots, V^\ell$ partitioned as in (H_1, H_2) and the vertices in $V^1, \dots, V^i$ as placed by the algorithm. More precisely, the hybrid partition (H_1^i, H_2^i) is defined as follows:

$$H_1^i \stackrel{\text{def}}{=} W_1^{i-1} \cup V_1^i$$

and

$$H_2^i \stackrel{\text{def}}{=} W_2^{i-1} \cup V_2^i,$$

where for $j \in \{1, 2\}$,

$$V_j^i \stackrel{\text{def}}{=} V_j^{\pi(\bar{U})} \cap V^i, \quad W_j^{i-1} \stackrel{\text{def}}{=} H_j^{i-1} \cap W^{i-1}, \quad \text{and} \quad W^{i-1} \stackrel{\text{def}}{=} V \setminus V^i$$

Note, that in particular, (H_1^ℓ, H_2^ℓ) is the partition $(V_1^{\pi(\bar{U})}, V_2^{\pi(\bar{U})})$. Since the partition of each V^i is determined by the choice of U^i and its partition, the i th hybrid partition is determined by the choice of $U^1, \dots, U^i$ and their partitions, *but not* by the choice nor the partitions of $U^{i+1}, \dots, U^\ell$. We shall show that for every $1 \leq i \leq \ell$, for any fixed choice and partitions of $U^1, \dots, U^{i-1}$, with

probability at least $1 - (\delta/2\ell)$ over the choice of U^i , there exists a partition (U_1^i, U_2^i) of U^i such that

$$\mu(H_1^i, H_2^i) \geq \mu(H_1^{i-1}, H_2^{i-1}) - \frac{3\epsilon}{4\ell}.$$

The lemma will directly follow.

For the $i - 1$ hybrid partition (H_1^{i-1}, H_2^{i-1}) , or more precisely, for the partition it induces on W^{i-1} , and a sample set U^i , let

$$U_1^i \stackrel{\text{def}}{=} W_1^{i-1} \cap U^i$$

and

$$U_2^i \stackrel{\text{def}}{=} W_2^{i-1} \cap U^i.$$

We say that U^i is *good* with respect to (W_1^{i-1}, W_2^{i-1}) and V^i if (U_1^i, U_2^i) is *representative* with respect to (W_1^{i-1}, W_2^{i-1}) and V^i . That is, (U_1^i, U_2^i) is such that for all but a fraction of $(1/8)\epsilon$ of the vertices v in V^i , the following holds:

$$\text{For each } j \in \{1, 2\}, \quad \frac{|\Gamma(v) \cap U_j^i|}{t} = \frac{|\Gamma(v) \cap W_j^{i-1}|}{N} \pm \frac{\epsilon}{32}. \quad (17)$$

(Recall that $a = b \pm c$ is a shorthand for $b - c \leq a \leq b + c$.) Assume that in fact for each i , the set U^i is good with respect to (W_1^{i-1}, W_2^{i-1}) and V^i . As was previously defined, we say that a vertex v is *unbalanced* with respect to (W_1^{i-1}, W_2^{i-1}) if

$$\text{for } j, j' \in \{1, 2\}, \quad j \neq j' \quad |\Gamma(v) \cap W_j^{i-1}| \geq |\Gamma(v) \cap W_{j'}^{i-1}| + \frac{1}{8}\epsilon N. \quad (18)$$

Thus, if $v \in V^i$ is an unbalanced vertex with respect to (W_1^{i-1}, W_2^{i-1}) for which Eq. (17) is satisfied, then $|\Gamma(v) \cap U_j^i| \geq |\Gamma(v) \cap U_{j'}^i| + (1/16)\epsilon t$. We are hence guaranteed (by Steps (2.a.i) and (2.a.ii) of the algorithm) that, when the partition (U_1^i, U_2^i) is used, then v is put opposite the majority of its neighbors in W^{i-1} (according to their position in (W_1^{i-1}, W_2^{i-1})). If v is balanced, then it might be placed on either side of the partition. The same is true for the (at most $\epsilon/8 \cdot N/\ell$) vertices for which Eq. (17) does not hold.

As was noted previously, the decrease in the size of the cut is only due to changes in the number of edges between vertices in V^i and vertices in W^{i-1} , and between pairs of vertices in V^i . In particular:

- (1) The number of cut edges between unbalanced vertices in V^i for which Eq. (17) is satisfied and vertices in W^{i-1} can not decrease.
- (2) The number of cut edges between unbalanced vertices in V^i for which Eq. (17) is *not* satisfied and vertices in W^{i-1} decreases by at most $\epsilon/8 \cdot |V^i| \cdot 2N \leq (\epsilon/4\ell)N^2$.
- (3) The number of cut edges between balanced vertices in V^i and vertices in W^{i-1} decreases by at most $|V^i| \cdot 2 \cdot (1/8)\epsilon N \leq (\epsilon/4\ell)N^2$.

- (4) The number of cut edges between pairs of vertices in V^i decreases by at most $|V^i|^2 = (1/\ell^2)N^2 \leq (\epsilon/4\ell)N^2$.

The total decrease is bounded by $(3\epsilon/4\ell)N^2$.

It remains to prove that with high probability a chosen set U^i is good (with respect to (W_1^{i-1}, W_2^{i-1}) and V^i). We first fix a vertex $v \in V^i$. Let $U^i = \{u_1, \dots, u_t\}$. Recall that U^i is chosen uniformly in $W^{i-1} \stackrel{\text{def}}{=} V \setminus V^i$. For $j \in \{1, 2\}$, and for $1 \leq k \leq t$, define a 0/1 random variable, ξ_j^k , which is 1 if u_k is a neighbor of v and $u_k \in W_j^{i-1}$, and is 0 otherwise. By definition, for each j , the sum of the ξ_j^k s is simply the number of neighbors v has in $U_j^i (= U^i \cap W_j^{i-1})$ and the probability that $\xi_j^k = 1$ is $(1/N)|\Gamma(v) \cap W_j^{i-1}|$. By an additive Chernoff bound (see Appendix B), and our choice of t , for each $j \in \{1, \dots, k\}$,

$$\Pr_{U^i} \left[\left| \frac{|\Gamma(v) \cap U_j^i|}{t} - \frac{|\Gamma(v) \cap W_j^{i-1}|}{N} \right| > \frac{\epsilon}{32} \right] = \exp(-\Omega(\epsilon^2 t)) \leq \frac{\epsilon \delta}{32\ell}.$$

By Markov's inequality (see Appendix B), for each $j \in \{1, 2\}$, with probability at least $1 - (\delta/4\ell)$ over the choice of U^i , for all but $(1/8)\epsilon$ of the vertices in V^i , Eq. (17) holds (for that j), and thus with probability at least $1 - (\delta/2\ell)$, U^i is good as required. $\square$

Applying Lemma 8.1.2.3.1 to a maximum cut of G , we get:

COROLLARY 8.1.2.3.2. *With probability at least $1 - (\delta/2)$ over the choice of $\bar{U}$, we have $\mu(V_1^{\bar{\pi}(\bar{U})}, V_2^{\bar{\pi}(\bar{U})}) \geq \mu(G) - (3/4)\epsilon$, where $(V_1^{\bar{\pi}(\bar{U})}, V_2^{\bar{\pi}(\bar{U})})$ is as defined in Step (3) of the Graph Partitioning Algorithm (Figure 4).*

8.1.3. The Max-Cut Approximation Algorithm. Given the graph partitioning algorithm described above, the Max-Cut approximation algorithm is quite straightforward. We uniformly choose a set S of vertices of size

$$m = \Theta \left(\frac{\ell \cdot t + \log(1/\delta)}{\epsilon^2} \right),$$

and run the graph partitioning algorithm restricted to this set. The only small difference from what might be expected is that we do not necessarily output the size of largest cut among the cuts defined by the resulting partitions of S (i.e., those determined by the sequences of partitions $\pi(\bar{U})$). Instead, we view $S = \{s_1, \dots, s_m\}$ as a multiset of $m/2$ (ordered) pairs of vertices, (i.e., $\{(s_1, s_2), \dots, (s_{m-1}, s_m)\}$) and we choose the cut that maximizes the number of such pairs that are edges in the cut. This is done for technical reasons since it ensures a certain independence in the probabilistic analysis. The exact code is given in Figure 5.

LEMMA 8.1.3.1. *For any fixed $\bar{U}$, with probability at least $1 - \delta/2$ over the choice of S , $\hat{\mu}(S_1^{\bar{\pi}(\bar{U})}, S_2^{\bar{\pi}(\bar{U})}) = \mu(V_1^{\bar{\pi}(\bar{U})}, V_2^{\bar{\pi}(\bar{U})}) \pm (1/4)\epsilon$, where $(S_1^{\bar{\pi}(\bar{U})}, S_2^{\bar{\pi}(\bar{U})})$ and $\hat{\mu}(\cdot, \cdot)$ are as defined in Step (4) of the Max-Cut Approximation Algorithm.*

PROOF. Consider first a particular sequence of partitions, $\pi(\bar{U})$. The key observation is that for every $s \in S$, and for $j \in \{1, 2\}$, $s \in S_j^{\pi(\bar{U})}$ if and only if $s \in V_j^{\pi(\bar{U})}$. Thus for each sequence of partitions $\pi(\bar{U})$, we are effectively sampling from $(V_1^{\pi(\bar{U})}, V_2^{\pi(\bar{U})})$. Furthermore, by viewing S as consisting of $m/2$

Max-Cut Approximation Algorithm

- (1) As Step (1) of Figure 4.
- (2) Uniformly choose a set $S = \{s_1, \dots, s_m\}$ of size $m = \Theta\left(\frac{\ell \cdot t + \log(1/\delta)}{\epsilon^2}\right)$. For $1 \leq i \leq \ell$, let $S^i \stackrel{\text{def}}{=} V^i \cap S$.
- (3) Analogously to Step (2) of Figure 4, for each of the sequences of partitions $\pi(\bar{U}) = \langle (U_1^1, U_2^1), \dots, (U_1^\ell, U_2^\ell) \rangle$, partition each S^i into two disjoint sets S_1^i and S_2^i , and let $S_j^{\pi(\bar{U})} = \bigcup_{i=1}^\ell S_j^i$ (for $j = 1, 2$).
- (4) For each partition $(S_1^{\pi(\bar{U})}, S_2^{\pi(\bar{U})})$, compute the fraction of cut edges between pairs of vertices (s_{2k-1}, s_{2k}) . More precisely, define

$$\hat{\mu}(S_1^{\pi(\bar{U})}, S_2^{\pi(\bar{U})}) \stackrel{\text{def}}{=} \frac{\left| \left\{ (s_{2k-1}, s_{2k}) \in E\left((S_1^{\pi(\bar{U})}, S_2^{\pi(\bar{U})}) \cup E\left(S_2^{\pi(\bar{U})}, S_1^{\pi(\bar{U})}\right)\right\} \right|}{m/2}$$

Let $(\tilde{S}_1^{\pi(\bar{U})}, \tilde{S}_2^{\pi(\bar{U})})$ be a partition for which this fraction is maximized, and output $\hat{\mu}(\tilde{S}_1^{\pi(\bar{U})}, \tilde{S}_2^{\pi(\bar{U})})$.

FIG. 5. Max-Cut Approximation Algorithm.

pairs of vertices (s_{2k-1}, s_{2k}) , and counting the number of such pairs that are on opposite sides of the partition and have an edge in between, we are able to approximate the density of the cut edges. For $1 \leq k \leq m/2$, let ξ_k be a 0/1 random variable which is 1 if $(s_{2k-1}, s_{2k}) \in E(G)$, and for $j \neq j'$, $s_{2k-1} \in S_j^{\pi(\bar{U})}$ and $s_{2k} \in S_{j'}^{\pi(\bar{U})}$. Then, by definition, $\hat{\mu}(S_1^{\pi(\bar{U})}, S_2^{\pi(\bar{U})}) = (2/m) \sum_k^{m/2} \xi_k$, and the probability that $\xi_k = 1$ is $\mu(V_1^{\pi(\bar{U})}, V_2^{\pi(\bar{U})})$. Hence, by an additive Chernoff bound and our choice of m ,

$$\begin{aligned} \Pr_S \left[\left| \hat{\mu}(S_1^{\pi(\bar{U})}, S_2^{\pi(\bar{U})}) - \mu(V_1^{\pi(\bar{U})}, V_2^{\pi(\bar{U})}) \right| > \frac{1}{8} \epsilon \right] \\ = \exp(-\Omega(\epsilon^2 m)) \\ = O(\delta \cdot 2^{-\ell \cdot t}) \end{aligned} \quad (19)$$

Since there are $2^{\ell \cdot t}$ sequences of partitions of $\bar{U}$, with probability at least $1 - (\delta/2)$, for every sequence of partitions $\pi(\bar{U})$, $\hat{\mu}(S_1^{\pi(\bar{U})}, S_2^{\pi(\bar{U})}) = \mu(V_1^{\pi(\bar{U})}, V_2^{\pi(\bar{U})}) \pm (1/8)\epsilon$, and hence $\hat{\mu}(S_1^{\tilde{\pi}(\bar{U})}, S_2^{\tilde{\pi}(\bar{U})}) = \mu(V_1^{\tilde{\pi}(\bar{U})}, V_2^{\tilde{\pi}(\bar{U})}) \pm (1/4)\epsilon$. $\square$

Combining Corollary 8.1.2.3.2 and Lemma 8.1.3.1, Part (1) of Theorem 8.1.1 follows. Part (2) of the theorem is proved below.

8.1.4. An Improved Graph Partitioning Algorithm. The improved (in terms of running time) graph partitioning algorithm starts by invoking the Max-Cut Approximation Algorithm of Figure 5, and recording the sequence of sets $\bar{U}$ uniformly selected in Step (1) and the sequence of partitions $\tilde{\pi}(\bar{U})$ selected in Step (4). Using this specific sequence $\tilde{\pi}(\bar{U})$, the algorithm executes a *single* iteration of Step (2) of the Graph Partitioning Algorithm of Figure 4 and obtains the partition $(V_1^{\tilde{\pi}(\bar{U})}, V_2^{\tilde{\pi}(\bar{U})})$. Since this improved algorithm only partitions all vertices according to $\tilde{\pi}(\bar{U})$ and does not even compute the size of the resulting

cut (as the original algorithm did), its running time is linear in N instead of quadratic. More precisely, it is $O(t \cdot N) = O(\log(1/(\epsilon\delta)/e^2)) \cdot N$. As for its correctness, by Lemma 8.1.2.3.1, we have that with probability at least $1 - \delta/2$ over the choice of $\bar{U}$, there exists a sequence of partitions $\pi(\bar{U})$, such that $\mu(V_1^{\pi(\bar{U})}, V_2^{\pi(\bar{U})}) \geq \mu(G) - (3/4)\epsilon$. From the proof of Lemma 8.1.3.1, we have that for a fixed $\bar{U}$, with probability at least $1 - (\delta/2)$ over the choice of S , $\hat{\mu}(S_1^{\pi(\bar{U})}, S_2^{\pi(\bar{U})})$ is within $\epsilon/8$ from $\mu(V_1^{\pi(\bar{U})}, V_2^{\pi(\bar{U})})$ for every sequence of partitions $\pi(\bar{U})$. It follows that with probability at least $1 - \delta$, the recorded partition $\tilde{\pi}(\bar{U})$ is such that $\mu(V_1^{\tilde{\pi}(\bar{U})}, V_2^{\tilde{\pi}(\bar{U})}) \geq \mu(G) - \epsilon$, as required. Part (2) of Theorem 8.1.1 follows.

8.1.5. Generalization to k -way Cuts. For a k -way partition $(V_1, \dots, V_k)$ of $V(G)$, we denote by $\mu_k(V_1, \dots, V_k)$ the edge density of the cut defined by $(V_1, \dots, V_k)$. Namely,

$$\mu_k(V_1, \dots, V_k) \stackrel{\text{def}}{=} \frac{|\{(v, v') \in E(G) : \text{for } j \neq j', v \in V_j \text{ \& } v' \in V_{j'}\}|}{N^2}.$$

Let $\mu_k(G)$ denote the edge density of the largest k -way cut in G .

THEOREM 8.1.5.1

- (1) *There exists an algorithm that on input k , ϵ , and δ , and oracle access to a graph G , with probability at least $1 - \delta$, outputs a value $\hat{\mu}_k$ such that $|\hat{\mu}_k - \mu_k(G)| \leq \epsilon$. The algorithm has query complexity and running time*

$$O\left(\frac{\log^2(k/(\epsilon\delta))}{\epsilon^7}\right) \quad \text{and} \quad \exp\left(O\left(\frac{\log^2(k/(\epsilon\delta))}{\epsilon^3}\right)\right),$$

respectively.

- (2) *There exists an algorithm that on input k , ϵ , and δ , and oracle access to G , runs in time*

$$\exp\left(O\left(\frac{\log^2(k/(\epsilon\delta))}{\epsilon^3}\right)\right) + O\left(\frac{\log(k/(\epsilon\delta))}{\epsilon^2}\right) \cdot N$$

and with probability at least $1 - \delta$ outputs a k -way partition $(V_1, \dots, V_k)$ of $V(G)$ such that $\mu_k(V_1, \dots, V_k) \geq \mu_k(G) - \epsilon$.

PROOF. The graph k -way partitioning algorithm (respectively, Max- k -way-Cut approximation algorithm and testing algorithms), are obtained from the 2-way partitioning algorithm by the following simple modifications

- Instead of considering all two-way partitions of each U^i , we consider all its k -way partitions.
- For each such partition, $(U_1^i, \dots, U_k^i)$, we partition V^i (into k disjoint sets) as follows. For each vertex $v \in V^i$, and for $1 \leq j \leq k$, let $d_j^i(v) = |\Gamma(v) \cap U_j^i|$, and let $j_{\min} = \arg\min_j \{d_j^i(v)\}$. Then, we put v in $V_{j_{\min}}^i$.

The following (minor) changes suffice for adapting the analysis to the modified algorithm. Let $(H_1^0, \dots, H_k^0)$ be a k -way partition of $V(G)$ that defines a

maximum k -way cut. For a given choice and k -way partitions of $U^1, \dots, U^{i-1}$, let $(H_1^{i-1}, \dots, H_k^{i-1})$ be the $i-1$ hybrid k -way partition (defined analogously to the two-way cut case) which is determined by this choice and partitions of $U^1, \dots, U^{i-1}$. Using the same notation introduced in the two-way cut case, for a set U^i and for each $j \in \{1, \dots, k\}$ let $U_j^i \stackrel{\text{def}}{=} W_j^{i-1} \cap U^i$ where $W_j^{i-1} \stackrel{\text{def}}{=} H_j^{i-1} \setminus V^i$. We shall say that U^i is *good* with respect to $(W_1^{i-1}, \dots, W_k^{i-1})$ and V^i , if for all but $(1/8)\epsilon$ of the vertices v in V^i ,

$$\text{For each } j \in \{1, \dots, k\}, \frac{|\Gamma(v) \cap U_j^i|}{t} = \frac{|\Gamma(v) \cap W_j^{i-1}|}{N} \pm \frac{\epsilon}{32}. \quad (20)$$

It follows that in order to ensure that each U^i be good (with respect to $(W_1^{i-1}, \dots, W_k^{i-1})$ and V^i), we need to choose $t = |U^i|$ to be $\log k$ times larger than in the two-way cut case.

The notion of *balanced* and *unbalanced* vertices is generalized as follows: Consider a partition $(W_1^{i-1}, \dots, W_k^{i-1})$ of V . For vertex v and $j \in \{1, \dots, k\}$, let $\Gamma_j(v) \stackrel{\text{def}}{=} \Gamma(v) \cap W_j^{i-1}$ be the subset of v 's neighbors that belong to W_j^{i-1} . Let $\min(v) \stackrel{\text{def}}{=} \min_j \{|\Gamma_j(v)|\}$ be the minimum size among these neighbor sets, and let

$$J(v) \stackrel{\text{def}}{=} \left\{ j \in \{1, \dots, k\} : \left| \Gamma_j(v) \right| \leq \min(v) + \frac{\epsilon}{16} N \right\}$$

be the *balanced* (index) set of v . Note that in particular, for $k = 2$, either $|J(v)| = 1$, so that v is *unbalanced*, or $|J(v)| = k (= 2)$, so that v is *balanced*. Consider the case in which Eq. (20) holds for v . In such a case, v is placed in a component j' of the partition such that $j' \in J(v)$. Namely, it is placed in a component in which its number of neighbors is not much far from the minimum, $\min(v)$. The key point is that if v is moved from component j (in $(H_1^{i-1}, \dots, H_k^{i-1})$ to component j' (in $(H_1^i, \dots, H_k^i)$), all edges that it has with vertices in components $j'' \neq j, j'$ remain cut edges, and only the number of cut-edges between v and vertices in components j and j' might change. Thus, when $j' \in J(v)$, the number of cut edges between v and vertices in $V \setminus V^i$ does not decrease by much. As in the case of $k = 2$, we are essentially “giving up” on all cut edges between pairs of vertices in V^i , and cut edges that are incident to vertices for which Eq. (20) does *not* hold. Finally, since the number of k -way partitions of all the U^i 's is $k^{t \cdot \ell}$, we must choose m (the size of S in the Max- k -way-Cut approximation algorithm) to be $\approx \Theta(\ell t / \epsilon^2 \cdot \log k)$ rather than $\approx \Theta(\ell t / \epsilon^2)$ (as our choice in the case of two-way cuts). $\square$

8.2. TESTING BISECTION. In this section, we study a variant of the Max-Cut problem in which both sides of the partition are required to be of equal size. Namely, using the notation presented in Section 8.1, let

$$\mu^{(1/2)}(G) \stackrel{\text{def}}{=} \max_{V_1 \subset V(G), |V_1| = N/2} \mu(V_1, V(G) \setminus V_1).$$

A partition (V_1, V_2) of $V(G)$, such that $|V_1| = |V_2| = N/2$ is called a *bisection*.²² For sake of the exposition (due to the similarity to Max-Cut), we first consider

²² We assume throughout this section that N is even. In case N is odd, one may require $|V_1| = |V_2| + 1 = (N + 1)/2$.

the less standard problem of maximizing the (number of edges crossing the) bisection. The (more standard) case of minimization is handled analogously (see Section 8.2.4). The main result of this section is

THEOREM 8.2.1

- (1) *There exists an algorithm that on input ϵ and δ and oracle access to a graph G , with probability at least $1 - \delta$, outputs a value $\hat{\mu}^{(1/2)}$ such that $|\hat{\mu}^{(1/2)} - \mu^{(1/2)}(G)| \leq \epsilon$. The algorithm has query complexity and running time*

$$O\left(\frac{\log^2(1/(\epsilon\delta))}{\epsilon^8}\right) \quad \text{and} \quad \exp\left(O\left(\frac{\log(1/(\epsilon\delta))}{\epsilon^2}\right)\right),$$

respectively.

