

Testing Basic Boolean Formulae *

Michal Parnas[†]

Dana Ron[‡]

Alex Samorodnitsky[§]

Abstract

We consider the problem of determining whether a given function $f : \{0,1\}^n \rightarrow \{0,1\}$ belongs to a certain class of Boolean functions \mathcal{F} or whether it is *far* from the class. More precisely, given query access to the function f and given a distance parameter ϵ , we would like to decide whether $f \in \mathcal{F}$ or whether it differs from every $g \in \mathcal{F}$ on more than an ϵ -fraction of the domain elements. The classes of functions we consider are singleton (“dictatorship”) functions, monomials, and monotone DNF functions with a bounded number of terms. In all cases we provide algorithms whose query complexity is independent of n (the number of function variables), and linear in $1/\epsilon$.

1 Introduction

The newly founded country of Eff is interested in joining the international organization Pea. This organization has one rule: It does not admit dictatorships. Eff claims it is not a dictatorship but is unwilling to reveal the procedure by which it combines the votes of its government members into a final decision. However, it agrees to allow Pea’s special envoy, Tee, to perform a small number of experiments with its voting method. Namely, Tee may set the votes of the government members (using Eff’s advanced electronic system) in any possible way, and obtain the final decision given these votes. Tee’s mission is not to actually identify the dictator among the government members (if such exists), but only to discover *whether* such a dictator exists. Most importantly, she must do so by performing as few experiments as possible. Given this constraint, Tee may decline Eff’s request to join Pea even if Eff is not exactly a dictatorship but only behaves like one most of the time.

The above can be formalized as a *Property Testing Problem*: Let $f : \{0,1\}^n \rightarrow \{0,1\}$ be a fixed but unknown function, and let \mathcal{P} be a fixed property of functions. We would like to determine, by querying f , whether f has the property \mathcal{P} , or whether it is ϵ -*far* from having the property for a given distance parameter ϵ . By ϵ -*far* we mean that more than an ϵ -fraction of its values should be modified so that it obtains the property \mathcal{P} . For example, in the above setting we would like to

*An extended abstract of this work (under the slightly different title: “Proclaiming Dictators and Juntas or Testing Boolean Formulae”) appeared in the proceedings of RANDOM02. The results in this version of the paper improve on those presented in RANDOM02.

[†]The Academic College of Tel-Aviv-Yaffo, 4 Antokolsky st., Tel-Aviv, Israel. Tel. 972-3-5211864, Fax. 972-3-5211871. E-mail michalp@mta.ac.il.

[‡]Department of EE – Systems, Tel-Aviv University, Ramat Aviv, Israel. Tel. 972-3-6406917, Fax. 972-3-6407095. E-mail danar@eng.tau.ac.il. Supported by the Israel Science Foundation (grant number 32/00-1).

[§]School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem, Israel. Tel. 972-2-6586176, Fax. 972-2-6585439. E-mail salex@cs.huji.ac.il. Work done while at the Institute for Advanced Study, Princeton, NJ. Supported by NSF (grant number CCR-9987845) and by a State of New Jersey grant.

test whether a given function f is a “dictatorship function”. That is, whether there exists an index $1 \leq i \leq n$, such that $f(x) = x_i$ for every $x \in \{0, 1\}^n$.

Previous work on testing properties of functions mainly focused on algebraic properties (e.g., [BLR93, RS96, Rub99]), or on properties defined by relatively rich families of functions such as the family of all monotone functions [GGL⁺00, DGL⁺99]. Here we are interested in studying the most basic families of Boolean functions: singletons, monomials, and DNF functions.

One possible approach to testing whether a function f has a certain property \mathcal{P} , is to try and actually *find* a good approximation for f from within the family of functions $\mathcal{F}_{\mathcal{P}}$ having the tested property \mathcal{P} . For this task we would use a *learning algorithm* that performs queries and works under the uniform distribution. Such an algorithm ensures that if f has the property (that is, $f \in \mathcal{F}_{\mathcal{P}}$), then with high probability the learning algorithm outputs a *hypothesis* $h \in \mathcal{F}_{\mathcal{P}}$ such that $\Pr[f(x) \neq h(x)] \leq \epsilon$, where ϵ is a given distance (or error) parameter. The testing algorithm would run the learning algorithm, obtain the hypothesis $h \in \mathcal{F}_{\mathcal{P}}$, and check that h and f in fact differ only on a small fraction of the domain. This last step is performed by taking a sample of size $\Theta(1/\epsilon)$ from $\{0, 1\}^n$ and comparing f and h on the sample. Thus, if f has the property \mathcal{P} then it will be accepted with high probability, and if f is ϵ -far from having \mathcal{P} , so that $\Pr[f(x) \neq h(x)] > \epsilon$ for every $h \in \mathcal{F}_{\mathcal{P}}$, then it will be rejected with high probability.