- (2) *There exists an algorithm that on input ϵ , and δ , and oracle access to a graph G , runs in time*

$$\exp\left(O\left(\frac{\log(1/(\epsilon\delta))}{\epsilon^3}\right)\right) + O\left(\frac{\log(1/(\epsilon\delta))}{\epsilon^3}\right) \cdot N$$

and with probability at least $1 - \delta$ outputs a bisection (V_1, V_2) of $V(G)$ such that $\mu(V_1, V_2) \geq \mu^{(1/2)}(G) - \epsilon$.

Item (1) yields a property tester for the class $\mathcal{M}_\rho^{\mathcal{C}_p^{(1/2)}} \stackrel{\text{def}}{=} \{G : \mu^{(1/2)}(G) \geq \rho\}$, for every $0 \leq \rho \leq 1$, with query complexity $O(\log^2(1/(\epsilon\delta))/\epsilon^8)$ and running time $\exp(O(\log(1/(\epsilon\delta))/\epsilon^3))$. If $\rho > 1/2$ then the tester rejects G since for $\rho > 1/2$ the class $\mathcal{M}_\rho^{\mathcal{C}_p^{(1/2)}}$ is empty. Otherwise the tester runs the approximation algorithm referred to in Item (1) with approximation parameter $\epsilon/2$ and confidence parameter δ (where ϵ and δ are the distance and confidence parameters, respectively, of the testing algorithm). The tester accepts G if and only if $\hat{\mu}^{(1/2)} \geq \rho - \epsilon/2$. If $G \in \mathcal{M}_\rho^{\mathcal{C}_p^{(1/2)}}$ (i.e., $\mu^{(1/2)}(G) \geq \rho$), then by Item (1) it is accepted with probability at least $1 - \delta$. Conversely, if G is accepted with probability greater than δ , then $\mu^{(1/2)}(G) \geq \rho - \epsilon$. That is, there exists an equal partition (V_1, V_2) of G such that $\mu(V_1, V_2) \geq \rho - \epsilon$. Therefore, $\text{dist}(G, \mathcal{M}_\rho^{\mathcal{C}_p^{(1/2)}}) \leq \epsilon$, since by adding at most ϵN^2 edges between V_1 and V_2 we can obtain a graph $G' \in \mathcal{M}_\rho^{\mathcal{C}_p^{(1/2)}}$.

A more natural property tester follows as in previous cases:

COROLLARY 8.2.2. *Let $m = \text{poly}(\epsilon^{-1} \log(1/\delta))$ and let R be a uniformly selected set of m vertices in $V(G)$. Let G_R be the subgraph (of G) induced by R . Then,*

- if $G \in \mathcal{M}_\rho^{\mathcal{C}_p^{(1/2)}}$, then $\Pr_R[\mu^{(1/2)}(G_R) > \rho - \epsilon/2] > 1 - \delta$.
- if $\text{dist}(G, \mathcal{M}_\rho^{\mathcal{C}_p^{(1/2)}}) > \epsilon$, then $\Pr_R[\mu^{(1/2)}(G_R) \leq \rho - \epsilon/2] > 1 - \delta$.

Our proof of Theorem 8.2.1 follows the outline of the proof of Theorem 8.1.1 (i.e., the analysis of the Max-Cut algorithms). However, there is one crucial difference between the problem of constructing (respectively, approximating the size of) a maximum cut and the problem of constructing (respectively, approximating the size of) a bisection: In a bisection, both sides of the cut must be of equal size. This has the following consequence. Recall that in case of Max-Cut, it is always beneficial (and possible) to relocate a vertex so that it is on the side

opposite the majority of its neighbors. In contrast, when restricted to bisections, this property no longer holds: a maximum size bisection may have vertices which belong to the same side of the bisection as the majority of their neighbors. Thus, when partitioning a subset V^i , we can not simply put all vertices (or all unbalanced vertices) on the side opposite the majority of their neighbors.

However, we can still use information concerning the unbalance of vertices with respect to a given bisection (and some additional information) in order to define a new bisection whose cut is not much smaller. Consider an arbitrary bisection (H_1, H_2) , and an arbitrary set of vertices X of size $O(\epsilon N)$. Assume we are told how many neighbors each vertex in X has in $W_1 \stackrel{\text{def}}{=} H_1 \setminus X$ and how many in $W_2 \stackrel{\text{def}}{=} H_2 \setminus X$. Further assume that we know $|H_1 \cap X|$ and $|H_2 \cap X|$. Let us relate with each vertex in X an *unbalance* value, which is simply the fraction of neighbors it has in W_2 (among all N possible neighbors) minus the fraction of neighbors it has in W_1 . This value (which ranges from 1 to -1) tries to capture our “preference” of placing a vertex on side 1.

Assume we now repartition the vertices in X so that there are $|H_1 \cap X|$ vertices on side 1 (and $|H_2 \cap X|$ on side 2) and all vertices on side 1 have unbalance value which is greater or equal to the unbalance value of any vertex on side 2. Then the resulting partition is clearly a bisection and the size of the cut decreases by at most $|X|^2 = O(\epsilon^2 N^2)$. The latter is due to the fact that for any other partition of X with $|H_1 \cap X|$ vertices on one side (and the rest on the other), there cannot be more cut edges between vertices in X and vertices in $V \setminus X$ than in the partition defined above. The decrease in the size of the cut is hence due to the decrease in the number of cut edges between pairs of vertices in X . Our graph bisection algorithm is based on this observation.

8.2.1. The Preliminary Bisection Algorithm. As was done in Section 8.1, we start by describing an algorithm that is aided by certain oracles, and then show how to simulate these oracles. Similarly to the (oracle-aided) graph partitioning algorithm for Max-Cut, the (oracle-aided) graph bisection algorithm proceeds in $\ell = O(1/\epsilon)$ iterations where in the i th iteration the vertices in V^i are partitioned into two subsets, V_1^i and V_2^i . Here too we think of the algorithm as defining *hybrid* partitions. Starting from the zero hybrid partition $\mathcal{H}^0 = (H_1^0, H_2^0)$, which is a maximum bisection, the i th hybrid partition $\mathcal{H}^i = (H_1^i, H_2^i)$ is a hybrid of the partition (V_1^i, V_2^i) of V^i (constructed in the i th iteration), and the partition $\mathcal{W}^{i-1} = (W_1^{i-1}, W_2^{i-1})$ of $W^{i-1} \stackrel{\text{def}}{=} V \setminus V^i$ induced by the $i-1$ hybrid partition, $\mathcal{H}^{i-1} = (H_1^{i-1}, H_2^{i-1})$. However, differently from the Max-Cut graph partitioning algorithm, here we might place vertices of V^i which are *unbalanced* with respect to $\mathcal{W}^{i-1}$ on the same side of the partition as the majority of their neighbors. This is done so to maintain the desired proportion (of vertices belonging to V^i) on each side of the new hybrid. That is, for each $i \in \{1, \dots, \ell\}$ let

$$\beta^i \stackrel{\text{def}}{=} \frac{|V^i \cap H_1^{i-1}|}{N/\ell} \quad (21)$$

be the fraction of vertices in V^i which belong to H_1^{i-1} . Assume we knew all β^i 's. If in each iteration, i , we make sure to put β^i of the vertices in V^i on side 1 and $(1 - \beta^i)$ on side 2, then since $\mathcal{H}^0$ is a bisection, so will be each hybrid partition, and in particular the final partition which the algorithm outputs.

Further assume that in each iteration of the algorithm we knew *exactly* how many neighbors each vertex in V^i has on each side of the partition. In such a case, we could compute for each vertex v its *unbalance* value:

$$\text{ub}(v) \stackrel{\text{def}}{=} \frac{|\Gamma(v) \cap W_2^{i-1}| - |\Gamma(v) \cap W_1^{i-1}|}{N}. \quad (22)$$

Let $L \stackrel{\text{def}}{=} N/\ell$ (where for simplicity we assume ℓ divides N), and $v_1, \dots, v_L$ be an ordering of the vertices in V^i according to their unbalance value; that is, $\text{ub}(v_k) \geq \text{ub}(v_{k+1})$. Consider the following partition (V_1^i, V_2^i) of V^i : $V_1^i \stackrel{\text{def}}{=} \{v_1, \dots, v_{\beta^i L}\}$, and $V_2^i \stackrel{\text{def}}{=} V^i \setminus V_1^i$. Let $(H_1^i, H_2^i) \stackrel{\text{def}}{=} (W_1^{i-1} \cup V_1^i, W_2^{i-1} \cup V_2^i)$. Then the number of cut edges in (H_1^i, H_2^i) between vertices in V^i and vertices in W^{i-1} is at least as large as in (H_1^{i-1}, H_2^{i-1}) . This is true since our partition of V^i , by definition, maximizes the number of such cut edges among all partitions which are a hybrid between (W_1^{i-1}, W_2^{i-1}) and a partition of V^i into two subsets of size $\beta^i L$ and $(1 - \beta^i)L$ respectively. The decrease in the size of the cut is hence at most $|V^i|^2 = (N/\ell)^2 = O((\epsilon/\ell)N^2)$.

We next remove the assumptions that we know β^i as well as $\text{ub}(v)$, for every $i \in \{1, \dots, \ell\}$ and $v \in V^i$. Firstly, we note that approximations (up to $O(\epsilon)$) to these values are good enough. Actually, we can afford having bad approximations of $\text{ub}(v)$ for an $O(\epsilon)$ fraction of the vertices in V^i . Similarly to the Max-Cut partitioning algorithm, we use sample sets U^i to obtain approximations $\widehat{\text{ub}}(v)$ to $\text{ub}(v)$. Namely, for $v \in V^i$ and for a partition (U_1^i, U_2^i) of U^i (where we consider all such partitions), we let

$$\widehat{\text{ub}}(v) \stackrel{\text{def}}{=} \frac{|\Gamma(v) \cap U_1^i| - |\Gamma(v) \cap U_2^i|}{t}, \quad (23)$$

where t is the size of each U^i . The approximations $\hat{\beta}^i = \beta^i \pm O(\epsilon)$ are obtained by simply trying all integer multiples of $\epsilon/8$ which sum up to $1/2$.²³ Each possible setting of $\hat{\beta}^1, \dots, \hat{\beta}^\ell$ and sequence of partitions of $U^1, \dots, U^\ell$ gives rise to a different bisection of V , and we choose the resulting bisection whose cut is maximized.

We show that with high probability over the choice of the U^i 's, there exist partitions of these sets, and there always exists a setting of the $\hat{\beta}^i$'s, so that at least one of the resulting bisection is close to having the maximum number of crossing edges. In particular, let $\hat{v}_1, \dots, \hat{v}_L$ be an ordering of the vertices in V^i according to $\widehat{\text{ub}}(v)$. As we prove in Lemma 8.2.1.1, if we put the vertices $\{\hat{v}_1, \dots, \hat{v}_{\lfloor \hat{\beta}^i L \rfloor}\}$ on side 1, and the vertices $\{\hat{v}_{\lfloor \hat{\beta}^i L \rfloor + 1}, \dots, \hat{v}_L\}$ on side 2, then the number of crossing edges in the resulting hybrid partition is not much smaller than that defined by the previous hybrid partition.

A detailed description of the graph bisection algorithm is given in Figure 6, and its formal analysis is provided in Lemma 8.2.1.1.

²³ This is always possible in case $1/\epsilon$ is an integer. Otherwise, we can try all integer multiples of $\epsilon'/8$ which sum up to $1/2$, where $\epsilon' = 1/\lceil 1/\epsilon \rceil$. Since $\epsilon' > \epsilon/2$, for simplicity we assume that $1/\epsilon$ is in fact an integer.

Graph Bisection Algorithm

- (1) Choose $\ell = \lceil \frac{1}{\epsilon} \rceil$ sets $U^1, \dots, U^\ell$ each of size $t = \Theta\left(\frac{1}{\epsilon^2} \log \frac{1}{\epsilon\delta}\right)$, where U^i is chosen uniformly in $V \setminus V^i$.
- (2) For each sequence of partitions $\Pi = \langle (U_1^1, U_2^1), \dots, (U_1^\ell, U_2^\ell) \rangle$ and for each of the ℓ -tuples $\hat{\beta} = (\hat{\beta}^1, \dots, \hat{\beta}^\ell)$, where each $\hat{\beta}^i \in [0, 1]$ is an integer multiple of $\frac{\epsilon}{8}$, and $\sum_i \hat{\beta}^i = \frac{1}{2}$, construct a bisection^a $(V_1^{\Pi, \hat{\beta}}, V_2^{\Pi, \hat{\beta}})$ as follows:
 - (a) For $i = 1 \dots \ell$ do:
 - i. For each $v \in V^i$ let $\widehat{ub}(v) \stackrel{\text{def}}{=} \frac{1}{t} (|\Gamma(v) \cap U_2^i| - |\Gamma(v) \cap U_1^i|)$;
 - ii. Let $\hat{v}_1, \dots, \hat{v}_L$ be an ordering of the vertices in V^i such that $\widehat{ub}(\hat{v}_k) \geq \widehat{ub}(\hat{v}_{k+1})$ (ties are broken according to lexicographical order).
 - iii. $V_1^i \leftarrow \{\hat{v}_1, \dots, \hat{v}_{\lfloor \hat{\beta}^i L \rfloor}\}$, and $V_2^i \leftarrow \{\hat{v}_{\lfloor \hat{\beta}^i L \rfloor + 1}, \dots, \hat{v}_L\}$
 - (b) Let $V_1^{\Pi, \hat{\beta}} = \bigcup_{i=1}^\ell V_1^i$, and let $V_2^{\Pi, \hat{\beta}} = \bigcup_{i=1}^\ell V_2^i$.
- (3) Among all bisections, $(V_1^{\Pi, \hat{\beta}}, V_2^{\Pi, \hat{\beta}})$, let $(V_1^{\Pi, \hat{\beta}}, V_2^{\Pi, \hat{\beta}})$ be the one with maximum number of crossing edges.

^aDue to questions of integrability (i.e. $\hat{\beta}^i L$ not being an integer) we might not get an exact bisection. However this can easily be treated by moving (at most ℓ) vertices between sides once the partition is constructed.

FIG. 6. Graph Bisection Algorithm.

LEMMA 8.2.1.1. *Let (H_1, H_2) be a fixed bisection of $V(G)$. Then, with probability at least $1 - \delta/2$ over the choice of $\bar{U}$, there exists a sequence of partitions Π of $\bar{U}$, and an ℓ -tuple $\hat{\beta}$, such that $\mu(V_1^{\Pi, \hat{\beta}}, V_2^{\Pi, \hat{\beta}}) \geq \mu(H_1, H_2) - (3/4)\epsilon$.*

PROOF. Let $\hat{\beta}$ be such that for every i , $\hat{\beta}^i = \beta^i \pm \epsilon/16$, where $\beta^i \stackrel{\text{def}}{=} |H_1 \cap V^i|$. The existence of such $\hat{\beta}^i$ follows from the resolution of the values taken by $\hat{\beta}^i$ and the fact that all possibilities (to within this resolution) were tried. For a fixed $\bar{U}$, a fixed sequence of partitions Π of $\bar{U}$, and the above choice of $\hat{\beta}$, let the i th hybrid partition (H_1^i, H_2^i) determined by Π and $\hat{\beta}$ be defined as follows. For $i = 0$, the partition (H_1^0, H_2^0) equals (H_1, H_2) . For $i > 0$ and $j \in \{1, 2\}$, let $W_j^{i-1} \stackrel{\text{def}}{=} H_j^{i-1} \setminus V^i$. As done in the Graph Bisection Algorithm (see Figure 6), let $\hat{v}_1, \dots, \hat{v}_L$ be an ordering of the vertices in V^i according to their $\widehat{ub}(\cdot)$ values. Let $V_1^i \stackrel{\text{def}}{=} \{\hat{v}_1, \dots, \hat{v}_{\lfloor \hat{\beta}^i L \rfloor}\}$, and $V_2^i \stackrel{\text{def}}{=} \{\hat{v}_{\lfloor \hat{\beta}^i L \rfloor + 1}, \dots, \hat{v}_L\}$. Then $H_j^i \stackrel{\text{def}}{=} W_j^{i-1} \cup V_j^i$.

Similarly to what was observed in Lemma 8.1.2.3.1, for $\hat{\beta}$ fixed as above, an i th hybrid partition is actually determined by the partitions of $U^1, \dots, U^i$ (which determine the $\widehat{ub}(\cdot)$ values as in Eq. (23)), and does not depend on the partitions of $U^{i+1}, \dots, U^\ell$. Similarly to the analysis of the graph partition algorithm for Max-Cut, we shall show that for every $1 \leq i \leq \ell$, and for a fixed choice and partitions of $U^1, \dots, U^{i-1}$, with probability at least $1 - \delta/2\ell$ over the choice of U^i , there exists a partition (U_1^i, U_2^i) of U^i such that

$$\mu(H_1^i, H_2^i) \geq \mu(H_1^{i-1}, H_2^{i-1}) - \frac{3\epsilon}{4\ell}. \quad (24)$$

By induction on i we have that the ℓ th hybrid partition (which is necessarily a bisection since $\sum_i \hat{\beta}^i = 1/2$) has a cut whose size is at most $(3\epsilon/4)N^2$ smaller than the 0th hybrid partition.

Let us define a *good* sample set U^i and a *representative partition* (U_1^i, U_2^i) , similarly to the way they were defined in Lemma 8.1.2.3.1 except that here we make the stronger quantitative requirement that for all but $\epsilon/32$ of the vertices v in V^i

$$\text{For each } j \in \{1, 2\}, \quad \frac{|\Gamma(v) \cap U_j^i|}{t} = \frac{|\Gamma(v) \cap W_j^{i-1}|}{N} \pm \frac{\epsilon}{64}. \quad (25)$$

As was shown in the proof of Lemma 8.1.2.3.1 for our choice of $t = |U^i|$, with probability at least $1 - \delta/2$, $U^1, \dots, U^\ell$ are good with respect to the respective partitions. Assume from now on that $U^1, \dots, U^\ell$ are good, and for each i let $\widehat{ub}(v)$ be defined with respect to the representative partition of U^i .

For a fixed i , we shall bound the difference between the size of the cut determined by the i th hybrid partition and that determined by the $(i-1)$ hybrid partition via two auxiliary hybrid partitions. Let $v_1, \dots, v_L$ be the ordering of the vertices in V^i according to their (correct) unbalance value $ub(v)$ with respect to (W_1^{i-1}, W_2^{i-1}) . Consider first the “ideal” partition $(V_{1,\text{id}}^i, V_{2,\text{id}}^i)$ of V^i , where $V_{1,\text{id}}^i \stackrel{\text{def}}{=} \{v_1, \dots, v_{\beta^i L}\}$, and $V_{2,\text{id}}^i \stackrel{\text{def}}{=} V^i \setminus V_{1,\text{id}}^i$. (Namely, this partition uses both the correct unbalance values and the correct β^i). Let $(H_{1,\text{id}}^i, H_{2,\text{id}}^i)$ be the corresponding ideal hybrid partition (namely, $H_{j,\text{id}}^i \stackrel{\text{def}}{=} W_j^{i-1} \cup V_{j,\text{id}}^i$). As was noted previously, the size of the cut determined by this ideal hybrid partition is at most $|V^i|^2 = (\epsilon/4\ell)N^2$ smaller than the cut determined by the previous hybrid partition (H_1^{i-1}, H_2^{i-1}) .

Next consider the following “almost-ideal” partition $(V_{1,\text{id}'}^i, V_{2,\text{id}'}^i)$ of V^i , where $V_{1,\text{id}'}^i \stackrel{\text{def}}{=} \{v_1, \dots, v_{\lfloor \hat{\beta}^i L \rfloor}\}$, and $V_{2,\text{id}'}^i \stackrel{\text{def}}{=} V^i \setminus V_{1,\text{id}'}^i$. Let $(H_{1,\text{id}'}^i, H_{2,\text{id}'}^i)$ be the corresponding almost-ideal hybrid partition (namely, $H_{j,\text{id}'}^i \stackrel{\text{def}}{=} W_j^{i-1} \cup V_{j,\text{id}'}^i$). Since $|\hat{\beta}^i - \beta^i| \leq \epsilon/16$, then the only difference between $(H_{1,\text{id}'}^i, H_{2,\text{id}'}^i)$ and $(H_{1,\text{id}}^i, H_{2,\text{id}}^i)$ is the placement of at most $(\epsilon/16)L$ vertices. By the above, and applying what we know about the ideal hybrid partition we have,

$$\mu(H_{1,\text{id}'}^i, H_{2,\text{id}'}^i) \geq \mu(H_{1,\text{id}}^i, H_{2,\text{id}}^i) - \frac{\epsilon}{16} L \cdot 2N \geq \mu(H_1^{i-1}, H_2^{i-1}) - \frac{3\epsilon}{8\ell} N^2. \quad (26)$$

In what follows, we bound the difference between the size of the cut determined by the i th hybrid partition (in which V^i is partitioned by the algorithm), and the size of the cut determined by the almost-ideal hybrid partition. This difference is due to the algorithm’s use of approximate unbalance values (i.e., the values $\widehat{ub}(\cdot)$ used by the algorithm induce a different order on V^i than the “correct” order used in the almost-ideal partition). Let Y^i be the set of *misplaced* vertices in V^i which are put on a different side in (H_1^i, H_2^i) than in $(H_{1,\text{id}'}^i, H_{2,\text{id}'}^i)$. Namely, $Y^i \stackrel{\text{def}}{=} (V_{1,\text{id}'}^i \cap V_2^i) \cup (V_{2,\text{id}'}^i \cap V_1^i)$.

CLAIM 8.2.1.2. *There exists a value $y \in [-1, 1]$ such that for all but at most $(\epsilon/16)L$ of the vertices v in Y^i , $ub(v) = y \pm (\epsilon/16)$.*

We prove the claim momentarily, but first derive a bound on $\mu(H_{1,\text{id}'}^i, H_{2,\text{id}'}^i) - \mu(H_1^i, H_2^i)$ based on the claim. By definition of the almost-ideal hybrid partition, we know that in the actual algorithm we put exactly the same number of vertices from V^i on each side of the partition as in the almost-ideal partition. It follows that the number of misplaced vertices on each side of the partition is the same, and we can pair the misplaced vertices and view these pairs as having switched sides. Whenever we switch sides between pairs of vertices whose unbalance value differs by at most $\epsilon/8$, the decrease in the number of cut edges between these two vertices and vertices in W^{i-1} is at most $(\epsilon/4)N$. The contribution of all such pairs is at most $(L/2) \cdot (\epsilon/4)N = (\epsilon/8\ell)N^2$. The number of cut edges between W^{i-1} and the at most $(\epsilon/16)L$ vertices in Y^i which differ significantly in their unbalance value from the rest, decreases by at most $(\epsilon/16)L \cdot 2N = (\epsilon/8\ell)N^2$. Thus,

$$\mu(H_1^i, H_2^i) \geq \mu(H_{1,\text{id}'}^i, H_{2,\text{id}'}^i) - 2 \cdot \frac{\epsilon}{8\ell} N^2. \quad (27)$$

Combining Eq. (26) and Eq. (27), we have

$$\begin{aligned} \mu(H_1^i, H_2^i) &\geq \mu(H_1^{i-1}, H_2^{i-1}) - \frac{3\epsilon}{8\ell} N^2 - \frac{\epsilon}{4\ell} N^2 \\ &= \mu(H_1^{i-1}, H_2^{i-1}) - \frac{5\epsilon}{8\ell} N^2 \end{aligned} \quad (28)$$

and the lemma follows. $\square$

PROOF OF CLAIM 8.2.1.2. Consider a grouping of the vertices in V^i into $2/\epsilon'$ *unbalance bins* according to their (correct) unbalance value, where $\epsilon' = \epsilon/32$. For $k = -(1/\epsilon'), \dots, (1/\epsilon') - 1$, the k th bin, denoted B_k , is defined as follows:

$$B_k \stackrel{\text{def}}{=} \{v \in V^i : \text{ub}(v) \in [k \cdot \epsilon' N, (k+1) \cdot \epsilon' N)\}.$$

Let g be the index of the bin which $v_{\lfloor \beta^i L \rfloor}$ belongs to. By definition of the bins, all vertices in B_g have approximately the same unbalance value. Since we only have approximations of the unbalance values we also group the vertices according to their approximated unbalance values. Namely, for $k = -(1/\epsilon'), \dots, (1/\epsilon') - 1$,

$$\hat{B}_k \stackrel{\text{def}}{=} \{v \in V^i : \hat{\text{ub}}(v) \in [k \cdot \epsilon' N, (k+1) \cdot \epsilon' N)\}.$$

By our assumption on the representativeness of (U_1^i, U_2^i) , at most $\epsilon/32$ of the vertices in V^i belong to a bin $\hat{B}_k$, whose index differs by more than 1 from their correct bin B_k (and vertices in the same, or in neighboring bins have approximately the same unbalance value).

Let $\hat{v}_1, \dots, \hat{v}_L$ be an ordering of the vertices in V^i according to their approximate unbalance value $\hat{\text{ub}}(v)$, and let $\hat{g}$ be the index of the bin which $\hat{v}_{\beta^i L}$ belongs to. We consider two cases.

Bisection Approximation Algorithm

- (1) As Step (1) of Figure 6.
- (2) Uniformly choose a set $S = \{s_1, \dots, s_m\}$ of size $m = \Theta\left(\frac{\ell^2 t + \ell^2 \cdot \log(1/(\epsilon\delta))}{\epsilon^2}\right)$. For $1 \leq i \leq \ell$, let $S^i \stackrel{\text{def}}{=} V^i \cap S$.
- (3) Analogously to Step (2) of Figure 6, for each of the sequence of partitions Π and for each of the ℓ -tuples $\hat{\beta} = (\hat{\beta}^1, \dots, \hat{\beta}^\ell)$, (where each $\hat{\beta}^i \in [0, 1]$ is an integer multiple of $\frac{\epsilon}{16\ell}$, and $\sum_i \hat{\beta}^i = 1/2$), construct a partition $(S_1^{\Pi, \hat{\beta}}, S_2^{\Pi, \hat{\beta}})$ of S . Specifically, in the i^{th} iteration of Substep (a), S_1^i is assigned the $\lfloor \hat{\beta}^i |S^i| \rfloor$ vertices with the biggest $\widehat{\text{ub}}(\cdot)$ value.
- (4) For each resulting partition $(S_1^{\Pi, \hat{\beta}}, S_2^{\Pi, \hat{\beta}})$, compute the fraction of cut edges between pairs of vertices (s_{2k-1}, s_{2k}) . More precisely, define

$$\hat{\mu}(S_1^{\Pi, \hat{\beta}}, S_2^{\Pi, \hat{\beta}}) \stackrel{\text{def}}{=} \frac{\left| \left\{ (s_{2k-1}, s_{2k}) \in E(S_1^{\Pi, \hat{\beta}}, S_2^{\Pi, \hat{\beta}}) \cup E(S_2^{\Pi, \hat{\beta}}, S_1^{\Pi, \hat{\beta}}) \right\} \right|}{m/2}.$$

Let $(\tilde{S}_1^{\Pi, \hat{\beta}}, \tilde{S}_2^{\Pi, \hat{\beta}})$ be a partition of S for which this fraction is maximized, and output $\hat{\mu}^{(1/2)}(G) = \hat{\mu}(\tilde{S}_1^{\Pi, \hat{\beta}}, \tilde{S}_2^{\Pi, \hat{\beta}})$.

FIG. 7. Bisection Approximation Algorithm.

Case 1. $\hat{g} = g \pm 1$. In this case all but at most $2 \cdot (\epsilon/32)L$ of the misplaced vertices (belong to the bins B_g and $B_{\hat{g}}$) and have unbalance value ranging between $(g - i) \cdot \epsilon' N = g \cdot \epsilon' N - (\epsilon/32)N$ and $(g + 1) \cdot \epsilon' \cdot N + (\epsilon/32)N = g \cdot \epsilon' \cdot N + (\epsilon/16)N$, as required.

Case 2. $\hat{g} \geq g + 2$ (the case $\hat{g} \leq g - 2$ is analogous). In such a case, necessarily, all but $(\epsilon/32)L$ of the vertices $v_1, \dots, v_{\lfloor \hat{\beta}^i L \rfloor}$ (which belong to bins $B_1, \dots, B_{\hat{g}}$), are put on side 1 (as they should). But since we put exactly $\lfloor \hat{\beta}^i L \rfloor$ vertices on side 1, the total number of misplaced vertices is bounded by $2 \cdot (\epsilon/32)L$, and the claim follows. $\square$

Applying Lemma 8.2.1.1 to a maximum cut of G , we get:

COROLLARY 8.2.1.3. *With probability at least $1 - \delta/2$ over the choice of $\bar{U}$, there exists a sequence of partitions Π of $\bar{U}$, and an ℓ -tuple $\hat{\beta}$ such that*

$$\mu(V_1^{\Pi, \hat{\beta}}, V_2^{\Pi, \hat{\beta}}) \geq \mu^{(1/2)}(G) - \frac{3}{4} \epsilon.$$

Thus, with probability at least $1 - (\delta/2)$, the Graph Bisection Algorithm (described in Figure 6) outputs a bisection $(V_1^{\Pi, \hat{\beta}}, V_2^{\Pi, \hat{\beta}})$ such that $\mu(V_1^{\Pi, \hat{\beta}}, V_2^{\Pi, \hat{\beta}}) \geq \mu^{(1/2)}(G) - (3/4)\epsilon$.

8.2.2. The Bisection Approximation Algorithm. Similarly to the Max-Cut Approximation Algorithm, the Bisection Approximation Algorithm (described in Figure 7) performs the same steps as the algorithm described in Figure 6, but does so only on a small sample S . The analysis of this Bisection Approximation Algorithm, given the correctness of the Graph Bisection Algorithm, is similar to

that of the Max-Cut Approximation Algorithm (Lemma 8.1.3.1) except for the following detail. Here we need to take into account that it is *not* necessarily the case that for a given $\bar{U}$, a sequence of partitions Π of $\bar{U}$ and $\hat{\beta}$, for each $s \in S$, $s \in S_j^{\Pi, \hat{\beta}}$, if and only if $s \in V_j^{\Pi, \hat{\beta}}$. This unfortunate phenomena is due to the possibility that for some vertices $v \in S^i$, vertex v appears before (respectively, after) the $\hat{\beta}^i N$ th vertex in the ordering of V^i , but after (respectively, before) the $\hat{\beta}^i m$ th vertex in the ordering of S^i . To deal with this we prove the following lemma using arguments analogous to Lemma 7.3.1 (which deals with an analogous phenomena in the analysis of the ρ -Clique Tester).

LEMMA 8.2.2.1. *For a fixed $\bar{U}$, Π , and $\hat{\beta}$, let $(V_1, V_2) = (V_1^{\Pi, \hat{\beta}}, V_2^{\Pi, \hat{\beta}})$ be as defined in the Graph Bisection Algorithm. Let S be a uniformly chosen sample of size*

$$m = \Omega\left(\frac{\ell(\ell t + \log(1/(\epsilon\delta)))}{\epsilon^2}\right)$$

such that that for each i , $|S^i| \geq m/(2\ell)$, and let $(S_1, S_2) = (S_1^{\Pi, \hat{\beta}}, S_2^{\Pi, \hat{\beta}})$ and $\hat{\mu}(S_1, S_2)$ be as defined in the Bisection Approximation Algorithm. Then

$$\Pr_S \left[\left| \hat{\mu}(S_1, S_2) - \mu(V_1, V_2) \right| > \frac{\epsilon}{4} \right] < \frac{\delta}{2} \cdot 2^{-\ell \cdot t} \cdot \left(\frac{\epsilon}{16} \right)^\ell.$$

PROOF. Let $\delta' \stackrel{\text{def}}{=} \delta/2 \cdot 2^{-\ell \cdot t} \cdot (\epsilon/16)^\ell$. For each $i \in \{1, \dots, \ell\}$, let V_1^i and V_2^i be as defined in the Graph Bisection Algorithm, and let S_1^i and S_2^i be as defined in the Bisection Approximation Algorithm (for the fixed $\bar{U}$, Π and $\hat{\beta}$ considered in the lemma). Let $\epsilon_1 \stackrel{\text{def}}{=} \epsilon/40$, and let $\bar{V}_1^i$ be the first $\lfloor (\hat{\beta}^i - \epsilon_1)L \rfloor$ vertices in V^i , and $\bar{V}_2^i$ the last $\lceil (1 - \hat{\beta}^i - \epsilon_1)L \rceil$ vertices in V^i . For each i , let $m_i \stackrel{\text{def}}{=} |S^i|$, where by the lemma's hypothesis, $m_i \geq m/(2\ell)$. By Claim 7.3.2, for each i ,

$$\Pr_S \left[\left| \frac{|S^i \cap \bar{V}_1^i|}{m_i} - (\hat{\beta}^i - \epsilon_1) \right| > \epsilon_1 \right] < \frac{\delta'}{4\ell}$$

and similarly,

$$\Pr_S \left[\left| \frac{|S^i \cap \bar{V}_2^i|}{m_i} - (1 - \hat{\beta}^i - \epsilon_1) \right| > \epsilon_1 \right] < \frac{\delta'}{4\ell}.$$

Note that the effects of rounding quantities such as $\hat{\beta}^i \cdot L$ (and $\hat{\beta}^i \cdot m_i$) are negligible since they effect the placement of at most one vertex from each V^i (respectively, S^i), and since $1/m_i$ (and certainly $1/L$) are much smaller than ϵ_1 we may ignore these effects. Thus, putting aside an error probability of $\delta/2$, assume from now on that for each i :

- (1) $|S^i \cap \bar{V}_1^i| \leq \hat{\beta}^i \cdot m_i$, and $|S^i \cap \bar{V}_2^i| \leq (1 - \hat{\beta}^i) \cdot m_i$ from which it follows (by definition of S_1^i and S_2^i) that for each $j \in \{1, 2\}$, $S_j^i \cap \bar{V}_j^i = S^i \cap \bar{V}_j^i$;
- (2) $|S^i \cap \bar{V}_1^i| \geq (\hat{\beta}^i - 2\epsilon_1) \cdot m_i$, and $|S^i \cap \bar{V}_2^i| \geq (1 - \hat{\beta}^i - 2\epsilon_1) \cdot m_i$. Combining this with Item (1) it follows that $|S_1^i \cap \bar{V}_1^i| \geq (\hat{\beta}^i - 2\epsilon_1) \cdot m_i$ and $|S_2^i \cap \bar{V}_2^i| \geq (1 - \hat{\beta}^i - 2\epsilon_1) \cdot m_i$.

Let $\bar{V}_1$ be the union of the $\bar{V}_1^i$'s and let $\bar{V}_2$ be the union of the $\bar{V}_2^i$'s. Then

$$\begin{aligned} \{(s_{2k-1}, s_{2k}) \in E(S_1, S_2)\} &= \{(s_{2k-1}, s_{2k}) \in E(S_1 \cap \bar{V}_1, S_2 \cap \bar{V}_2)\} \\ &\cup \{(s_{2k-1}, s_{2k}) \in (E(S_1 \setminus \bar{V}_1, S_2) \cup E(S_1, S_2 \setminus \bar{V}_2))\} \end{aligned} \quad (29)$$

By Item (1) above and an additive Chernoff bound we get that

$$\Pr_S \left[\left| \frac{|\{(s_{2k-1}, s_{2k}) \in E(S_1 \cap \bar{V}_1, S_2 \cap \bar{V}_2)\}|}{m/2} - \frac{|E(\bar{V}_1, \bar{V}_2)|}{N^2} \right| > \frac{\epsilon}{10} \right] < \frac{\delta'}{2}. \quad (30)$$

By definition of $\bar{V}_1$ and $\bar{V}_2$, we have that

$$\frac{|E(V_1, V_2)|}{N^2} - \frac{|E(\bar{V}_1, \bar{V}_2)|}{N^2} \leq 2\epsilon_1 = \frac{\epsilon}{20}. \quad (31)$$

By Item (2) above, we know that for each $i \in \{1, \dots, \ell\}$ and $j \in \{1, 2\}$, $|S_j^i \setminus \bar{V}_j^i| \leq 2\epsilon_1 \cdot m_i$, and so

$$\begin{aligned} &\frac{|\{(s_{2k-1}, s_{2k}) \in (E(S_1 \setminus \bar{V}_1, S_2) \cup E(S_1, S_2 \setminus \bar{V}_2))\}|}{m/2} \\ &\leq \frac{|\{k : s_{2k-1} \in S_1 \setminus \bar{V}_1 \text{ or } s_{2k} \in S_2 \setminus \bar{V}_2\}|}{m/2} \leq 4\epsilon_1 = \frac{\epsilon}{10} \end{aligned} \quad (32)$$

Summing up the probabilities of errors and combining the bounds of Eqs. (30)–(32), the lemma follows. $\square$

Part (1) of Theorem 8.2.1 follows by combining Corollary 8.2.1.3 and Lemma 8.2.2.1. We only need to observe that: (1) by a multiplicative Chernoff Bound, with very high probability in fact $|S^i| \geq m/(2\ell)$ for each i ; and (2) the number of sequences of partitions Π is $2^{\ell t}$ and the number of settings of $\hat{\beta}$ is less than $(16/\epsilon)^\ell$.

8.2.3. The Improved Graph-Bisection Algorithm. Similarly to the improved graph-partitioning algorithm for Max-Cut, the improved graph-bisection algorithm (whose running time is as stated in Theorem 8.2.1, Part (2)) starts by invoking the Bisection approximation algorithm of Figure 7, and recording the sequence of sets $\bar{U}$ uniformly selected in Step (1), and the sequence of partitions $\bar{\Pi}$ and the ℓ -tuple $\tilde{\beta}$, selected in Step (4). Using these specific $\bar{\Pi}$ and $\tilde{\beta}$, the algorithm executes a single iteration of Step (2) of the Graph Bisection Algorithm in Figure 6. Since it does not check the resulting partition, it saves a multiplicative factor of N in its running time. More precisely, it has running time $O(t \cdot N) = O(\log(1/(\epsilon\delta)/e^2)) \cdot N$ (on top of the running time of the approximation algorithm).

As for its correctness, by Corollary 8.2.1.3, we have that with probability at least $1 - \delta/2$ over the choice of $\bar{U}$, there exists Π and $\hat{\beta}$ such that $\mu(V_1^{\Pi, \hat{\beta}}, V_2^{\Pi, \hat{\beta}}) \geq \mu(G) - (3/4)\epsilon$. From Lemma 8.2.2.1, we have that for a fixed $\bar{U}$, with probability at least $1 - \delta/2$ over the choice of S , $\hat{\mu}(S_1^{\Pi, \hat{\beta}}, S_2^{\Pi, \hat{\beta}})$ is within $\epsilon/4$ from

$\mu(V_1^{\Pi, \hat{\beta}}, V_2^{\Pi, \hat{\beta}})$ for every sequence of partitions Π and $\bar{U}$, and every $\hat{\beta}$. It follows that with probability at least $1 - \delta$ over the choice of $\bar{U}$ and S , the recorded $\hat{\Pi}$ and $\hat{\beta}$ (from the Bisection Approximation Algorithm) are such that $\mu(V_1^{\hat{\Pi}, \hat{\beta}}, V_2^{\hat{\Pi}, \hat{\beta}}) \geq \mu(G) - \epsilon$, as required.

8.2.4. Variations

8.2.4.1. BISECTION MINIMIZATION. An easy modification suffices for finding (respectively, approximating the size of) a nearly minimum bisection rather than a nearly maximum one. In each iteration of the algorithm(s), instead of placing in side 1 the first $\hat{\beta}^i$ vertices in decreasing order of (approximate) unbalance, we would do the opposite. Namely, since we would like to minimize the size of the cut, we try and put vertices on the size opposite the *minority* of their neighbors. Although we might not be able to do so for all vertices (since we are restricted to constructing a bisection), analogously to the maximization problem, there exists one side in which all vertices have a smaller (i.e., more negative) unbalance value than all those on the other side. Thus, in the i th iteration, we order all vertices in V^i (or S^i in the approximation algorithm) according to *increasing* unbalance value $\text{ub}(v)$ (where $\text{ub}(v)$ is as defined in the maximization algorithms), and put the first $\hat{\beta}^i$ vertices on side 1 and the rest on side 2.

8.2.4.2. OTHER RESTRICTIONS ON THE PARTITION. We can also easily generalize the algorithms to construct (respectively, approximate the size of) partitions with other predetermined proportion of vertices on each side. A key observation is that the main steps of our algorithms are oblivious of the bisection requirement other than in asking that $\sum_i \hat{\beta}^i = 1/2$. In fact, the main steps can produce partitions with maximum (respectively, minimum) number of crossing edges per each proportion of vertices on each side (up to some resolution). Therefore, all we need to do is modify the restriction on the sum of the $\hat{\beta}^i$'s to allow either a different fixed proportion or a range of proportions. Thus, for example, we can approximate quantities such as $\text{opt}_{|V_1|=N/3} \{\mu(V_1, V(G) \setminus V_1)\}$, or $\text{opt}_{N/3 \leq |V_1| \leq 2N/3} \{\mu(V_1, V(G) \setminus V_1)\}$, where $\text{opt} \in \{\max, \min\}$.

Testing algorithms for properties corresponding to the above optimization problems essentially follow from the approximation algorithms. In particular, consider first properties corresponding to minimization problems. When the property is defined as having a cut of density *at most* ρ (subject to certain constraints on the partition defining the cut) then the corresponding class is never empty (in particular, the empty graph belongs to the class). Furthermore, for any α , if a graph G has a partition (V_1, V_2) such that $|V_1| = \alpha N$ and $\mu(V_1, V_2) \leq \rho + \epsilon$, then we can always remove at most ϵN^2 edges to obtain $\mu(V_1, V_2) = \rho$. Therefore, in the case of cut-minimization problems of the type discussed in this subsection, the corresponding testing algorithms follow directly from the approximation algorithms.

The situation is slightly more involved when dealing with maximization problems. For $0 \leq \alpha \leq 1/2$ and a graph G , let $\mu^{(\alpha)}(G)$ denote the maximum edge density among all cuts $(V_1, V(G) \setminus V_1)$ such that $|V_1| = \alpha N$. Let $\mathcal{MC}_\rho^{(\alpha)} \stackrel{\text{def}}{=} \{G : \mu^{(\alpha)}(G) \geq \rho\}$. To test whether a graph G belongs to the class $\mathcal{MC}_\rho^{(\alpha)}$, we first check whether the class is empty, in which case we reject G . Namely, if $2\alpha(1 - \alpha) < \rho$, then the class must be empty, since the maximum number of edges connecting vertices in a set of size αN to vertices in a set of size $(1 - \alpha)N$ is

$2\alpha(1 - \alpha)N^2$. If $2\alpha(1 - \alpha) \geq \rho$ (and so the class is not empty), the testing algorithm follows from the approximation algorithm mentioned above, analogously to the way the ρ -Bisection algorithm follows from the Bisection approximation algorithm (see discussion following Theorem 8.2.1).