Hence, provided that there exists a learning algorithm for the tested family $\mathcal{F}_{\mathcal{P}}$, we obtain a testing algorithm whose complexity is of the same order of that of the learning algorithm. To be more precise, the learning algorithm should be a *proper* learning algorithm. That is, the hypothesis h it outputs must belong to $\mathcal{F}_{\mathcal{P}}$.¹

A natural question that arises is whether we can do better by using a different approach. Recall that we are not interested in actually finding a good approximation for f in $\mathcal{F}_{\mathcal{P}}$, but we only want to know whether such an approximation *exists*. Therefore, perhaps we can design a different and more efficient testing algorithm than the one based on learning. In particular, the complexity measure we would like to improve is the *query complexity* of the algorithm.

As we show below, for all the properties we study, we describe algorithms whose query complexity is linear in $1/\epsilon$, where ϵ is the given distance parameter, and *independent* of the input size n .² As we discuss shortly, the corresponding proper learning algorithms have query complexities that depend on n , though only polylogarithmically. We believe that our results are of interest both because they completely remove the dependence on n in the query complexity, and also because in certain aspects they are inherently different from the corresponding learning algorithms. Hence they may shed new light on the structure of the properties studied.

1.1 Our Results

We present the following testing algorithms:

- An algorithm that tests whether f is a singleton function. That is, whether there exists an index $1 \leq i \leq n$, such that $f(x) = x_i$ for every $x \in \{0, 1\}^n$, or $f(x) = \bar{x}_i$ for every $x \in \{0, 1\}^n$. This algorithm has query complexity $O(1/\epsilon)$.
- An algorithm that tests whether f is a monomial with query complexity $O(1/\epsilon)$.

¹This is as opposed to *non-proper* learning algorithms that given query access to $f \in \mathcal{F}_{\mathcal{P}}$ are allowed to output a hypothesis h that belongs to a more general hypothesis class $\mathcal{F}' \supset \mathcal{F}_{\mathcal{P}}$. Non-proper learning algorithms are not directly applicable for our purposes.

²The running times of the algorithms are all linear in the number of queries performed and in n . This dependence on n in the running time is clearly unavoidable, since even writing down a query takes time n .

- An algorithm that tests whether f is a monotone DNF having at most ℓ terms, with query complexity $\tilde{O}(\ell^2/\epsilon)$.

We note that the above results improve on those presented in an extended abstract of this work [PRS01a].

Techniques Our algorithms for testing singletons and for testing monomials have a similar structure. In particular, they combine two tests. One test is a “natural” test that arises from an exact logical characterization of these families of functions. In the case of singletons this test uniformly selects pairs $x, y \in \{0, 1\}^n$ and verifies that $f(x \wedge y) = f(x) \wedge f(y)$, where $x \wedge y$ denotes the bitwise ‘and’ of the two strings. The corresponding test for monomials performs a slight variant of this test. The other test in both cases is a seemingly less evident test with an algebraic flavor. In the case of singletons it is a linearity test [BLR93] and in the case of monomials it is an affinity test. This test ensures that if f passes it then it has or is close to having a certain structure. This structure aids us in analyzing the logical test. We note that our current analysis of the affinity test differs from the one presented in previous versions of this work [PRS01a, PRS01b]. In particular, in previous versions we used the Discrete Fourier Transform, while here we build on basic probabilistic arguments.

The testing algorithm for monotone DNF functions uses the test for monomials as a sub-routine. Recall that a DNF function is a disjunction of monomials (the terms of the function). If f is a DNF function with a bounded number of monotone terms, then the test will isolate the different terms of the function and test that each is in fact a monotone monomial. If f is far from being such a DNF function, then at least one of these tests will fail with high probability.

1.2 Related Work

1.2.1 Property Testing

Property testing was first defined and applied in the context of algebraic properties of functions [RS96], and has since been extended to various domains, perhaps most notably those of graph properties (e.g. [GGR98, GR97, AFKS99]). (For surveys see [Ron01, Fis01]). The relation between testing and learning is discussed at length in [GGR98]. In particular, that paper suggests that testing may be applied as a preliminary stage to learning. Namely, efficient testing algorithms can be used in order to help in determining what hypothesis class should be used by the learning algorithm.

Linearity Testing and its Variants As noted above, we use linearity testing [BLR93] in our test for singletons, and affinity testing, which can be viewed as an extension of linearity testing, for testing monomials. Other works in which improvements and variants of linearity testing are analyzed include [BCH⁺95, AHR99].

Testing the Long-Code We note that a test which is very similar to our test for singletons was applied to testing the *Long-Code* [BGS98]. Specifically, Bellare, Goldreich and Sudan [BGS98] consider the following task. For an integer ℓ , let \mathcal{F}_ℓ be the set of all Boolean functions over $\{0, 1\}^\ell$. Let G be a function from \mathcal{F}_ℓ to $\{0, 1\}$. The goal is to test whether G has the following property: there exists some fixed $a \in \{0, 1\}^\ell$ such that $G(f) = f(a)$ for every $f \in \mathcal{F}_\ell$. In such a case G is the code-word of length 2^{2^ℓ} corresponding to the plain-text a .