When the sizes of sides of the partition are not required to be fixed but rather are allowed to be within a certain range, the testing algorithm is a little less straightforward. For $0 \leq \alpha_1 \leq \alpha_2 \leq 1/2$ and a graph G , let $\mu^{(\alpha_1, \alpha_2)}(G)$ denote the maximum edge density among all cuts $(V_1, V(G) \setminus V_1)$ such that $\alpha_1 N \leq |V_1| \leq \alpha_2 N$. Let $\mathcal{M}_{\rho}^{(\alpha_1, \alpha_2)} \stackrel{\text{def}}{=} \{G : \mu^{(\alpha_1, \alpha_2)}(G) \geq \rho\}$. Note that the class of graphs having a cut of size at least ρ (i.e., $\mathcal{M}_{\rho} = \mathcal{M}_{\rho}^{(0, 1/2)}$), which was considered in Section 8.1, is indeed a special case. If $2\alpha_2(1 - \alpha_2) < \rho$ then the class $\mathcal{M}_{\rho}^{(\alpha_1, \alpha_2)}$ is empty. Otherwise, let $\alpha'_1 \geq \alpha_1$ be the minimum value such that $2\alpha'_1(1 - \alpha'_1) \geq \rho$. We run the algorithm for approximating $\mu^{(\alpha_1, \alpha_2)}(G)$ as described above with approximation parameter $\epsilon/2$ (where ϵ is the distance parameter of the testing algorithm), while requiring that $\alpha'_1 \leq \Sigma_i \hat{\beta}^i \leq \alpha_2$. We accept the graph if and only if the approximate value obtained is at least $\rho - \epsilon/2$ (as in the proof of Corollary 8.1.2).

The following example best illustrates why we introduce the restriction that $\Sigma_i \hat{\beta}^i \geq \alpha'_1$ instead of just using $\Sigma_i \hat{\beta}^i \geq \alpha_1$. Consider the case in which $\alpha_1 = 0$, $\alpha_2 = 1/2$, and $\rho = 1/2$ (and so we are simply asking whether the graph G has any cut of density $1/2$). Suppose that G is a complete bipartite graph between a set of vertices V_1 , and a set of vertices V_2 , such that $|V_1| = (1/2 - \sqrt{\epsilon/2})$ and $|V_2| = (1/2 + \sqrt{\epsilon/2})$. Thus, $\mu^{(0, 1/2)}(G) = \rho - \epsilon$, and so with fairly high probability the approximation algorithm would output an approximate value that is close to ρ . However, G is $\sqrt{\epsilon}$ -far from having the desired property, and should be rejected. Note that the difficulty is not with the approximation algorithm but rather with the relation between approximation and testing. In this case, setting $\alpha'_1 = \alpha_2 = 1/2$ (as suggested above, since $2\alpha'_1 \cdot (1 - \alpha'_1) \geq 1/2$ implies $\alpha'_1 = 1/2$), we restrict our algorithm to consider only partitions for which $\Sigma_i \hat{\beta}^i \geq \alpha'_1$, and so the algorithm will detect that the graph should be rejected.

9. The General Partition Problem

The following framework of a general partition problem generalizes all properties considered in previous sections. In particular, it captures any graph property that requires the existence of partitions satisfying certain fixed density constraints. These constraints may refer both to the number of vertices in each component of the partition and to the number of edges between each pair of components.

Let $\Phi \stackrel{\text{def}}{=} \{\rho_j^{\text{LB}}, \rho_j^{\text{UB}}\}_{j=1}^k \cup \{\varrho_{j,j'}^{\text{LB}}, \varrho_{j,j'}^{\text{UB}}\}_{j,j'=1}^k$ be a set of non-negative parameters so that $\rho_j^{\text{LB}} \leq \rho_j^{\text{UB}}$ ($\forall j$) and $\varrho_{j,j'}^{\text{LB}} \leq \varrho_{j,j'}^{\text{UB}}$ ($\forall j, j'$). (LB stands for Lower Bound, and UB stands for Upper Bound.) Let $\mathcal{GP}_{\Phi}$ be the class of graphs that have a k -way partition $(V_1, \dots, V_k)$ with the following conditions being satisfied.

$$\forall j \quad \rho_j^{\text{LB}} \cdot N \leq |V_j| \leq \rho_j^{\text{UB}} \cdot N \quad (33)$$

and

$$\forall j, j' \quad \varrho_{j,j'}^{\text{LB}} \cdot N^2 \leq |E(V_j, V_{j'})| \leq \varrho_{j,j'}^{\text{UB}} \cdot N^2, \quad (34)$$

where recall that $E(V_j, V_{j'})$ is the set of edges between vertices in V_j and vertices in $V_{j'}$ (where we include edges going in both directions). That is, Eq. (33) places lower and upper bounds on the relative sizes of the various components of the partition; whereas Eq. (34) imposes lower and upper bounds on the density of edges among the various pairs of components.

Remark 9.1 (A Tedious One.) To avoid integrability problems, we consider generalized (fractional) k -way partitions in which up to $k - 1$ vertices may be split among several parts. Had we not followed this convention, the set of N -vertex graphs in $\mathcal{GP}_\Phi$ could be empty for some values of N and non-empty for others. For example, if $\rho_1^{\text{LB}} = \rho_1^{\text{UB}} = 1/3$ then only graphs with $3M$ vertices may be in $\mathcal{GP}_\Phi$. In such a case, the tester must reject any graph with $3M + 1$ vertices (as the class of graphs with $3M + 1$ vertices having the property defined by the parameters is empty), whereas it must accept some $3M$ -vertex graphs. Consequently, such a tester must count the number of vertices in the graph. These integrability problems have nothing to do with the combinatorial structure that we wish to investigate and thus we avoid them by taking the above somewhat unnatural convention.

In this section, we describe a testing algorithm for the class $\mathcal{GP}_\Phi$ (for any given set of parameters $\Phi = \{\rho_j^{\text{UB}}, \rho_j^{\text{LB}}\} \cup \{\varrho_{j,j'}^{\text{UB}}, \varrho_{j,j'}^{\text{LB}}\}$). Similarly to the testing algorithms described in Sections 7 and 8, the testing algorithm of this section is based on a randomized *partitioning* algorithm for the related partition problem. Namely, given a graph G , a set of parameters Φ , an approximation parameter ϵ and a confidence parameter δ , so that G has a k -way partition which obeys Eqs. (33) and (34), the partitioning algorithm constructs a partition $(V_1, \dots, V_k)$ of G for which the following hold with probability at least $1 - \delta$:

$$\forall j, \quad (\rho_j^{\text{LB}} - \epsilon) \cdot N \leq |V_j| \leq (\rho_j^{\text{UB}} + \epsilon) \cdot N, \quad (35)$$

and

$$\forall j, j', \quad (\varrho_{j,j'}^{\text{LB}} - \epsilon) \cdot N^2 \leq |E(V_j, V_{j'})| \leq (\varrho_{j,j'}^{\text{UB}} + \epsilon) \cdot N^2, \quad (36)$$

A partition obeying (35) and (36) is called an ϵ -approximation for the partitioning problem defined by the set of parameters Φ .

As already indicated in the special case of ρ -Cut, and in the generalization of the Bisections testing algorithm (Section 8.2.4), the relationship between having an ϵ -approximation for the general partitioning problem, and being ϵ -close to the class of graphs having the property is not completely straightforward. In particular, a graph may have a partition that is an ϵ -approximation for the partitioning problem defined by the set of parameters Φ , but is $\Omega(\sqrt{\epsilon})$ -far from the class. $\mathcal{GP}_\Phi$. We shall deal with this difficulty when designing the testing algorithm. Jumping ahead we mentioned that instead of checking whether the tested graph has a partition that is an ϵ -approximation for the partitioning problem (or an $f(\epsilon, k)$ -approximation, for some function f of ϵ and k), we directly check whether the graph is ϵ -close to $\mathcal{GP}_\Phi$.

As stated above, all properties considered in previous sections can be cast as special cases of the general partition problem. For example, k -Colorability is expressed by setting $\varrho_{j,j}^{\text{UB}} = 0$ for every j (and placing no other constraints which means setting $\rho_j^{\text{LB}} = 0$, $\rho_j^{\text{UB}} = 1$, and similarly setting the $\varrho_{j,j'}^{\text{UB}}$'s and $\varrho_{j,j'}^{\text{LB}}$'s for

$j' \neq j$). In case we are interested in maximizing or minimizing a parameter (e.g., maximizing $E(V_1, V_2)$ in the case of Max-Cut) we can simply run the general partitioning (respectively, testing) algorithm on all values of this parameter which are multiples of ϵ , and find the maximum/minimum value attainable.²⁴ However, as can be seen from our theorems below (and the table in Figure 1), this generality has a price: The query complexity and running times of our algorithms (for the general partition problem) are quite large. More efficient algorithms for specific problems such as k -Coloring, Max-Clique, and Max-Cut, were presented in previous sections. Although we cannot exploit the problem-specific properties as done in the previous sections, we nonetheless apply some of the ideas used in the above algorithms.

THEOREM 9.2. *There exists an algorithm $\mathcal{A}$ such that for every given set of parameters Φ , algorithm $\mathcal{A}$ is a property testing algorithm for the class $\mathcal{GP}_\Phi$ with query complexity and running time*

$$\log^2\left(\frac{1}{\epsilon\delta}\right) \cdot \left(\frac{O(k^2)}{\epsilon}\right)^{2k+8} \quad \text{and} \quad \exp\left(\log\left(\frac{1}{\epsilon\delta}\right) \cdot \left(\frac{O(k^2)}{\epsilon}\right)^{k+1}\right),$$

respectively.

We note that in the running time of the algorithm, we ignore a factor that is polynomial in the length of the description of Φ , as we view this length as a fixed constant and not a parameter to the problem. Furthermore, in any reasonable application, it is much smaller than all other factors. As in previous sections, we also obtain an analogous graph partitioning algorithm.

THEOREM 9.3. *There exists a graph partitioning algorithm that on input Φ , ϵ , and δ , and oracle access to a graph G , runs in time*

$$\exp\left(\log\left(\frac{1}{\epsilon\delta}\right) \cdot \left(\frac{O(1)}{\epsilon}\right)^{k+1}\right) + O\left(\frac{\log(k/(\epsilon\delta))}{\epsilon^2}\right) \cdot N$$

and if G has a k -way partition satisfying Eqs. (33) and (34), then with probability at least $1 - \delta$ the graph partitioning algorithm outputs a partition which satisfies Eqs. (35) and (36).

We start by describing a less efficient graph partitioning algorithm with running time $\exp(\log(1/(\epsilon\delta)) \cdot (O(1)/\epsilon)^{k+1}) \cdot N^2$. Based on this algorithm, we shall obtain the tester postulated in Theorem 9.2 and finally derive the (more efficient) graph partitioning algorithm (postulated in Theorem 9.3).

9.1. HIGH LEVEL DESCRIPTION OF THE PARTITIONING ALGORITHM. The algorithm is based on the following observation, which generalizes an observation applied in the Bisection algorithm (Section 8.2). Let $\mathcal{H} = (H_1, \dots, H_k)$ be any fixed partition of V . In particular, we may want to consider a partition which

²⁴ Actually, our partitioning algorithm works by producing a set of partitions of the graph vertices and then searching among them for one which is an ϵ -approximation of the partitioning problem. The testing algorithm runs a similar procedure on a sample set of vertices. The procedure for producing these partitions depends on k , ϵ , and δ , but not on the particular set of parameters Φ , and, therefore, we do not actually need to run the algorithm more than once.

obeys Eqs. (33) and (34). Let X be a set of vertices of small size (i.e., of size $O(\epsilon N)$) and suppose that all but $O(\epsilon|X|)$ of the vertices in X have approximately the same (i.e., $\pm O(\epsilon N)$) number of neighbors in each individual $H_j \setminus X$. Namely, for each $j \in \{1, \dots, k\}$, there exists a value β_j such that for all but $O(\epsilon|X|)$ of the vertices v in X , we have

$$\frac{|\Gamma(v) \cap (H_j \setminus X)|}{N} = \beta_j \pm O(\epsilon). \quad (37)$$

(Recall that $a = b \pm c$ is a shorthand for $b - c \leq a \leq b + c$.) The observation is that if we arbitrarily redistribute the vertices of X among the k components (i.e., H_j 's) while maintaining the *number of vertices* in each component, then the *number of edges* between every pair of components is approximately maintained.

More precisely, let $(X_1, \dots, X_k)$ be an arbitrary partition of X so that $|X_j| = |X \cap H_j| \pm O(\epsilon|X|)$, and let $H'_j = (H_j \setminus X) \cup X_j$. Let $\mathcal{H}' \stackrel{\text{def}}{=} (H'_1, \dots, H'_k)$. Then by our assumption on the X_j 's, we have $||H_j| - |H'_j|| = O(\epsilon|X|)$ for every j . Furthermore, by our assumption concerning the “neighbor-profile” of vertices in X (Eq. (37)), for every j, j' ,

$$||E(H_j, H_{j'})| - |E(H'_j, H'_{j'})|| = O(|X|^2 + |X| \cdot \epsilon N) = O(|X| \cdot \epsilon N) \quad (38)$$

where the second equality is due to the size of X . The first equality, namely the bound on the difference of the number of edges, is proved in detail in Lemma 9.2.1. The first crucial observation is that edges with both endpoints *not* in X are in the same component in $\mathcal{H}'$ as they were in $\mathcal{H}$, and thus are not effected by the redistribution of X . As for edges with at least one endpoint in X , there are $O(|X|^2)$ edges with both endpoints in X (accounted for in the first term of Eq. (38)), and the changes in the number of edges with exactly one endpoint in X (due to Eq. (37) and $|X_j| = |X \cap H_j| \pm O(\epsilon|X|)$) can be bounded by $O(|X| \cdot \epsilon N)$.

So far, we have dealt with a single set of vertices such that all vertices in the set have approximately the same “neighborhood profile” with respect to a given partition. In general, we wish to handle the case in where vertices have arbitrary neighborhood profile with respect to the partition. The idea is to pack vertices into *clusters* according to their neighborhood profile. Specifically, let $\mathcal{H}$ be as defined above, let Y be any given set of vertices of size $O(\epsilon N)$, and let $\mathcal{W} = (W_1, \dots, W_k)$ be defined by $W_j \stackrel{\text{def}}{=} H_j \setminus Y$ (for each j). We first *cluster* the vertices in Y according to the approximate number of neighbors they have in each W_j . That is, in each (disjoint) cluster all vertices have approximately the same number of neighbors in each W_j . Suppose we now partition the vertices in each cluster X into k parts, $(X_1, \dots, X_k)$ in an arbitrary way so that the number of vertices in each X_j is approximately $|X \cap H_j|$, and add each X_j to W_j , defining a hybrid partition, $\mathcal{H}' = (H'_1, \dots, H'_k)$. Then, by the above discussion, which may be viewed as concerning a single cluster, for every j ,

$$||H_j| - |H'_j|| = \sum_{\text{clusters } X \subseteq Y} O(\epsilon|X|) = O(\epsilon|Y|) = O(\epsilon^2 N) \quad (39)$$

and for every j, j' ,

$$||E(H_j, H_{j'})| - |E(H'_j, H'_{j'})|| = \sum_{\text{clusters } X \subseteq Y} O(|X| \cdot \epsilon N) = O(\epsilon^2 N^2). \quad (40)$$

Similarly to the analysis of the graph partitioning algorithms for Max-Cut and Bisection, we shall use the above observation to define a sequence of $O(1/\epsilon)$ hybrid partitions. It will follow that for the final hybrid, the differences in vertex and edge densities as compared to the initial one (which obeys Eqs. (33) and (34)) is at most ϵ .

9.1.1. The Oracle-Aided Procedure. In view of the above observation, we are almost ready to describe the partitioning algorithm. Similarly to the graph-partitioning algorithms for Max-Cut and Bisection, we first describe a mental experiment in which we assume the algorithm has access to certain oracles (which it actually does not have direct access to). We later show how we can approximately simulate these oracles. The algorithm works in ℓ iterations, using as before a *fixed* partition into ℓ equal-size sets $V^1, \dots, V^\ell$, where $\ell = 4/\epsilon$ and $L = N/\ell$. In the i th iteration, the algorithm partitions the set V^i into $(V^i_1, \dots, V^i_k)$. Let $\mathcal{H}^0 = (H^0_1, \dots, H^0_k)$ be a k -way partition that obeys Eqs. (33) and (34), and for each $i > 0$, let $\mathcal{H}^i = (H^i_1, \dots, H^i_k)$ be the i th *hybrid partition*, where $H^i_j \stackrel{\text{def}}{=} (H^{i-1}_j \setminus V^i) \cup V^i_j$. Let $\mathcal{W}^{i-1} = (W^{i-1}_1, \dots, W^{i-1}_k)$ be the partition induced on $V \setminus V^i$ by $\mathcal{H}^{i-1}$. That is, for each j , $W^{i-1}_j \stackrel{\text{def}}{=} H^{i-1}_j \setminus V^i$. For any given vertex $v \in V^i$ and for every $j \in \{1, \dots, k\}$, let $\gamma^i_j(v) \stackrel{\text{def}}{=} |\Gamma(v) \cap W^{i-1}_j|/N$. The k -tuple $(\gamma^i_1(v), \dots, \gamma^i_k(v))$ is called the *neighborhood profile* of v , referred to in the above discussion.

We assume that the algorithm has an oracle that for each i , given a vertex $v \in V^i$ and $j \in \{1, \dots, k\}$, returns a value $\hat{\gamma}^i_j(v)$ so that for all but $O(\epsilon L)$ of the vertices v in V^i it holds that for every j , $\hat{\gamma}^i_j(v) = \gamma^i_j(v) \pm \epsilon/32$. Using this oracle, the algorithm clusters the vertices in V^i according to the number of neighbors they have in each component of $\mathcal{W}^{i-1}$ as approximated by the oracle: For every possible $\vec{\alpha} = \langle \alpha_1, \dots, \alpha_k \rangle$ where each α_j ranges over all integer multiples of $\epsilon/16$, let

$$V^{i,\vec{\alpha}} \stackrel{\text{def}}{=} \left\{ v \in V^i : \forall j, \alpha_j - \frac{\epsilon}{32} < \hat{\gamma}^i_j(v) \leq \alpha_j + \frac{\epsilon}{32} \right\}.$$

We refer to the $\vec{\alpha}$'s as the *cluster names*, since each $\vec{\alpha}$ uniquely defines a different cluster of V^i .

Assume further that the algorithm also had access to an oracle which for every cluster $V^{i,\vec{\alpha}}$ and for each j , returns an approximation, up to an error of $\epsilon/16$, of the fraction of vertices in $V^{i,\vec{\alpha}}$ which belong to H^{i-1}_j . Let this approximate fraction be denoted $\beta^{i,\vec{\alpha}}_j$. That is,

$$\beta^{i,\vec{\alpha}}_j = \frac{|V^{i,\vec{\alpha}} \cap H^{i-1}_j|}{|V^{i,\vec{\alpha}}|} \pm \frac{\epsilon}{16}.$$

We refer to $\langle \beta^{i,\vec{\alpha}}_1, \dots, \beta^{i,\vec{\alpha}}_k \rangle$ as the *quantitative partition* of $V^{i,\vec{\alpha}}$, since it only determines how *many* vertices from $V^{i,\vec{\alpha}}$ should be in each component of the partition (and does not specify *which* vertices should be in each component).

However, by our observation concerning redistribution of vertices belonging to the same cluster, the quantitative partition of $V^{i,\tilde{\alpha}}$, is all that matters, and we may as well partition $V^{i,\tilde{\alpha}}$ in an arbitrary way as long as the quantitative partition is satisfied. Let $(V_1^{i,\tilde{\alpha}}, \dots, V_k^{i,\tilde{\alpha}})$ be such a partition; that is, $\lfloor \beta_j^{i,\tilde{\alpha}} |V^{i,\tilde{\alpha}}| \rfloor \leq |V_j^{i,\tilde{\alpha}}| \leq \lceil \beta_j^{i,\tilde{\alpha}} |V^{i,\tilde{\alpha}}| \rceil$, for every j .

Let $(V_1^i, \dots, V_k^i)$ be defined by $V_j^i \stackrel{\text{def}}{=} \cup_{\tilde{\alpha}} V_j^{i,\tilde{\alpha}}$, for each j . By our previous discussion we know that for each i , the changes in vertex and edge densities between the $i - 1$ and i th hybrid partitions is $O(\epsilon^2)$ (see Eqs. (39) and (40)). Combining this with our hypothesis concerning $\mathcal{H}^0$ (i.e., that $\mathcal{H}^0$ obeys Eqs. (33) and (34)) we conclude that $\mathcal{H}^\ell = \{\cup_i V_1^i, \dots, \cup_i V_k^i\}$ is an $O(\epsilon)$ -approximation of the partitioning problem (as defined by Eqs. (33) and (34)).

9.1.2. Simulating the Oracles. We next get rid of the oracles used in each iteration. It is not hard to see that we do not actually need an oracle to give us the $\beta_j^{i,\tilde{\alpha}}$'s. Instead, we try all possibilities. Recall that i takes on $\ell = O(1/\epsilon)$ values, and for each i , there are $O(1/\epsilon)^k$ possible values of $\tilde{\alpha}$ (i.e., clusters). Finally, for each i , $\tilde{\alpha}$, and $j \in \{1, \dots, k\}$, there are $O(1/\epsilon)$ possible values of $\beta_j^{i,\tilde{\alpha}}$. Thus, we simply try all possible

$$\begin{aligned} O((1/\epsilon)^k)^{O(1/\epsilon)^k \cdot \ell} &< \exp(k \cdot O(1/\epsilon)^{k+1} \cdot k \log(1/\epsilon)) \\ &< \exp\left(\left(\frac{O(1)}{\epsilon}\right)^{k+1} \cdot \log(1/\epsilon)\right). \end{aligned} \quad (41)$$

(vectors of) values for the quantitative partitions of the clusters. Each one gives rise to a different partition of V , and we can search among these partitions for an ϵ -approximation of the partitioning problem.

In order to approximately simulate the oracles for $\gamma_j^i(v)$, we apply the same technique used in previous sections. Namely, we uniformly select ℓ sets, $U^1, \dots, U^\ell$, where $U^i \subset V \setminus V^i$ (each of size $t = \Theta(\epsilon^{-2} \log(k/(\epsilon\delta)))$), and for each i we consider all k -way partitions $(U_1^i, \dots, U_k^i)$, of U^i . For each possible sequence of partitions (i.e., one partition per U^i), when we partition V^i , for each $v \in V^i$ we use the approximation $\hat{\gamma}_j^i(v) \stackrel{\text{def}}{=} |\Gamma(v) \cap U_j^i|/t$ for $\gamma_j^i(v)$. As we prove in detail in Lemma 9.2.1, with high probability over the choice of $U^1, \dots, U^\ell$, there exist *representative* partitions of the U^i , such that $\hat{\gamma}_j^i(v)$ is in fact a good approximation of $\gamma_j^i(v)$ for almost all vertices.

9.2. THE PRELIMINARY PARTITIONING ALGORITHM. A detailed description of the graph partitioning algorithm is given in Figure 8, and its formal analysis is provided in Lemma 9.2.1.

LEMMA 9.2.1. *Let $\mathcal{H} = (H_1, \dots, H_k)$ be a fixed partition of $V(G)$. Then with probability at least $1 - \delta$ over the choice of $\bar{U}$, there exists a partition $\Pi = (U^1, \dots, U^\ell)$ of $\bar{U}$, and a setting of $\tilde{\beta}$, such that*

$$\forall j, \left| \frac{|V_j^{\Pi, \tilde{\beta}}| - |H_j|}{N} \right| \leq \epsilon,$$

Graph Partitioning Algorithm for $\mathcal{GP}_\Phi$

Let $\Lambda = \{\vec{\alpha} : \vec{\alpha} = \langle \alpha_1, \dots, \alpha_k \rangle, \alpha_i \in [0, 1] \text{ is an integer multiple of } \epsilon/16\}$.

- (1) Choose $\ell = \lceil \frac{4}{\epsilon} \rceil$ sets $U^1, \dots, U^\ell$ each of size $t = \Theta\left(\frac{1}{\epsilon^2} \log \frac{k}{\epsilon \delta}\right)$, where U^i is chosen uniformly in $V \setminus V^i$. Let $\bar{U} = \langle U^1, \dots, U^\ell \rangle$.
- (2) For every setting of $\vec{\beta} = \langle \vec{\beta}^1, \dots, \vec{\beta}^\ell \rangle$, where $\vec{\beta}^i = \langle \beta_1^{i, \vec{\alpha}}, \dots, \beta_k^{i, \vec{\alpha}} \rangle_{\vec{\alpha} \in \Lambda}$ and $\beta_j^{i, \vec{\alpha}} \in [0, 1]$ is an integer multiple of $\epsilon/8$, and for each of the sequence of partitions $\Pi = \langle (U_1^1, U_k^1), \dots, (U_1^\ell, U_k^\ell) \rangle$, construct a partition $\mathcal{V}^{\Pi, \vec{\beta}} = (V_1^{\Pi, \vec{\beta}}, \dots, V_k^{\Pi, \vec{\beta}})$ as follows:
 - (a) For $i = 1 \dots \ell$ do:
 - i. For each $v \in V^i$, and for $j \in \{1, \dots, k\}$, let $\hat{\gamma}_j^i(v) \stackrel{\text{def}}{=} \frac{1}{t} \cdot |\Gamma(v) \cap U_j^i|$.
 - ii. For each $\vec{\alpha} \in \Lambda$ let the $\vec{\alpha}$ -cluster of V^i , denoted $V^{i, \vec{\alpha}}$, be defined as follows:
 $V^{i, \vec{\alpha}} \stackrel{\text{def}}{=} \{v \in V^i : \forall j, \alpha_j - \epsilon/32 < \hat{\gamma}_j^i(v) \leq \alpha_j + \epsilon/32\}$.
 - iii. For each cluster $V^{i, \vec{\alpha}}$, $\vec{\alpha} \in \Lambda$, let $V_1^{i, \vec{\alpha}}$ be the first $\beta_1^{i, \vec{\alpha}} \cdot |V^{i, \vec{\alpha}}|$ vertices in $V^{i, \vec{\alpha}}$, let $V_2^{i, \vec{\alpha}}$ be the next $\beta_2^{i, \vec{\alpha}} \cdot |V^{i, \vec{\alpha}}|$ vertices in $V^{i, \vec{\alpha}}$, and in general, let $V_j^{i, \vec{\alpha}}$ be the first $\beta_j^{i, \vec{\alpha}} \cdot |V^{i, \vec{\alpha}}|$ vertices in $V^{i, \vec{\alpha}} \setminus \bigcup_{j' < j} V_{j'}^{i, \vec{\alpha}}$ (where in case the resulting numbers are not integers, they are rounded, alternating between rounding up and down).
 That is, the $\vec{\alpha}$ -cluster of V^i is partitioned according to the (guess of the) corresponding quantitative partition (i.e., the $\beta_j^{i, \vec{\alpha}}$'s).
 - iv. For each $j \in \{1, \dots, k\}$, let $V_j^i = \bigcup_{\vec{\alpha} \in \Lambda} V_j^{i, \vec{\alpha}}$.
 - (b) For each $j \in \{1, \dots, k\}$, let $V_j = \bigcup_{i=1}^\ell V_j^i$.
 Since each V_j actually depends on Π and $\vec{\beta}$ we denote it by $V_j^{\Pi, \vec{\beta}}$.
 - (c) If the partition $\mathcal{V}^{\Pi, \vec{\beta}}$ obeys Equations (35) and (36) then Output $\mathcal{V}^{\Pi, \vec{\beta}}$.

FIG. 8. Graph Partitioning Algorithm.

and

$$\forall j, j', \left| \frac{|E(V_j^{\Pi, \vec{\beta}}, V_{j'}^{\Pi, \vec{\beta}})| - |E(H_j, H_{j'})|}{N^2} \right| \leq \epsilon.$$

PROOF. For a fixed partition Π of $\bar{U}$ and a fixed setting $\vec{\beta} = \langle \vec{\beta}^1, \dots, \vec{\beta}^\ell \rangle$, we consider the following $\ell + 1$ hybrid partitions: The hybrid $\mathcal{H}^0 = (H_1^0, \dots, H_k^0)$ is simply $\mathcal{H}$; the hybrid partition $\mathcal{H}^i = (H_1^i, \dots, H_k^i)$ is defined as follows:

$$H_j^i \stackrel{\text{def}}{=} W_j^{i-1} \cup V_j^i,$$

where the partition $W^{i-1} = (W_1^{i-1}, \dots, W_k^{i-1})$ of $V \setminus V^i$, is defined by $W_j^{i-1} \stackrel{\text{def}}{=} H_j^{i-1} \setminus V^i$, and V_j^i 's are determined as in Figure 8. We shall show that for every $1 \leq i \leq \ell$, for a given choice and partitions of $U^1, \dots, U^{i-1}$, and for a given setting of $\vec{\beta}^1, \dots, \vec{\beta}^{i-1}$, there always exists a setting of $\vec{\beta}^i$ (referred to as the *correct* setting), and with probability at least $1 - \delta/\ell$ over the choice of U^i , there exists a partition $(U_1^i, \dots, U_k^i)$ of U^i such that

$$\forall j, \left| \frac{|H_j^i| - |H_j^{i-1}|}{N} \right| \leq \frac{\epsilon}{16\ell}, \quad (42)$$

and

$$\forall j, j', \left| \frac{|E(H_j^i, H_{j'}^i)| - |E(H_j^{i-1}, H_{j'}^{i-1})|}{N^2} \right| \leq \frac{\epsilon}{\ell}. \quad (43)$$

The lemma follows by induction on i .

Let $(U_1^i, \dots, U_k^i)$ be the partition of U^i induces by $\mathcal{W}^{i-1}$. Namely, for each j , $U_j^i \stackrel{\text{def}}{=} W_j^{i-1} \cap U^i$. We say that U^i is *good* with respect to $(W_1^{i-1}, \dots, W_k^{i-1})$ and V^i if the following holds. For all but at most an $\epsilon/8$ fraction of the vertices v in V^i ,

$$\text{For each } j \in \{1, \dots, k\} \frac{|\Gamma(v) \cap U_j^i|}{t} = \frac{|\Gamma(v) \cap W_j^{i-1}|}{N} \pm \frac{\epsilon}{32}. \quad (44)$$

If the above holds then we say that $(U_1^i, \dots, U_k^i)$ is *representative* with respect to $(W_1^{i-1}, \dots, W_k^{i-1})$ and V^i .

9.2.1. Proving Eqs. (42) and (43) for a Good U^i . Assume first that $(U_1^i, \dots, U_k^i)$ is representative with respect to $(W_1^{i-1}, \dots, W_k^{i-1})$ and V^i , and let the clusters $V^{i,\tilde{\alpha}}$ be as defined in Figure 8, Step (2.a.ii), where $\hat{\gamma}_j^i(v)$ (for every $v \in V^i$ and j) is defined based on V_j^i as in Step (2.a.i). Let $\tilde{\beta}^i = \beta_1^{i,\tilde{\alpha}}, \dots, \beta_k^{i,\tilde{\alpha}} \rangle_{\tilde{\alpha} \in \Lambda}$ be such that for every $\alpha \in \Lambda$, and for each j ,

$$\beta_j^{i,\tilde{\alpha}} = \frac{|H_j^{i-1} \cap V^{i,\tilde{\alpha}}|}{|V^{i,\tilde{\alpha}}|} \pm \frac{\epsilon}{16}.$$

Since each $\beta_j^{i,\tilde{\alpha}}$ takes on all values which are multiples of $\epsilon/8$, there is indeed such a setting of $\tilde{\beta}^i$'s. It follows that when the vertices in V^i are partitioning using this (approximately correct) $\tilde{\beta}^i$, then Eq. (42) holds. That is,

$$\begin{aligned} ||H_j^i| - |H_j^{i-1}|| &= \left| \sum_{\tilde{\alpha}} |H_j^{i-1} \cap V^{i,\tilde{\alpha}}| - \sum_{\tilde{\alpha}} \beta_j^{i,\tilde{\alpha}} \cdot |V^{i,\tilde{\alpha}}| \right| \\ &\leq \sum_{\tilde{\alpha}} \frac{\epsilon}{16} \cdot |V^{i,\tilde{\alpha}}| \\ &= \frac{\epsilon}{16} \cdot \frac{N}{\ell}. \end{aligned}$$

We now show that Eq. (43) holds as well. Towards this end, we fix arbitrary j, j' and consider the contribution of three types of edges to the left-hand side of Eq. (43):

- (1) The contribution of edges with both endpoints NOT in V^i . Since $v \in H_j^i$ iff $v \in H_j^{i-1}$, for every $v \notin V^i$, such edges do not contribute to the difference (i.e., to the left-hand side of Eq. (43)).
- (2) The contribution of edges with both endpoints IN V^i . There are at most $|V^i|^2$ such edges. Using $|V^i| = N/\ell = (\epsilon/4)N$, the potential contribution of these edges is bounded by $(\epsilon/4\ell)N^2$.

- (3) The contribution of edges with exactly ONE endpoint in V^i . We distinguish two cases.
- (a) Edges incident to vertices in V^i for which Eq. (44) does NOT hold. Since $(U_1^i, \dots, U_k^i)$ is representative with respect to $(W_1^{i-1}, \dots, W_k^{i-1})$ and V^i , there are at most $(\epsilon/8\ell)N$ such vertices. Thus, the contribution of these edges to the left-hand side of Eq. (43) is bounded above by $(\epsilon/8\ell)N \cdot 2N = (\epsilon/4\ell)N^2$
 - (b) Edges incident to vertices in V^i for which Eq. (44) DOES hold. The contribution of these edges is due to two types of approximation errors. The first approximation error is due to the clustering itself. That is, vertices in V^i that belong to the same cluster, might differ by $(\epsilon/8)N$ in the number of neighbors they have in each W_j^{i-1} . Specifically, for every v, v' which belong to the same cluster $V^{i,\tilde{\alpha}}$ and for each j , we have

$$\begin{aligned}
 |\gamma_j^i(v) - \gamma_j^i(v')| &\leq |\hat{\gamma}_j^i(v) - \hat{\gamma}_j^i(v')| + |\hat{\gamma}_j^i(v) - \gamma_j^i(v)| + |\hat{\gamma}_j^i(v') - \gamma_j^i(v')| \\
 &\leq \frac{\epsilon}{16} + 2 \cdot \frac{\epsilon}{32} \\
 &= \frac{\epsilon}{8}
 \end{aligned}$$

The second approximation error is due to the fact that the fractional partition is specified with bounded precision (i.e., $\beta_j^{i,\tilde{\alpha}} = |\mathbf{V}_j^{i-1} \cap V^{i,\tilde{\alpha}}|/|V^{i,\tilde{\alpha}}| \pm \epsilon/16$), and so at most $(\epsilon/16\ell) \cdot N$ vertices might be misplaced, contributing at most $(\epsilon/8\ell) \cdot N^2$ edges. Thus, the contribution of both cases to the left-hand side of Eq. (43) is bounded above by $(3\epsilon/8\ell)N^2$.

Summing all cases we get a contribution bounded above by $(1/4 + 1/4 + 3/8) \cdot (\epsilon/\ell)N^2 < (\epsilon/\ell)N^2$.

9.2.2. Bounding the Probability that U^i is not Good. We first fix a vertex $v \in V^i$. Let $U^i = \{u_1, \dots, u_t\}$. For $j \in \{1, \dots, k\}$, and for $1 \leq s \leq t$, define a 0/1 random variable, ξ_j^s , which is 1 if u_s is a neighbor of v and $u_s \in W_j^{i-1}$, and is 0 otherwise. By definition, for each particular j , the sum of the ξ_j^s 's is simply the number of neighbors v has in U_j^i , and the probability that $\xi_j^s = 1$ is $1/N |\Gamma(v) \cap W_j^{i-1}|$. By an additive Chernoff bound (see Appendix B) and our choice of t , for each $j \in \{1, \dots, k\}$,

$$\Pr_{U^i} \left[\left| \frac{|\Gamma(v) \cap U_j^i|}{t} - \frac{|\Gamma(v) \cap W_j^{i-1}|}{N} \right| > \frac{\epsilon}{32} \right] = \exp(-\Omega(\epsilon^2 t)) = \frac{\epsilon \delta}{8k\ell}.$$

By Markov's inequality (see Appendix B), for each $j \in \{1, \dots, k\}$, with probability at least $1 - \delta/k\ell$ over the choice of U^i , for all but $(1/8)\epsilon$ of the vertices in V^i , Eq. (44) holds (for that j), and thus with probability at least $1 - (\delta/\ell)$, U^i is good as required. This concludes the proof of Lemma 9.2.1. $\square$

By considering a partition $\mathcal{H}$ for which Eqs. (33) and (34) hold, Lemma 9.2.1 implies that if $G \in \mathcal{GP}_\Phi$, then with high probability the *Graph Partitioning*

Testing Algorithm for $\mathcal{GP}_\Phi$

Let $\epsilon' \stackrel{\text{def}}{=} \frac{\epsilon}{3k^2}$.

(1) Choose $\ell = \frac{4}{\epsilon'}$ sets $U^1, \dots, U^\ell$ each of size $t = \Theta\left(\left(\frac{1}{\epsilon'}\right)^2 \log \frac{k}{\epsilon'\delta}\right)$, where U^i is chosen uniformly in $V \setminus V^i$. Let $\bar{U} = \langle U^1, \dots, U^\ell \rangle$.

(2) Uniformly choose a set $S = \{s_1, \dots, s_m\}$ of size $m = \Theta\left(\left(\frac{O(1)}{\epsilon'}\right)^{2k+5} \cdot \log(1/(\epsilon'\delta))\right)$.

For $1 \leq i \leq \ell$, let $S^i \stackrel{\text{def}}{=} V^i \cap S$.

(3) For every setting of $\vec{\beta} = \langle \vec{\beta}^1, \dots, \vec{\beta}^\ell \rangle$ and for each of the sequence of partitions $\Pi = \langle (U_1^1, U_k^1), \dots, (U_1^\ell, U_k^\ell) \rangle$ (as in Figure 8),

construct a partition $S^{\Pi, \vec{\beta}} = (S_1^{\Pi, \vec{\beta}}, \dots, S_k^{\Pi, \vec{\beta}})$ as follows:

(a) For $i = 1 \dots \ell$ do:

- i. For each $v \in S^i$, and for $j \in \{1, \dots, k\}$, let $\hat{\gamma}_j^i(v) \stackrel{\text{def}}{=} \frac{1}{t} (|\Gamma(v) \cap U_j^i|)$.
- ii. For each $\vec{\alpha} \in \Lambda$ let the $\vec{\alpha}$ -cluster of S^i , denoted $S^{i, \vec{\alpha}}$ be defined as follows:
 $S^{i, \vec{\alpha}} \stackrel{\text{def}}{=} \{v \in S^i : \forall j, \alpha_j - \epsilon'/32 < \hat{\gamma}_j^i(v) \leq \alpha_j + \epsilon'/32\}$.
- iii. For each cluster $S^{i, \vec{\alpha}}$, $\vec{\alpha} \in \Lambda$, let $S_1^{i, \vec{\alpha}}$ be the first $\beta_1^{i, \vec{\alpha}} \cdot |S^{i, \vec{\alpha}}|$ vertices in $S^{i, \vec{\alpha}}$, let $S_2^{i, \vec{\alpha}}$ be the next $\beta_2^{i, \vec{\alpha}} \cdot |S^{i, \vec{\alpha}}|$ vertices in $S^{i, \vec{\alpha}}$ and in general, for $j < k$ let $S_j^{i, \vec{\alpha}}$ be the first $\beta_j^{i, \vec{\alpha}} \cdot |S^{i, \vec{\alpha}}|$ vertices in $S^{i, \vec{\alpha}} \setminus \bigcup_{j' < j} S_{j'}^{i, \vec{\alpha}}$, and let $S_k^{i, \vec{\alpha}}$ be the remaining vertices in $S^{i, \vec{\alpha}}$ (where in case the resulting numbers are not integers, they are rounded, alternating between rounding up and down).

iv. For each $j \in \{1, \dots, k\}$, let $S_j^i = \bigcup_{\vec{\alpha} \in \Lambda} S_j^{i, \vec{\alpha}}$.

(b) For each $j \in \{1, \dots, k\}$, let $S_j = \bigcup_{i=1}^\ell S_j^i$.

(4) For each partition, $S^{\Pi, \vec{\beta}}$, let $\rho_j^{\Pi, \vec{\beta}} \stackrel{\text{def}}{=} \frac{1}{m} \cdot |S^{\Pi, \vec{\beta}}|$ ($\forall j$)

and $\varrho_{j, j'}^{\Pi, \vec{\beta}} \stackrel{\text{def}}{=} \frac{1}{m/2} \left| \{(s_{2r-1}, s_{2r}) \in E(S_j^{\Pi, \vec{\beta}}, S_{j'}^{\Pi, \vec{\beta}}) \cup E(S_{j'}^{\Pi, \vec{\beta}}, S_j^{\Pi, \vec{\beta}})\} \right|$ ($\forall j, j'$).

If for some $\Pi, \vec{\beta}$ the set of densities $\{\rho_j^{\Pi, \vec{\beta}}\} \cup \{\varrho_{j, j'}^{\Pi, \vec{\beta}}\}$ is $2\epsilon'$ -close to an admissible set w.r.t. Φ then output Accept, otherwise output Reject.

FIG. 9. Testing Algorithm for $\mathcal{GP}_\Phi$.

Algorithm described in Figure 8 will find an ϵ -approximation for the partitioning problem defined by Φ . As noted previously, the number of sequences of partitions considered is bounded by $k^{\ell t} = \exp(\ell t \log k) = \exp(O(\log k / \epsilon^3 \log(k/\epsilon\delta)))$, and the number of settings of $\vec{\beta}$ is $\exp(O(1/\epsilon)^{k+1} \cdot \log(1/\epsilon))$ (see Eq. (41)). Hence, the number of iterations (in Step (2)), is $\exp(O(1/\epsilon)^{k+1} \cdot \log(1/(\epsilon\delta)))$. In each iteration, the running time is governed by Step (2.c) that takes $O(N^2)$ time. As was mentioned previously, we later show how to produce such a partitioning in a more efficient way by first running the testing algorithm for $\mathcal{GP}_\Phi$.

9.3. THE TESTING ALGORITHM. The testing algorithm for $\mathcal{GP}_\Phi$ (described in Figure 9) essentially performs the same steps as the graph partitioning algorithm on a small sample, S , and it is analyzed below. An important technical detail is that (in Step (4)) the tester checks whether the sampled densities are close to an *admissible set of densities* (defined below), rather than testing if they satisfy inequalities analogous to Eqs. (35) and (36). As noted previously, a graph may have a partition that satisfies Eqs. (35) and (36), but still be more than ϵ -far from

the class $\mathcal{GP}_\Phi$. By checking whether the sampled densities are close to an admissible set, we are verifying directly whether the tested graph is close to the class $\mathcal{GP}_\Phi$.

In the following definition, for each j , α_j corresponds to the fraction of vertices in the j th component of the partition (so in particular the α_j 's must sum to 1), and for each j, j' , $b_{j,j'}$ corresponds to the fraction of edges (among all vertex pairs) between components j and j' . The first two items in the definition merely state conditions on these α_j 's and $b_{j,j'}$'s that must hold in any graph. The last item refers to the desired bounds with respect to Φ itself.

Definition 9.3.1. (Admissible Set of Densities). A set of (non-negative) reals, $\{a_j\} \cup \{b_{j,j'}\}$, is called *admissible* with respect to $\Phi (= \{\rho_j^{\text{LB}}, \rho_j^{\text{UB}}\} \cup \{\varrho_{j,j'}^{\text{LB}}, \varrho_{j,j'}^{\text{UB}}\})$ if it satisfies the following inequalities

$$\sum_{j=1}^k a_j = 1 \quad \text{and} \quad \sum_{j,j'=1}^k b_{j,j'} \leq 1 \quad (45)$$

$$b_{j,j} \leq a_j^2 \quad (\forall j) \quad \text{and} \quad b_{j,j'} \leq 2 \cdot a_j \cdot a_{j'} \quad (\forall j \neq j') \quad (46)$$

and

$$\rho_j^{\text{LB}} \leq a_j \leq \rho_j^{\text{UB}} \quad (\forall j) \quad \text{and} \quad \varrho_{j,j'}^{\text{LB}} \leq b_{j,j'} \leq \varrho_{j,j'}^{\text{UB}} \quad (\forall j, j'). \quad (47)$$

In Step (4) of Figure 9, we check that a set of densities $\{\rho_j^{\Pi, \tilde{\beta}}\} \cup \{\varrho_{j,j'}^{\Pi, \tilde{\beta}}\}$ is $2\epsilon'$ -close to an admissible set where $\epsilon' \stackrel{\text{def}}{=} \epsilon/3k^2$. That is, given the $\rho_j^{\Pi, \tilde{\beta}}$'s and $\varrho_{j,j'}^{\Pi, \tilde{\beta}}$'s, we ask whether there are nonnegative a_j 's and $b_{j,j'}$'s which, in addition to the above admissibility conditions, also satisfy

$$|a_j - \rho_j^{\Pi, \tilde{\beta}}| \leq 2\epsilon' \quad (\forall j) \quad \text{and} \quad |b_{j,j'} - \varrho_{j,j'}^{\Pi, \tilde{\beta}}| \leq 2\epsilon' \quad (\forall j, j')$$

We note that the problem of whether a set of densities is $2\epsilon'$ -close to an admissible set for Φ can be solved in $\exp(\text{poly}(k) \cdot L)$ -time, where L is the length of the encoding (in binary) of Φ and ϵ (see Appendix C). Note that this time-bound is dominated by the number of partitions examined in Step (4) (of Figure 9).

PROOF OF THEOREM 9.2. Recall that $\epsilon' = \epsilon/3k^2$. We first note that if $G \in \mathcal{GP}_\Phi$, then by Lemma 9.2.1, with probability $1 - (\delta/2)$ over the choice of $\bar{U}$, there exist Π and $\tilde{\beta}$ such that $\mathcal{V}^{\Pi, \tilde{\beta}}$ is an ϵ' -approximation of the partition problem. Furthermore, the densities related to this partition are ϵ' -close to an admissible set for Φ (that is, to the set of densities defined by a partition $(H_1, \dots, H_k)$ for which Eqs. (33) and (34) hold). On the other hand, if G is ϵ -far from $\mathcal{GP}_\Phi$, then no k -way partition of V has densities which are $3\epsilon'$ -close to an admissible set with respect to Φ .²⁵

²⁵ Here is where the notion of admissible solution plays an important role. As noted in the beginning of this section, it would not suffice to check whether no k -way partition of V is an $O(\epsilon/k^2)$ -good approximation (of the partition problem). This is true since the existence of a good approximation does not ensure that the distance of G to the class is of the same order as the approximation.

Suppose, on the contrary, that V has a k -way partition with densities that are $3\epsilon'$ -close to an admissible set with respect to Φ . Then we can modify this partition in two steps as follows:

- Step 1.** We move up to $3k\epsilon' \cdot N$ vertices (according to the guaranteed admissible vertex densities), so that two conditions hold. First, the number of vertices in each component of the partition is between the required lower and upper bounds. Second, for each j, j' , the number of vertex pairs, one in component j and one in component j' , is at least as required by the edge-density lower bound $\varrho_{j,j'}^{\text{LB}}$.
- Step 2.** We omit/add up to $3k^2\epsilon' \cdot N^2 = \epsilon N^2$ edges (according to the required edge-densities upper and lower bounds). Hence, we obtain a graph in $\mathcal{GP}_\Phi$ in contradiction to the hypothesis.

Theorem 9.2 follows from the lemma below, which analogously to Lemmas 7.3.1 and 8.2.2.1 deals with the performance of the partitions of S as an “approximation” to the corresponding partitions of V . For a fixed $\bar{U}$ (which we assume is good, as defined in the proof of Lemma 9.2.1), we now consider the partitions $\mathcal{V}^{\Pi, \tilde{\beta}}$ s which would result in running the Graph Partitioning Algorithm with an approximation parameter $\epsilon' = \epsilon/3k^2$, and a confidence parameter $\delta/2$. We relate these partitions to the partitions $\mathcal{G}^{\Pi, \tilde{\beta}}$ s that are generated by the Testing Algorithm when using the same $\bar{U}$. Let n_p , n_a , and n_b be upper bounds on the number of sequences of partitions Π , number of clusters (one per each i and each setting of $\tilde{\alpha}$), and the number of settings of $\tilde{\beta}$, respectively. Thus,

$$n_p = k^{\ell_t},$$

$$n_a = \ell \cdot \left(\frac{16}{\epsilon'} \right)^k = \left(\frac{O(1)}{\epsilon'} \right)^{k+1},$$

and

$$n_b = \left(\frac{8}{\epsilon'} \right)^{k \cdot n_a} = \exp \left(\log \left(\frac{1}{\epsilon'} \right) \cdot \left(\frac{O(1)}{\epsilon'} \right)^{k+1} \right).$$

Note that $n_p \cdot n_b$ is the number of partitions considered by each algorithm.

LEMMA 9.3.2. *For a fixed $\bar{U}$, Π , and $\tilde{\beta}$, let $\mathcal{V}^{\Pi, \tilde{\beta}}$ be defined as in the Graph Partitioning Algorithm, when executed with approximation parameter $\epsilon' = \epsilon/3k^2$. Let S be a uniformly chosen sample of size*

$$m = \Omega \left(\frac{n_a}{\epsilon'} \cdot \frac{\log(k \cdot n_p \cdot n_a \cdot n_b)/\delta}{(\epsilon')^2} \right)$$

and let $\mathcal{G}^{\Pi, \tilde{\beta}}$, $\{\rho_j^{\Pi, \tilde{\beta}}\}$, and $\{\varrho_{j,j'}^{\tilde{\beta}}\}$ be as defined by the Testing algorithm. Then

$$\Pr_S \left[\exists j : \left| \rho_j^{\Pi, \tilde{\beta}} - \frac{|V_j^{\Pi, \tilde{\beta}}|}{N} \right| > \epsilon' \right] < \frac{\delta}{4 \cdot n_p \cdot n_b}$$

and

$$\Pr_S \left[\exists j, j' : \varrho_{j,j'}^{\Pi, \tilde{\beta}} - \frac{|E(V_j^{\Pi, \tilde{\beta}}, V_{j'}^{\Pi, \tilde{\beta}})|}{N^2} > \epsilon' \right] < \frac{\delta}{4 \cdot n_p \cdot n_b}.$$

PROOF. Let $\delta' \stackrel{\text{def}}{=} \delta/(4 \cdot n_p \cdot n_b)$ and let $\epsilon_1 \stackrel{\text{def}}{=} \epsilon'/24$. For each $i \in \{1, \dots, \ell\}$ and $\tilde{\alpha}$, let $V_j^{i, \tilde{\alpha}}$ and $V_j^{i, \tilde{\alpha}}$ be as defined in the Graph Partitioning Algorithm, and let $S_j^{i, \tilde{\alpha}}$ and $S_j^{i, \tilde{\alpha}}$ be as defined in the Testing algorithm (for the fixed $\bar{U}$, Π and $\tilde{\beta}$ considered in the lemma). By an additive Chernoff bound, for each cluster,

$$\Pr_S \left[\left| \frac{|S_j^{i, \tilde{\alpha}}|}{m} - \frac{|V_j^{i, \tilde{\alpha}}|}{N} \right| > \frac{\epsilon'}{2} \right] < \exp(-\Omega(m \cdot (\epsilon')^2)) < \frac{\delta'}{4n_a}. \quad (48)$$

Let V_j be the union of the $V_j^{i, \tilde{\alpha}}$'s and let S_j be the union of the $S_j^{i, \tilde{\alpha}}$'s (as defined by the algorithm). Recall that for each j , $\lfloor \tilde{\beta}^{i, \tilde{\alpha}} \cdot |V_j^{i, \tilde{\alpha}}| \rfloor \leq |V_j^{i, \tilde{\alpha}}| \leq \lceil \tilde{\beta}^{i, \tilde{\alpha}} \cdot |V_j^{i, \tilde{\alpha}}| \rceil$, and similarly, $\lfloor \tilde{\beta}^{i, \tilde{\alpha}} \cdot |S_j^{i, \tilde{\alpha}}| \rfloor \leq |S_j^{i, \tilde{\alpha}}| \leq \lceil \tilde{\beta}^{i, \tilde{\alpha}} \cdot |S_j^{i, \tilde{\alpha}}| \rceil$. Thus the effects of rounding can be taken into account by assuming an arbitrary placement of $2 \cdot n_a$ vertices. Equation (48) together with this bound on the rounding effects implies that for every j , with probability at least $1 - \delta'/4$,

$$\frac{|S_j|}{m} = \sum_{i, \tilde{\alpha}} \beta_j^{i, \tilde{\alpha}} \cdot \frac{|S_j^{i, \tilde{\alpha}}|}{m} \pm \frac{n_a}{m} = \sum_{i, \tilde{\alpha}} \beta_j^{i, \tilde{\alpha}} \cdot \left(\frac{|V_j^{i, \tilde{\alpha}}|}{N} \pm \frac{\epsilon'}{2} \right) \pm \frac{2n_a}{m} = \frac{|V_j|}{N} \pm \epsilon',$$

so we have proved the first part of the lemma.

We say that a cluster $V_j^{i, \tilde{\alpha}}$ is *significant*, if it contains at least $\epsilon_1 \cdot N/n_a$ vertices. Otherwise, it is insignificant. Note that the total number of vertices belonging to insignificant clusters is at most $\epsilon_1 \cdot N$. For each significant cluster and each j , let $\bar{V}_j^{i, \tilde{\alpha}}$ be the same as $V_j^{i, \tilde{\alpha}}$ except that it does not contain the first and the last $\epsilon_1 \cdot |V_j^{i, \tilde{\alpha}}|$ vertices in the cluster (where first and last are with respect to the lexicographic order on vertices). For each i and $\tilde{\alpha}$, let $m_{i, \tilde{\alpha}} \stackrel{\text{def}}{=} |S_j^{i, \tilde{\alpha}}|$. By a multiplicative Chernoff bound, with probability at least $1 - \delta'/12$, for every significant cluster,

$$m_{i, \tilde{\alpha}} \geq \frac{1}{2} \cdot m \cdot \frac{|V_j^{i, \tilde{\alpha}}|}{N} \geq \frac{1}{2} \cdot m \cdot \frac{\epsilon_1}{n_a} = \Omega\left(\frac{\log(k \cdot n_b/\delta')}{(\epsilon')^2}\right).$$

Therefore, by an additive Chernoff bound, for each significant cluster, and each j ,

$$\Pr_S \left[\left| \frac{|S_j^{i, \tilde{\alpha}} \cap \bar{V}_j^{i, \tilde{\alpha}}|}{m_{i, \tilde{\alpha}}} - (\tilde{\beta}_j^{i, \tilde{\alpha}} - 2\epsilon_1) \right| > 2\epsilon_1 \right] < \frac{\delta'}{12 \cdot k \cdot n_a}. \quad (49)$$

Thus, putting aside an error probability of $\delta'/12$, assume from now on that in fact for every significant cluster and every j ,

- (1) $|S_j^{i, \tilde{\alpha}} \cap \bar{V}_j^{i, \tilde{\alpha}}| \leq \tilde{\beta}_j^{i, \tilde{\alpha}} \cdot m_{i, \tilde{\alpha}}$ from which it follows (by definition of $S_j^{i, \tilde{\alpha}}$) that for every j , $S_j^{i, \tilde{\alpha}} \cap \bar{V}_j^{i, \tilde{\alpha}} = S_j^{i, \tilde{\alpha}} \cap \bar{V}_j^{i, \tilde{\alpha}}$;
- (2) $|S_j^{i, \tilde{\alpha}} \cap \bar{V}_j^{i, \tilde{\alpha}}| \geq (\tilde{\beta}_j^{i, \tilde{\alpha}} - 4\epsilon_1) \cdot m_{i, \tilde{\alpha}}$, from which it follows that $|S_j^{i, \tilde{\alpha}} \cap \bar{V}_j^{i, \tilde{\alpha}}| \geq (\tilde{\beta}_j^{i, \tilde{\alpha}} - 4\epsilon_1) \cdot m_{i, \tilde{\alpha}}$.

Next observe that for each j, j' ,

$$\begin{aligned} \{(s_{2k-1}, s_{2k}) \in E(S_j, S_{j'})\} &= \{(s_{2k-1}, s_{2k}) \in E(S_j \cap \bar{V}_j, S_{j'} \cap \bar{V}_{j'})\} \\ &\cup \{(s_{2k-1}, s_{2k}) \in (E(S_j \setminus \bar{V}_j, S_{j'}) \cup E(S_j, S_{j'} \setminus \bar{V}_{j'}))\}. \end{aligned} \quad (50)$$

By Item (1) above and an additive Chernoff bound we get that for each j, j'

$$\begin{aligned} \Pr_S \left[\left| \frac{|\{(s_{2k-1}, s_{2i}) \in E(S_j \cap \bar{V}_j, S_{j'} \cap \bar{V}_{j'})\}|}{m/2} - \frac{|E(\bar{V}_j, \bar{V}_{j'})|}{N^2} \right| > \frac{\epsilon'}{4} \right] \\ < \frac{\delta'}{12k^2} \end{aligned} \quad (51)$$

By definition of $\bar{V}_j$ and $\bar{V}_{j'}$, we have that

$$\frac{|E(V_j, V_{j'})|}{N^2} - \frac{|E(\bar{V}_j, \bar{V}_{j'})|}{N^2} \leq 4\epsilon_1 < \frac{\epsilon'}{4} \quad (52)$$

By Item (2) above, we know that for every *significant* cluster and every j , $|S_j^{i,\tilde{\alpha}} \setminus \bar{V}_j^{i,\tilde{\alpha}}| \leq 4\epsilon_1 m_{i,\tilde{\alpha}}$. By a multiplicative Chernoff bound, with probability at least $1 - \delta'/12$, the sum of $m_{i,\tilde{\alpha}}$ taken over all *insignificant* clusters is at most $2\epsilon_1 \cdot m$. Thus,

$$\sum_{i,\tilde{\alpha}} |S_j^{i,\tilde{\alpha}} \setminus V_j^{i,\tilde{\alpha}}| \leq r\epsilon_1 \cdot \sum_{i,\tilde{\alpha}} m_{i,\tilde{\alpha}} + 2\epsilon_1 \cdot m = 6\epsilon_1 m.$$

Therefore,

$$\begin{aligned} &\frac{|\{(s_{2k-1}, s_{2k}) \in (E(S_j \setminus \bar{V}_j, S_{j'}) \cup E(S_j, S_{j'} \setminus \bar{V}_{j'}))\}|}{m/2} \\ &\leq \frac{|\{k : s_{2k-1} \in S_j \setminus \bar{V}_j \text{ or } s_{2k} \in S_{j'} \setminus \bar{V}_{j'}\}|}{m/2} \leq 12\epsilon_1 = \frac{\epsilon'}{2} \end{aligned} \quad (53)$$

Summing up all the probabilities of errors and combining the bounds of Eqs. (51)–(53), the lemma follows. $\square$

PROOF OF THEOREM 9.3. In order to get a partitioning algorithm whose running time is as stated in the theorem, we first run the Testing Algorithm and find a choice of Π and $\tilde{\beta}$ (in time independent of N) for which Step (4) of the Testing Algorithm accepts. Next we use this choice to partition all of V by running Step (2) of the Partition Algorithm only on this choice. As we have shown in the proof of Theorem 9.2, with high probability, a partition $\mathcal{P}^{\Pi,\tilde{\beta}}$ will cause the test to accept if $\mathcal{V}^{\Pi,\tilde{\beta}}$ is an ϵ' -approximation of the partitioning problem, and if $\mathcal{V}^{\Pi,\tilde{\beta}}$ is not an $3\epsilon'$ -approximation, then $\mathcal{P}^{\Pi,\tilde{\beta}}$ will cause the test to reject. Thus, with high probability if the testing algorithm accepted G due to a certain setting of Π and $\tilde{\beta}$, then $\mathcal{V}^{\Pi,\tilde{\beta}}$ is an ϵ -approximation of the partitioning problem. $\square$

10. Extensions, Limitations and Beyond

10.1. EXTENSIONS AND LIMITATIONS OF THE ABOVE ALGORITHMS. We have presented several Graph Property Testers that use queries and are evaluated with respect to the uniform distribution (on pairs of vertices). We now comment on several extensions, variations and considerations.

10.1.1. *Impossibility of Testing without Queries.* Proposition 6.1.2.1 shows that (edge) queries are essential for testing Bipartiteness. The construction used in the proof actually establishes the same for testing the class of graphs having cliques of density at least $1/2$, and for approximating the Max-Cut (of dense graphs up to an $N^2/8$ additive term). Furthermore, the proof can be easily modified to yield the same result for testing k -Colorability, for any $k \geq 3$.

10.1.2. *Extension to Directed Graphs.* Some of the problems we study have analogues in directed graphs. In particular this is true for a directed version of ρ -Cut, where we are interested in testing whether a directed graph has a two-way partition (V_1, V_2) such that the number of edges crossing from V_1 to V_2 is at least ρN^2 . Note that directed graphs in general do not have a symmetric adjacency matrix and the existence of a directed edge from v_1 to v_2 does not necessarily imply the existence of an edge from v_2 to v_1 . For any two disjoint sets of vertices, W_1 and W_2 , let $\mu(W_1, W_2)$ denote density of the edges crossing from W_1 to W_2 (i.e., the number of edges crossing divided by N^2). Similarly to the undirected case, the algorithm is essentially based on the following observation. Consider a vertex v and a partition (W_1, W_2) of $V \setminus \{v\}$. Let $E(v, W_2)$ be the set of edges going from v to vertices in W_2 , and let $E(W_1, v)$ be the set of edges going from vertices in W_1 to v . In case $|E(v, W_2)| > |E(W_1, v)|$, it is preferable to put v on side 1, that is, $\mu(W_1 \cup \{v\}, W_2) > \mu(W_1, W_2 \cup \{v\})$, and in case $|E(W_1, v)| > |E(v, W_2)|$, we should put v on side 2. The notion of *unbalance* is hence slightly modified, but as in the case of undirected cuts, the above observation generalizes to sets of vertices of size $O(\epsilon N)$. Thus the ρ -Directed-Cut testing algorithm is very similar to the ρ -Cut testing algorithm, the only difference is that when deciding where to put a vertex $v \in V^i$ given a fixed partition (U_1, U_2) of the uniformly selected set U , we compare the number of edges going from v to vertices in U_2 to the number of edges going from U_1 to v . The algorithms for k -way-Cut and Bisection are modified similarly.

It is also possible to extend the definition of the general partition problem to directed graphs by allowing the bounds $\varrho_{j,j'}^{\text{LB}}$, and $\varrho_{j,j'}^{\text{UB}}$, to differ from $\varrho_{j',j}^{\text{LB}}$ and $\varrho_{j',j}^{\text{UB}}$, respectively for $j \neq j'$. That is, there are separate requirements on the number of edges crossing from V_j to $V_{j'}$ and the number of edges crossing from $V_{j'}$ to V_j . In such a case, the graph partitioning and testing algorithms for directed graphs differ from the algorithms for undirected graphs only in their definition of *clusters* (but the clusters are partitioned by the algorithms as in the undirected case). A cluster of vertices, defined with respect to a fixed partition $(W_1, \dots, W_k)$, is a set of vertices Z such that for all $v_1, v_2 \in Z$, and for every $j \in \{1, \dots, k\}$, $E(v_1, W_j) \approx E(v_2, W_j)$, and $E(W_j, v_1) \approx E(W_j, v_2)$.

10.1.3. *Extension to Product Distributions.* Our algorithms for k -Colorability, ρ -Clique, and ρ -Cut can be easily extended to provide testers with respect to “product distributions.” We call the distribution $\Psi : V(G)^2 \mapsto [0, 1]$ a *product distribution* if there exists a distribution on vertices $\psi : V(G) \mapsto [0, 1]$ so that

$\Psi(u, v) = \psi(u) \cdot \psi(v)$. Recall that each of our algorithms takes a uniform sample of vertices, and queries the graph only on the existence of edges between these vertices. Instead, in case we need to test G with respect to the product distribution Ψ , we sample $V(G)$ according to the distribution ψ . Note that the algorithm need not “know” ψ ; it suffices that the algorithm has access to a source of vertices drawn from this distribution (or to a source of pairs drawn according to Ψ). To prove that this extension is valid consider an auxiliary graph G' consisting of vertex-components so that each component corresponds to a vertex in G . Complete bipartite graphs will be placed between pairs of components which correspond to edges in G . The size of the component will be related to the probability measure of the corresponding vertex in G . Thus, uniform distribution on $V(G')$ is almost the same as distribution Ψ on $V(G)$. For the analysis of the ρ -Clique tester, we also add edges between vertices residing in the same component. (This is certainly NOT done in analyzing the k -Colorability and Max-Cut testers.) It can be shown that testing G for any of the above mentioned properties with respect to the product distribution Ψ corresponds to testing G' (for the respective property) under the uniform distribution. Suppose, for example, that G is k -Colorable. Then so is G' , and thus, the k -Colorability Tester will always accept G' (and thus always accept G). On the other hand, if the k -Colorability Tester, run with parameters ϵ, δ , accepts G' with probability $1 - \delta$, then there is a k -Coloring of G' which violates less than an ϵ fraction of vertex pairs. To obtain a k -Coloring of G we randomly assign each vertex in G a color according to the proportions of colors assigned to the corresponding vertices in G' . It follows that the expected probability mass (according to Ψ) assigned to violated edges is less than ϵ .

10.1.4. Impossibility of Distribution-Free Testing. In contrast to the above extension, it is not possible to test any of the graph properties discussed above in a distribution-free manner (even with queries). For simplicity, let us consider the case of Bipartite Testing. Consider the following class of distributions on pairs of vertices of an N -vertex graph. Each distribution is defined by a partition of $\{1, \dots, N\}$ into $N/4$ 4-tuples. The distribution assigns probability $1/3N$ to each (ordered) pair of distinct vertices which belong to the same 4-tuple. Pairs residing in different 4-tuples are assigned probability 0. For each distribution, we consider two graphs. The first graph consists of $N/4$ paths of length 3, each residing in a different 4-tuple; whereas the second graph consists of $N/4$ triangles, each residing in a different 4-tuple. Clearly, the first graph is bipartite whereas the second is not. Furthermore, the second graph is $1/6$ -far (with respect to the above distribution) from being bipartite. Still, no tester which works in $o(\sqrt{N})$ time can tell these two graphs apart, even if it gets samples drawn from the distribution and is allowed queries. The reason is that t samples drawn from the distribution will, with probability at least $1 - 4t^2/N$, come from different 4-tuples. Furthermore, in this case, any query made by the tester, except if it is on a pair which has appeared in the sample, is likely to be on a pair which is not from the same 4-tuple and thus a nonedge (in both graphs).

10.1.5. On the Possibility of Working in $\text{poly}(1/\epsilon)$ Time. The algorithm for Bipartite Testing works in $\text{poly}(1/\epsilon)$ -time (see Theorem 6.1.2), whereas all the other testers presented in this paper work in time exponential in $\text{poly}(1/\epsilon)$. In fact, it seems that one cannot hope for much better. For example, we claim that

if one can test 3-Colorability with distance parameter ϵ in time $\text{poly}(1/\epsilon)$ then $\mathcal{NP} \subseteq \mathcal{BPP}$. To decide if a graph G is 3-Colorable, simply set $\epsilon = 1/|V(G)|^2$ (and $\delta = 1/3$) and run the property testing. Clearly, if G is 3-Colorable then the test will accept with probability at least $2/3$, whereas if G is not 3-Colorable it must be ϵ -far from being 3-Colorable and thus be rejected by the test with probability at least $2/3$. We remark that a similar argument can be made when using a relatively bigger value of ϵ . For example, to decide if G is 3-Colorable consider an auxiliary graph, G' , with $2^{|V(G)|}$ vertices which are grouped into $n \stackrel{\text{def}}{=} |V(G)|$ components each consisting of $2^n/n$ vertices. These huge components will correspond to vertices in G and complete bipartite graphs will be placed between pairs of components that correspond to edges in G . Thus, we can set $\epsilon = 1/n^2$ and apply the same reduction as above this time showing that deciding 3-Colorability of G reduces to 3-Colorability testing of G' with distance parameter $\epsilon = 1/\text{poly} \log |V(G')|$.

10.1.6. On One-Sided Failure Probability. The testers for Bipartite and k -Colorability always accept graphs which have the property. In contrast, all other testers presented in this paper may fail to accept a yes-instance (yet, with probability at most δ). We claim that non-zero failure probability, in these cases, is inevitable. Consider for example the execution of a potential Clique Tester given access to a graph with no edges at all. Clearly, the tester must reject with probability at least $1 - \delta$. Fix any sequence of coin tosses (for the tester) which makes it reject. This determines a sequence of queries into the graph (all queries are answered by 0). Assuming that the number of queries is less than $(1 - \rho)N$, there exists an N -vertex graph having a clique of size ρN in which all the queried pairs are non-edges. It follows that the Clique Tester rejects this yes-instance with positive probability.²⁶ We conclude that there is a fundamental difference between testing k -Colorability and testing ρ -Clique.

10.1.7. Impossibility of Learning with Queries. All the above classes of (graph-) functions are not learnable, not even under uniform distribution and when allowing (membership) queries. Furthermore, this holds also if we allow unlimited computing power as long as we restrict the number of queries to $o(N)$. Intuitively, there are “too” many graphs (functions) in each class. Formally, we may consider attempts to learn a random bipartite N -vertex graph and may even fix the 2-partition, say, place vertices $\{1, \dots, N/2\}$ on one side. Then, each uninspected pair (i, j) , with $i \leq N/2$ and $j > N/2$, is equally likely to be either an edge or a non-edge (i.e., be labeled 1 or 0 by the corresponding function).

10.1.8. Extension to Weighted Graphs (with Bounded Weights). Let $G = (V, E)$ be a (simple undirected) graph, and $\omega : E \mapsto [0, B]$ be a weight function assigning each edge $e \in E$ a non-negative value (bounded by B). We assume that B , the bound on the weights of the edges, is known. We associate with G a function $f_{G,\omega} : V \times V \rightarrow [0, B]$, where $f_{G,\omega}(v_1, v_2) = 0$ if $(v_1, v_2) \notin E$ and is $\omega(v_1, v_2)$ otherwise. When performing a query on a pair of vertices, (v_1, v_2) , our testing algorithms receive the value of $f_G(v_1, v_2)$. The notion of distance between graphs is slightly different from the unweighted case: We define the distance

²⁶ An analogous argument can be used to show that any Clique Tester must accept some no-instance with positive probability. (Start by considering an execution on the complete graph.)

between two weighted graphs (G, ω) and (G', ω') , or equivalently between $f_{G,\omega}$ and $f_{G',\omega'}$, to be $1/N^2 \sum_{v_1, v_2 \in V} |f_{G,\omega}(v_1, v_2) - f_{G',\omega'}(v_1, v_2)|$. The distance between a weighted graph and a class of weighted graphs is defined in the obvious manner. The generalization to weighted graphs may affect the testing problems in two ways:

10.1.8.1. AFFECT ON THE OBJECTIVE. Consider, for example, a generalization of ρ -Cut to weighted graphs. A weighted graph $G = (V, E, \omega)$ has a cut of weight at least ρ , if there exists a partition (V_1, V_2) of V such that $1/N^2 \sum_{v_1 \in V_1, v_2 \in V_2} 2f_G(v_1, v_2) \geq \rho$. Thus, the class ρ -Cut is a class of weighted graphs (rather than of graphs). Still, our algorithms for unweighted ρ -Cut are easily adapted to the weighted case; yet, the query complexity (respectively, running-time) of our tester will depend polynomially (respectively, exponentially) on the bound B .²⁷ The k -Cut, Bisection, and General Partition properties and algorithms generalize similarly.

10.1.8.2. AFFECT ON DISTANCE. Consider, for example, testing k -Colorability of weighted graphs. Clearly, the property of being k -Colorable has nothing to do with the weights of the edges; yet, the distance from the class of k -Colorable graph does depend on these weights. However, the latter dependency can be easily waived by replacing each nonzero weighted edge by an edge with weight B . Note that this replacement does not affect k -Colorability and that it only increases the distance of non- k -Colorable weighted graph from the class of k -Colorable weighted graphs. We stress that in the resulting graph all edges have weight B . Thus, testing a weighted graph for k -Colorability with distance parameter ϵ , reduces to testing the underlying (unweighted) graph (for k -Colorability) with distance parameter ϵ/B . Again, the bound B turns out to affect the query complexity and running-time as in the first case.

10.2. TESTING OTHER GRAPH PROPERTIES. The following remarks and observations are meant to indicate that providing a characterization of graph properties according to the complexity of testing them may be quite challenging.

We first recall that testing k -Colorability seems to be fundamentally different from testing ρ -Clique since the first can be done without ever rejecting yes-instances whereas the second cannot be done without two-sided failure probability.

10.2.1. Easy to Test Graph Properties. Next, we observe that any graph property which can be made to hold by adding or omitting few edges from any graph can be tested very easily. Namely,

²⁷ Specifically, all that need be changed is the notion of *balanced vertices*: We say that a vertex v is *unbalanced* with respect to a partition (W_1, W_2) if $1/N |\sum_{w_1 \in W_1} f_G(v, w_1) - \sum_{w_2 \in W_2} f_G(v, w_2)|$ is non-negligible. Similarly to the unweighted case we shall use partitions (U_1, U_2) of a uniformly selected sample U to approximate this difference (and use an additional sample S to approximate weights of cuts). The only difference in the analysis is that instead of using sums of 0/1 random variables to approximate expected values, we are using sums of random variables whose value lies in $[0, B]$, and hence we'll get a polynomial dependence on B in the sample complexity, and an exponential dependency in the running time.

PROPOSITION 10.2.1.1 (TESTING ALMOST TRIVIAL CLASSES). *Let $\alpha > 0$ be a constant and $\mathcal{C}$ a class of graphs so that for every graph, G ,*

$$\text{dist}(G, \mathcal{C}) \leq |V(G)|^{-\alpha}.$$

Then, $\mathcal{C}$ can be tested by using at most $\text{poly}(\epsilon^{-1} \log(1/\delta))$ labeled random examples, where ϵ is the distance parameter. Furthermore, the test always accepts graphs in $\mathcal{C}$.

Classes that satisfy the hypothesis of the proposition include: *Connected Graphs* (add $\leq |V(G)| - 1$ edges), *Eulerian Graphs* (make connected and add $\leq |V(G)|/2$ edges), *Hamiltonian Graphs* (add $\leq |V(G)| - 1$ edges), *Graphs with $\kappa(\cdot)$ -vertex Dominating Set*, for a given $\kappa(\cdot)$ (add $\leq |V(G)| - \kappa(|V(G)|)$ edges), *Graphs having Perfect Matching* (add $\leq |V(G)|/2$ edges), and *Graphs containing a subgraph H* (add $\leq |E(H)|$ edges).²⁸ We remark that the above also holds with respect to some of the complementing classes such as *UnConnected Graphs*, *Non-Hamiltonian Graphs*, *Non-Eulerian Graphs*, *Graphs without $\kappa(\cdot)$ -vertex Dominating Set*, and *Graphs not having Perfect Matching* (e.g., by removing edges to make one vertex isolated).

PROOF. Let $\Delta(N) \stackrel{\text{def}}{=} \max_{G: |V(G)|=N} \{\text{dist}(G, \mathcal{C})\}$ and suppose that $\Delta(N)$ or a good upper bound on it is known. (One may always use $\Delta(N) = N^{-\alpha}$.) Given oracle access to an N -vertex graph G (and a distance parameter ϵ), if $\epsilon > \Delta(N)$, then the tester always accepts. Otherwise, the tester inspects all $N^2 = \text{poly}(1/\epsilon)$ edges and decides accordingly. Actually, we may take a sample of $O(\log(N/\delta) \cdot N^2)$ labeled random examples and accept iff either the sample does not cover all possible vertex pairs or the sample “reveals” a graph in $\mathcal{C}$. The main point is that in case $\epsilon > \Delta(N)$ every N -vertex graph is ϵ -close to $\mathcal{C}$. $\square$

We stress that as long as the distance parameter (i.e., ϵ) is not too small, the tester is trivial (i.e., it accepts any graph without performing any checking). Slightly more is required in order to check any graph property which holds only on very sparse graphs. Here, as long as ϵ is not too small, the tester need only check that random examples of vertex-pairs have no edge between them (i.e., the graph is sufficiently sparse).

PROPOSITION 10.2.1.2 (TESTING CLASSES OF SPARSE GRAPHS). *Let $\alpha > 0$ be a constant and $\mathcal{C}$ a class of graphs so that $\mathcal{C} \subseteq \{G : |E(G)| < |V(G)|^{2-\alpha}\}$. Then, $\mathcal{C}$ can be tested by using at most $\text{poly}(\epsilon^{-1} \log(1/\delta))$ labeled random examples; where ϵ and δ are the distance and confidence parameters.*

Classes that satisfy the hypothesis of the proposition include: *Trees*, *Forests*, and *Planar Graphs* (all with $|E(G)| = O(|V(G)|)$).

PROOF. Let $\rho(N) \stackrel{\text{def}}{=} \max_{G \in \mathcal{C}: |V(G)|=N} \{|E(G)|/N^2\}$ and suppose that $\rho(N)$ or a good upper bound on it is known. (One may always use $\rho(N) = N^{-\alpha}$.) Given oracle access to an N -vertex graph G (and parameters ϵ, δ), if $\epsilon > 6\rho(N)$, then the tester takes a sample of $t \stackrel{\text{def}}{=} O(\epsilon^{-2} \log(1/\delta))$ labeled random examples and accepts iff it has seen at most $(\rho(N) + (\epsilon/3)) \cdot t$ edges. Otherwise, the

²⁸ The latter can be generalized to *Graphs containing a subgraph that can be contracted to one of the graphs in $\{H_1, \dots, H_r\}$* . The Non-Planar Graphs are a special case.

tester inspects all $N^2 = \text{poly}(1/\epsilon)$ edges and decides accordingly. (Again, the actual implementation is by a sample of $O(\log(N/\delta) \cdot N^2)$ edges that is very likely to cover all vertex pairs.) The main point is that the graphs in $\mathcal{C}$ have edge density at most $\rho(N)$, whereas graphs that are ϵ -away from $\mathcal{C}$ have edge density at least $\epsilon - \rho(N)$ which for $\epsilon > 6\rho(N)$ is more than $\rho(N) + 2\epsilon/3$. We conclude by noting that, with probability at least $1 - \delta$ the number of edges seen by the tester provides a good estimate (i.e., $\epsilon/3$ deviation) to the density of the graph. $\square$

We stress that both the above testers make no queries. In contrary, the following slightly more involved tester does make queries in order to check that vertices drawn at random have about the same degree. This algorithm is a tester for the class of Regular Graphs, and its correctness is established based on a theorem due to Noga Alon (private communication).

PROPOSITION 10.2.1.3 (TESTING REGULAR GRAPHS). *The class of regular graphs can be tested by using at most $\text{poly}(\epsilon^{-1} \log(1/\delta))$ queries, where ϵ and δ are the distance and confidence parameters.*

PROOF. Given oracle access to an N -vertex graph G (and parameters ϵ, δ), the test takes a sample, S , of $O(\epsilon^{-1} \log(1/\delta))$ many vertices and for each $v \in S$ makes $O(\epsilon^{-2} \log(1/(\epsilon\delta)))$ queries of the form “is $(v, w) \in E(G)$ ”, where $w \in V(G)$ is uniformly chosen. Thus, the test estimates the degrees of all vertices in S . The test accept iff all the estimated degrees (divided by N) are within $\epsilon/3$ of one another.

Clearly, if G is regular then, with probability at least $1 - \delta$, the test accepts it. Assume, on the other hand, that the test accepts with probability greater than δ . Then, for some ρ and $\epsilon' = \epsilon/13$, it must be the case that all but an ϵ' fraction of the vertices in G have degree $(\rho \pm \epsilon')N$. By omitting/adding edges to the few vertices with degree outside the above interval, we obtain a graph G' so that

- (1) $\text{dist}(G, G') \leq \epsilon'$.
- (2) every vertex in G' has degree $(\rho \pm 2\epsilon')N$.

At this point, we invoke a theorem due to Noga Alon (see Appendix D) which asserts that a graph G' in which the difference between the maximum and minimum degree as bounded above by $\epsilon''|V(G')|$ is at most $(3 + o(1)) \cdot \epsilon''$ -away from the class of regular graphs. Thus, G is at most ϵ -away from this class, and the proposition follows. $\square$

Estimating vertex degree also suffices to test that the minimum cut in the graph is above some threshold. That is,

PROPOSITION 10.2.1.4 (TESTING MIN-CUT). *The class of graphs G with minimum cut at least $\kappa = \kappa(|V(G)|)$ can be tested by using at most $\text{poly}(\epsilon^{-1} \log(1/\delta))$ queries, where ϵ and δ are the distance and confidence parameters.*

PROOF. If $\epsilon = O(\log(N)/N)$, then we examine the entire graph. Otherwise, we merely test via a $\text{poly}(1/\epsilon)\log(1/\delta)$ sample that all vertices seem to have degree (approximately) above $\kappa = \kappa(N)$. That is, to test that the minimum cut is at least κ , we sample sufficiently many vertices, approximate their degree according to the sample, and accept iff all estimated degrees are above, say, $\kappa - (\epsilon/2)N$. The analysis utilizes the observation that at most $O(N \log N)$ edges

must be added to an N -vertex graph of minimum degree d in order to make it have min-cut at least d . This observation is proved by a random construction.

The basic idea is to consider the immediate neighborhoods of each of the N vertices in the graph. We get a collection of subsets, each having cardinality at least d . Thus, all that is needed is to guarantee d edge-disjoint paths between each pair of such subsets. This can be done easily, by designating d special vertices in the graph, and randomly connecting each vertex in the graph to the designated set by $O(\log N)$ random edges. Consider one specific neighborhood (out of the N). With probability greater than $1 - (1/N)$, the random edges (from its vertices to the designated set) contain a d -matching. Thus, we obtain d edge-disjoint paths between each pair of neighborhoods, which implies that the augmented graph is d -edge-connected. $\square$

10.2.2. Testing Graph Properties Using the Regularity Lemma. As noted above, much less efficient testers for k -Colorability and other graph properties can be obtained by using the Regularity Lemma of Szemerédi [1978]. Interestingly, the Regularity Lemma yields the following result about testing, which we do not know to obtain without it. Let H be an arbitrary fixed graph (e.g., the triangle K_3) and consider the class of graphs which have no H subgraphs. Using the Regularity Lemma, Noga Alon (private communication) observed that there exist testers for H -freeness, with query complexity which is a tower of $\text{poly}(1/\epsilon)$ exponents. Recall that ϵ is our distance parameter (i.e., Alon's tester rejects a graph if it is ϵ -far from being H -free). Alon expressed the opinion that proving a result like this without the Regularity Lemma (and hence probably getting better bounds) would be, indeed, very challenging, and would probably have some very nice combinatorial applications.

10.2.3. Hard to Test Graph Properties. Analogously to Proposition 4.1.1, we show that there are graph properties requiring inspection of a constant fraction of all possible vertex-pairs.

PROPOSITION 10.2.3.1. *There exists a class of graphs, $\mathcal{G}$, for which any testing algorithm must inspect a constant fraction of the vertex pairs. This holds even for testing with respect to the uniform distributions, for any distance parameter $\epsilon < 1/2$ and confidence parameter $\delta < 1/2$, and when allowing the algorithm to make queries and use unlimited computing time.*

PROOF. In adapting the proof of Proposition 4.1.1, we introduce for each N a random subset of $2^{(1/20)N^2}$ N -vertex graphs. Each graph is specified by the lower triangle of the corresponding adjacency matrix. This allows at most $N!$ representations of the same graph (i.e., all its automorphism). The multiple representation only affects the first part of the proof; that is, the bound on the probability that a uniformly selected graph is ϵ -close to $\mathcal{G}$. However, the extra factor of $N!$ is easily eliminated by the probability $\exp(-\Omega(N^2))$ that a random graph is ϵ -close to a specific graph. The second part of the proof (i.e., the distance between the two observed distributions) remains almost unchanged. $\square$

Actually, a similar result holds with respect to graph properties which are in $\mathcal{NP}$; that is, classes of graphs that constitute $\mathcal{NP}$ sets.

PROPOSITION 10.2.3.2. *There exists an $\mathcal{NP}$ set of graphs, $\mathcal{G}$, for which any testing algorithm must inspect at least $\Omega(N^2)$ of the vertex pairs, where N is the number of*

vertices in the graph. This holds even for testing with respect to the uniform distributions, for any constant distance parameter $\epsilon < 1/2$ and confidence parameter $\delta < 1/2$, and when allowing the algorithm to make queries and use unlimited computing time.

This proposition subsumes Proposition 10.2.3.1.

PROOF. We adapt the proof of Proposition 10.2.3.1, by considering, for each N , all graphs which arise for particular “pseudorandom” sequences. Specifically, we consider $\binom{N}{2}$ -long sequences taken from a $1/2 \cdot 2^{-t}$ -biased sample space (cf., Naor and Naor [1993] or Alon et al. [1992]), where $t \stackrel{\text{def}}{=} (1/100)N^2$. Efficiently constructible sample spaces of size $(2^t \cdot N)^5$ having the above property can be found in Naor and Naor [1993] and Alon et al. [1992]. Graphs are now specified, as before, by letting each such sequence define (the lower triangle of) the corresponding adjacency matrix, and so there are $\exp((1/20)N^2)$ such graphs and each graph is specified by a sequence of length $O(t + \log N) = \text{poly}(N)$. The first part of the proof remains unchanged (since all that matters is the number of graphs in the class). The second part of the proof is actually simplified since any t observed bits in the random sequence as above deviates from the uniform distribution by at most $1/2$ (i.e., using the notation of Proposition 4.1, $\delta_S(\mathcal{G}) < 1/2$ for every S of size t). All which remains is to be convinced that we have constructed an NP set. This follows by letting the NP-witness of the membership of a graph in the set be the isomorphism to the canonical representation (i.e., the representation corresponding to the almost unbiased sequence). $\square$

Appendix A. Recurring Notation

For an integer b , we let $[b] \stackrel{\text{def}}{=} \{1, \dots, b\}$.

For any three rational numbers, α , β , and γ , we let $\alpha = \beta \pm \gamma$ stand for $\beta - \gamma \leq \alpha \leq \beta + \gamma$.

A graph is typically denoted $G = (V, E)$ and N typically denotes the number of vertices in G . Furthermore, typically $V = \{1, 2, \dots, N\}$.

Appendix B. Useful Inequalities

Below are several important inequalities that are used throughout the paper.

THEOREM B1 (MARKOV’S INEQUALITY). *Let X be a random variable assuming only non-negative values. Then for all $k \in \mathbb{R}^+$,*

$$\Pr[X \geq k \cdot \text{Exp}[X]] \leq \frac{1}{k}.$$

THEOREM B2 (CHERNOFF BOUNDS). *Let $X_1, X_2, \dots, X_m$ be m independent random variables where $X_i \in [0, 1]$. Let $p \stackrel{\text{def}}{=} (1/m) \sum_i \text{Exp}[X_i]$. Then, for every $\gamma \in [0, 1]$, the following bounds hold:*

—(Additive Form [Hoeffding 1963])

$$\Pr \left[\frac{1}{m} \cdot \sum_{i=1}^m X_i > p + \gamma \right] < \exp(-2\gamma^2 m)$$

and

$$\Pr \left[\frac{1}{m} \cdot \sum_{i=1}^m X_i < p - \gamma \right] < \exp(-2\gamma^2 m).$$

—(Multiplicative Form [Chernoff 1952])

$$\Pr \left[\frac{1}{m} \cdot \sum_{i=1}^m X_i > (1 + \gamma) \cdot p \right] < \exp \left(-\frac{\gamma^2 p m}{3} \right)$$

and

$$\Pr \left[\frac{1}{m} \cdot \sum_{i=1}^m X_i < (1 + \gamma) \cdot p \right] < \exp \left(-\frac{\gamma^2 p m}{3} \right)$$

Appendix C. Determining Closeness to an Admissible Set

In this appendix, we consider the computational problem of determining whether a sequence of densities is $2\epsilon'$ -close to an admissible set for Φ . That is,

Input:

- parameters $\rho_1^{\text{LB}}, \dots, \rho_k^{\text{LB}}, \rho_1^{\text{UB}}, \dots, \rho_k^{\text{UB}}, \varrho_{1,1}^{\text{LB}}, \dots, \varrho_{k,k}^{\text{LB}}, \varrho_{1,1}^{\text{UB}}, \dots, \varrho_{k,k}^{\text{UB}}$, and ϵ' .
- densities: $\rho_1, \dots, \rho_k$ and $\varrho_{1,1}, \dots, \varrho_{k,k}$.

Question. Does the following system of inequalities in x_i 's and $y_{i,j}$'s have a solution?

$$\sum_{i=1}^k x_i = 1 \quad \text{and} \quad \sum_{i,j=1}^k y_{i,j} \leq 1 \quad (54)$$

$$\rho_i^{\text{LB}} \leq x_i \leq \rho_i^{\text{UB}} (\forall i) \quad \text{and} \quad \varrho_{i,j}^{\text{LB}} \leq y_{i,j} \leq \varrho_{i,j}^{\text{UB}} (\forall i, j) \quad (55)$$

$$y_{i,i} \leq x_i^2 (\forall i) \quad \text{and} \quad y_{i,j} \leq 2 \cdot x_i \cdot x_j (\forall i \neq j) \quad (56)$$

$$|x_i - \rho_i| \leq 2\epsilon' \quad \text{and} \quad |y_{i,j} - \varrho_{i,j}| \leq 2\epsilon' (\forall i, j). \quad (57)$$

We first observe that the corresponding lower and upper bounds in Eqs. (55) and (57) can be combined. We also observe that the above system has a solution if and only if it has a solution in which the $y_{i,j}$ are set to be as small as possible. That is, a solution in which each $y_{i,j}$ has the minimum value that obeys the lower bounds in Eqs. (55) and (57). This yields the following system of inequalities,

where $L_i = \max\{\rho_i^{\text{LB}}, \rho_i - 2\epsilon'\}$, $U_i = \min\{\rho_i^{\text{UB}}, \rho_i + 2\epsilon'\}$, $L_{i,i} = \max\{\varrho_{i,i}^{\text{LB}}, \varrho_{i,i} - 2\epsilon'\}$, and $L_{i,j} = 1/2 \cdot \max\{\varrho_{i,j}^{\text{LB}}, \varrho_{i,j} - 2\epsilon'\}$ ($i \neq j$):

$$\sum_{i=1}^k x_i = 1 \quad (58)$$

$$L_i \leq x_i \leq U_i \quad (\forall i) \quad (59)$$

$$x_i \cdot x_j \geq L_{i,j} \quad (\forall i, j) \quad (60)$$

We first observe that Eqs. (58)–(60) constitute a Convex Program; furthermore, its feasibility region, in case it is not empty, is a t -dimensional convex set, where $t \leq k - 1$. Next, we observe that this convex set contains any simplex defined by $t + 1$ points of general position inside the convex set. This holds, in particular, for points which are on the intersection of t of the boundaries/inequalities (i.e., “vertices” of the body). It can be easily verified that for any such two points and for any coordinate, if the points are different along this coordinate, then their difference is bounded below by $2^{-k \cdot L}$, where L is the length of the encoding (in binary) of Φ and ϵ . It follows that the feasibility region, if not empty, contains a t -dimensional ball of radius $2^{-\text{poly}(k) \cdot L}$ (and is contained in $[0, 1]^t$). Thus, the feasibility problem can be solved by exhaustive search in $\exp(\text{poly}(k) \cdot L)$ -time: First, we reduce the problem to t dimensions, by guessing $k - t$ (“independent”) inequalities which are satisfied at equality. Next, we search the resulting t -dimensional space for a feasible solution, by examining all points which reside on a cubic integer lattice spanned by vectors of length $2^{-\text{poly}(k) \cdot L}$.

Remark C1. The feasibility problem may be solvable by the Ellipsoid Method (cf., Grötschel et al. [1988]) in $\text{poly}(k \cdot L)$ -time, but this saving has little effect on our application.

Appendix D. A Note on Regular Graphs (by Noga Alon)

The following theorem is due to Noga Alon. We stress that the theorem refers to simple undirected graphs (i.e., with no self-loops and no parallel edges).

THEOREM D1. *Let $G = (V, E)$ be a graph on N vertices with maximum degree D and minimum degree d , where $D - d \leq \epsilon N$. Then, there is a regular graph on N vertices obtained from G by omitting and adding at most $3\epsilon N^2 + 2N$ edges.*

PROOF. By replacing G , if needed, with its complement, we may and will assume that $D \leq (1/2)(1 + \epsilon)N$. It is convenient to first reduce to the case in which the maximum degree is a bit smaller than $0.5N$. This can be done as follows. Delete the maximum possible number of edges from G subject to keeping its minimum degree d . Let us denote the resulting graph by G' and the maximum degree in it by D' . In case $D' = d$ we are done. Otherwise, we consider the set of vertices of degree D' and observe that it is an independent set (since the existence of an edge between two vertices of degree greater than d violates the definition of G'). Thus, for each set A of vertices of degree D' in G' , $|\Gamma(A)| > |A|$, since the number of edges from A to $\Gamma(A)$ equals $|A| \cdot D'$ as well as is bounded above by $|\Gamma(A)| \cdot (D' - 1)$. Therefore, by Hall’s theorem, there

is a matching in G' that saturates all vertices of degree D' . Omitting such a matching, we obtain a graph, denoted G_1 , of maximum degree $D_1 = D' - 1$ and minimum degree $d_1 \in \{d, d - 1\}$. Iterating this procedure (of first omitting edges between vertices of maximum degree and then omitting an appropriate matching) we obtain a sequence of graphs $G_1, G_2, \dots, G_t$, stopping if either G_t is regular or if $t = (\epsilon/2)N + 2$. Let $(D_1, d_1), \dots, (D_t, d_t)$ be the corresponding sequence of maximum and minimum degrees in the G_i s. Clearly, $D_t \leq D - t$, and $D_t - d_t \leq \epsilon N$. Since $d_t \geq d - t$, we conclude that during this process we have omitted at most $(D - (d - t)) \cdot N \leq 1.5\epsilon N^2 + 2N$ edges. In case G_t is regular, we are done. Otherwise, we have $D_t \leq 0.5N - 2$. We let $H \stackrel{\text{def}}{=} G_t$, $D' \stackrel{\text{def}}{=} D_t$, $d' \stackrel{\text{def}}{=} d_t$, and let D'' be the smallest even integer which is at least D' (i.e., $D'' = D' + (\lceil D'/2 \rceil - \lfloor D'/2 \rfloor) \in \{D', D' + 1\}$).

To complete the proof we show how to modify H by omitting and adding to it at most $O(\epsilon n^2)$ edges so that the resulting graph will be D'' -regular. For each vertex v of H , define $s(v) = D'' - d(v)$, where $d(v)$ is the degree of v in H . Thus the sum $S = \sum_{v \in V} s(v)$ is even, and we have to increase the degree of each vertex v by $s(v)$. We do so in $S/2$ steps, where in each step we increase by 1 the degrees of two (not necessarily distinct) vertices with positive s values, and keep the other degrees invariant. Specifically, in each step, we either add one edge or add two edges and remove one edge. In either case, we update the relevant s values. Since $S \leq (D'' - d')N \leq \epsilon N^2$ the desired result follows. Following is a description of a typical step:

Case 1. There is a vertex v for which $s(v) \geq 2$. Let Y denote the set of all its non-neighbors. If some member of Y has a positive s value, we connect it to v and update their s values to complete the step. Else, each member of Y has degree D'' . Note that $|Y| > N/2$ (since $|Y| \geq N - (D'' - 2) \geq N - D' + 1 \geq 0.5N + 3$). Hence, there must be an edge between two members of Y , since otherwise $|Y| \cdot D'' \leq (N - |Y|) \cdot D''$. Let (u, w) be such an edge. Then we omit (u, w) from the graph, add the two edges (v, u) and (v, w) to the graph, update the s value of v , and complete the step.

Case 2. The only vertices v with positive s values have $s(v) = 1$. Since the sum of the s -values is even, there are at least two such vertices, say, u and v . Let Y be the set of all non-neighbors of u and Z the set of all non-neighbors of v . Note that, as before, if some vertex in either Y or Z has a positive s value, then we can connect it to either u or v and complete the step. We now claim that there must be an edge yz in the graph with $y \in Y$ and $z \in Z$, since otherwise all edges from Y lead to vertices in $V \setminus Z$, implying that $|Y| \cdot D'' \leq (N - |Z|) \cdot D''$, which is impossible, since both $|Y|$ and $|Z|$ are greater than $N/2$. We can now add the edges uy and zv , remove the edge yz , and update the s values of u and v , completing the step.

Observing the the maximum number of edge modifications in these $S/2$ steps is $(3/2)S \leq (3/2)\epsilon N^2$, the theorem follows. $\square$

ACKNOWLEDGMENTS. We wish to thank Noga Alon, Michel Goemans, Ravi Kannan, David Karger and Madhu Sudan for useful discussions. We are also grateful to two anonymous referees for their helpful comments.

REFERENCES

- ALON, N., DUKE, R. A., LEFMAN, H., RODL, V., AND YUSTER, R. 1994. The algorithmic aspects of the regularity lemma. *J. Algorithms* 16, 80–109.
- ALON, N., GOLDBREICH, O., HÅSTAD, J., AND PERALTA, R. 1992. Simple constructions of almost k -wise independent random variables. *J. Rand. Struct. Algorithms* 33, 289–304.
- ANGLUIN, D. 1978. On the complexity of minimum inference of regular sets. *Inf. Cont.* 39, 337–350.
- ARORA, S., FRIESE, A., AND KAPLAN, H. 1996. A new rounding procedure for the assignment problem with applications to dense graph arrangement problems. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, New York, pp. 21–30.
- ARORA, S., KARGER, D., AND KARPINSKI, M. 1995. Polynomial time approximation schemes for dense instances of NP-hard problems. In *Proceedings of the 27th Annual ACM Symposium on the Theory of Computing* (Las Vegas, Nev., May 29–June 1). ACM, New York, pp. 284–293.
- ARORA, S., LUND, C., MOTWANI, R., SUDAN, M., AND SZEGEDY, M. 1992. Proof verification and intractability of approximation problems. In *Proceedings of the 33rd Annual IEEE Symposium on Foundations of Computer Science*. IEEE, New York, pp. 14–23.
- ARORA, S., AND SAFRA, S. 1992. Probabilistic checkable proofs: A new characterization of NP. In *Proceedings of the 33rd Annual IEEE Symposium on Foundations of Computer Science*. IEEE, New York, pp. 1–13.
- BABAI, L., FORTNOW, L., LEVIN, L. A., AND SZEGEDY, M. 1991a. Checking computations in polylogarithmic time. In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing* (New Orleans, La., May 6–8). ACM, New York, pp. 21–31.
- BABAI, L., FORTNOW, L., AND LUND, C. 1991b. Non-deterministic exponential time has two-prover interactive protocols. *Computat. Complex.* 1, 1, 3–40.
- BELLARE, M., COPPERSMITH, D., HÅSTAD, J., KIWI, M., AND SUDAN, M. 1995a. Linearity testing in characteristic two. In *Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, New York, pp. 432–441.
- BELLARE, M., GOLDBREICH, O., AND SUDAN, M. 1995b. Free bits, PCPs, and nonapproximability—Towards tight results. In *Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, New York, pp. 422–431. (Full version available from ECCC, <http://www.eccc.uni-trier.de/eccc/>).
- BELLARE, M., GOLDWASSER, S., LUND, C., AND RUSSELL, A. 1993. Efficient probabilistically checkable proofs and applications to approximation. In *Proceedings of the 25th Annual ACM Symposium on the Theory of Computing* (San Diego, Calif., May 16–18). ACM, New York, pp. 294–304. (Erratum in *Proceedings of the 26th Annual ACM Symposium on the Theory of Computing*. ACM, New York, 1994, p. 820.)
- BELLARE, M., AND SUDAN, M. 1994. Improved non-approximability results. In *Proceedings of the 26th Annual ACM Symposium on the Theory of Computing* (Montréal, Que., Canada, May 23–25). ACM, New York, pp. 184–193.
- BEN-DAVID, S. 1992. Can finite samples detect singularities of real-valued functions? In *Proceedings of the 24th Annual ACM Symposium on the Theory of Computing* (Victoria, B.C., Canada, May 4–6). ACM, New York, pp. 390–399.
- BLUM, A., AND RIVEST, R. 1989. Training a 3-node neural network is NP-complete. In *Advances in Neural Information Processing Systems I*. Morgan-Kaufmann, San Mateo, Calif., pp. 494–501.
- BLUM, M., LUBY, M., AND RUBINFELD, R. 1993. Self-testing/correcting with applications to numerical problems. *J. Comput. Syst. Sci.* 47, 549–595.
- BLUMER, A., EHRENFEUCHT, A., HAUSSLER, D., AND WARMUTH, M. K. 1989. Learnability and Vapnik–Chervonenkis dimension. *J. ACM* 36, 4 (Oct.), 929–965.
- CANETTI, R., FEIGE, U., GOLDBREICH, O., AND NAOR, M. 1996. Adaptively secure multi-party computation. Tech Rep. TR-682. Laboratory of Computer Science, Massachusetts Institute of Technology, Cambridge, Mass. Extended abstract in *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing* (Philadelphia, Pa., May 22–24) ACM, New York, pp. 639–648.
- CHERNOFF, H. 1952. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Ann. Math. Stat.* 23, 493–507.
- COVER, T. M. 1973. On determining the rationality of the mean of a random variable. *Ann. Stat.* 1, 862–871.

- DE LA VEGA, W. F. 1994. MAX-CUT has a randomized approximation scheme in dense graphs. *Rand. Struct. and Algorithms* 8, 4, 187–198.
- EDWARDS, K. 1986. The complexity of colouring problems on dense graphs. *Theoret. Comput. Sci.* 43, 337–343.
- ERGUN, F., KANNAN, S., KUMAR, S. R., RUBINFELD, R., AND VISWANATHAN, M. 1998. Spot-checkers. In *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing*. ACM, New York, to appear.
- FEIGE, U., GOLDWASSER, S., LOVÁSZ, L., SAFRA, S., AND SZEGEDY, M. 1991. Approximating Clique is almost NP-complete. In *Proceedings of the 32nd Annual IEEE Symposium on Foundation of Computer Science*. IEEE, New York, pp. 2–12.
- FRIEZE, A., AND KANAN, R. 1996. The regularity lemma and approximation schemes for dense problems. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, New York, pp. 12–20.
- GEMMELL, P., LIPTON, R., RUBINFELD, R., SUDAN, M., AND WIGDERSON, A. 1991. Self-testing/correcting for polynomials and for approximate functions. In *Proceedings of the 23rd Annual ACM Symposium on the Theory of Computing* (New Orleans, La., May 6–8). ACM, New York, pp. 32–42.
- GOLD, M. E. 1978. Complexity of automation identification from given data. *Inf. Cont.* 37, 302–320.
- GOLDREICH, O. 1995. Foundations of Cryptography—Fragments of a Book. Available from ECCC, <http://www.eccc.uni-trier.de/eccc/>.
- GOLDREICH, O., AND RON, D. 1997. Property testing in bounded degree graphs. In *Proceedings of the 29th Annual ACM Symposium on the Theory of Computing* (El Paso, Tex., May 4–6). ACM, New York, pp. 406–415.
- GOLDREICH, O., AND RON, D. 1998. A sublinear bipartite tester for bounded degree graphs. In *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing*. ACM, New York, to appear.
- GRÖTSCHEL, M., LOVÁSZ, L., AND SCHRIJVER, A. 1988. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, New York.
- HAJNAL, P. 1991. An $\Omega(n^{4/3})$ lower bound on the randomized complexity of graph properties. *Combinatorica* 11, 2, 131–144.
- HÅSTAD, J. 1996a. Testing of the long code and hardness for clique. In *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing* (Philadelphia, Pa., May 22–24). ACM, New York, pp. 11–19.
- HÅSTAD, J. 1996b. Clique is hard to approximate within $n^{1-\epsilon}$. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, New York, pp. 627–636.
- HÅSTAD, J. 1997. Some optimal inapproximability results. In *Proceedings of the 29th Annual ACM Symposium on the Theory of Computing* (El Paso, Tex., May 4–6). ACM, New York, pp. 1–10.
- HOCHBAUM, D. S., AND SHMOYS, D. B. 1987. Using dual approximation algorithms for scheduling problems: Theoretical and practical results. *J. ACM* 34, 1 (Jan.), 144–162.
- HOCHBAUM, D. S., AND SHMOYS, D. B. 1988. A polynomial approximation scheme for machine scheduling on uniform processors: Using the dual approximation approach. *SIAM J. Comput.* 17, 3, 539–551.
- HOEFFDING, W., AND WOLFOWITZ, J. 1958. Distinguishability of sets of distributions. *Ann. Math. Stat.* 29, 700–718.
- KARGER, D. R., MOTWANI, R., AND SUDAN, M. 1994. Approximate graph coloring by semidefinite programming. In *Proceedings of the 35th Annual IEEE Symposium on the Foundation of Computer Science*. ACM, New York, pp. 2–13.
- KEARNS, M. J., MANSOUR, Y., RON, D., RUBINFELD, R., SCHAPIRE, R. E., AND SELLIE, L. 1994. On the learnability of discrete distributions. In *Proceedings of the 25th Annual ACM Symposium on the Theory of Computing* (Montréal, Que., Canada, May 23–25). ACM, New York, pp. 273–282.
- KEARNS, M., AND RON, D. 1998. Testing problems with sub-learning sample complexity. In *Proceedings of the 11th Annual ACM Symposium on Computational Learning Theory*. ACM, New York, to appear.
- KEARNS, M. J., SCHAPIRE, R. E., AND SELLIE, L. M. 1992. Toward efficient agnostic learning. In *Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory* (Pittsburgh, Pa., July 27–29). ACM, New York, pp. 341–351.
- KING, V. 1991. An $\Omega(n^{5/4})$ lower bound on the randomized complexity of graph properties. *Combinatorica* 11, 1, 23–32.

- KIWI, M. 1996. Probabilistically Checkable Proofs and the Testing of Hadamard-like Codes. Ph.D. dissertation. Massachusetts Institute of Technology, Cambridge, Mass.
- KULKARNI, S. R., AND ZEITOUNI, O. 1993. On probably correct classification of concepts. In *Proceedings of the 6th Annual ACM Conference on Computational Learning Theory* (Santa Cruz, Ca., July 26–28). ACM, New York, pp. 111–116.
- LOVASZ, L., AND YOUNG, N. 1991. Lecture notes on evasiveness of graph properties. Tech. Rep. TR-317-91. Computer Science Department, Princeton Univ., Princeton, N.J.
- NAOR, J., AND NAOR, M. 1993. Small-bias probability spaces: Efficient constructions and applications. *SIAM J. Comput.* 22, 4, 838–856.
- PETRANK, E. 1994. The hardness of approximations: Gap location. *Computat. Complex.* 4, 133–157.
- PITT, L., AND VALIANT, L. G. 1988. Computational limitations on learning from examples. *J. ACM* 35, 4 (Oct.), 965–984.
- PITT, L., AND WARMUTH, M. K. 1993. The minimum consistent DFA problem cannot be approximated within a polynomial. *J. ACM* 40, 1 (Jan.), 95–142.
- ROSENBERG, A. L. 1973. On the time required to recognize properties of graphs: A problem. *SIGACT News* 5, 15–16.
- RUBINFELD, R. 1994. Robust functional equations and their applications to program testing. In *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, New York, pp. 288–299.
- RUBINFELD, R., AND SUDAN, M. 1996. Robust characterization of polynomials with applications to program testing. *SIAM J. Comput.* 25, 2, 252–271.
- SCHWARTZ, J. T. 1980. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM* 27, 4, (Oct.), 701–717.
- SZEMERÉDI, E. 1978. Regular partitions of graphs. In *Proceedings of the Colloquium International CNRS*. pp. 399–401.
- TREVISAN, L. 1998. Recycling queries in PCPs in linearity tests. In *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing*, ACM, New York, to appear.
- VALIANT, L. G. 1984. A theory of the learnable. *Commun. ACM* 27, 11 (Nov.), 1134–1142.
- VAPNIK, V. N., AND CHERVOENKIS, A. Y. 1971. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Prob. Applic.* 17, 2, 264–280.
- YAMANISHI, K. 1995. Probably almost discriminative learning. *Mach. Learn.* 18, 23–50.
- YAO, A. C. C. 1987. Lower bounds to randomized algorithms for graph properties. In *Proceedings of the 28th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, New York, pp. 393–400.
- ZEITOUNI, O., AND KULKARNI, S. R. 1995. A general classification rule for probability measures. *Ann. Stat.* 23, 1393–1407.

RECEIVED NOVEMBER 1996; REVISED DECEMBER 1997; ACCEPTED FEBRUARY 1998