In order to better understand the relation to our problem, we view each string $a \in \{0, 1\}^\ell$ as an index between 0 and $2^\ell - 1$, and each function $f \in \mathcal{F}_\ell$ as a string of length 2^ℓ that corresponds to the truth-table of f . Then the test should accept G if there exists some index a such that for every f , $G(f) = f(a)$. The test should reject f if for every a , $G(f) \neq f(a)$ on more than an ϵ -fraction of the strings (functions) f . In other words, testing the Long-Code is equivalent to testing monotone singletons over $\{0, 1\}^n$ when $n = 2^\ell$. Thus we extend the Long-Code/singletons test to any n , while simplifying the analysis.

Testing k -Juntas Following the publication of the extended abstract of this work [PRS01a], a recent work [FKR⁺02] addresses the problem of testing whether a Boolean function depends on at most k variables, for a given parameter k . In [FKR⁺02] it is shown that this “ k -junta” property can be tested using a number of queries that is linear in $1/\epsilon$ and polynomial in k . It is also noted that the k -junta testing algorithm can be applied as a subroutine to testing monomials using a number of queries that is $\tilde{O}(1/\epsilon)$.

1.2.2 Learning Boolean Formulae

Singletons, and more generally monomials, can be easily learned under the uniform distribution. The learning algorithm uniformly selects a sample of size $\Theta(\log n/\epsilon)$ and queries the function f on all sample strings. It then searches for a monomial that is consistent with f on the sample. Finding a consistent monomial if such exists can be done in time linear in the sample size and in n . A simple probabilistic argument, which is a slight variant of Occam’s Razor [BEHW87]³, can be used to show that a sample of size $\Theta(\log n/\epsilon)$ is sufficient to ensure that with high probability any monomial that is consistent with the sample is an ϵ -good approximation of f .

There is a large variety of results on learning DNF functions, and in particular monotone DNF, in several different models. We restrict our attention to the model most relevant to our work, namely when membership queries are allowed and the underlying distribution is uniform. The best known algorithm results from combining the works of [BJT99] and [KS99], and builds on Jackson’s celebrated Harmonic Sieve algorithm [Jac97]. This algorithm has query complexity $\tilde{O}\left(r \cdot \left(\frac{\log^2 n}{\epsilon} + \frac{\ell^2}{\epsilon^2}\right)\right)$, where r is the number of variables appearing in the DNF formula, and ℓ is the number of terms. However, this algorithm does not output a DNF formula as its hypothesis. On the other hand, Angluin [Ang88] describes a proper learning algorithm for monotone DNF formulae that uses membership queries and works under arbitrary distributions. The query complexity of her algorithm is $\tilde{O}(\ell \cdot n + \ell/\epsilon)$. Using the same preprocessing technique as suggested in [BJT99], if the underlying distribution is uniform then the query complexity can be reduced to $\tilde{O}\left(\frac{r \cdot \log^2 n}{\epsilon} + \ell \cdot \left(r + \frac{1}{\epsilon}\right)\right)$. Recall that the query complexity of our testing algorithm has similar dependence on ℓ and $1/\epsilon$ but does not depend on n .

1.3 Organization

We start with some necessary preliminaries in Section 2. In Section 3 we present our algorithm for testing singleton functions. The algorithm for testing monomials is presented in Section 4, and the algorithm for testing monotone DNF in Section 5. In Section 6 we discuss a possible simpler alternative to the singleton test.

³Applying the theorem known as Occam’s Razor would give a stronger result in the sense that the underlying distribution may be arbitrary (that is, not necessarily uniform). This however comes at a price of a linear, as opposed to logarithmic, dependence of the sample/query complexity on n .

2 Preliminaries

We shall use the following definitions.

Definition 1 Let $x, y \in \{0, 1\}^n$, and let $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$.

- We denote by $|x|$ the number of ones in the vector x .
- We write $y \succeq x$ if in each coordinate $y_i \geq x_i$.
- Let $2^x \stackrel{\text{def}}{=} \{z \in \{0, 1\}^n : z \preceq x\}$. Hence, $|2^x| = 2^{|x|}$.
- Let $x \wedge y$ denote the string $z \in \{0, 1\}^n$ such that for every $i \in [n]$, $z_i = x_i \wedge y_i$.
- Let $x \oplus y$ denote the string $z \in \{0, 1\}^n$ such that for every $i \in [n]$, $z_i = x_i \oplus y_i$.

Definition 2 (Singletons, Monomials, and DNF functions) A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a singleton function, if there exists an $i \in [n]$ such that $f(x) = x_i$ for every $x \in \{0, 1\}^n$, or $f(x) = \bar{x}_i$ for every $x \in \{0, 1\}^n$.

We say that f is a monotone k -monomial for $1 \leq k \leq n$ if there exist k indices $i_1, \dots, i_k \in [n]$, such that $f(x) = x_{i_1} \wedge \dots \wedge x_{i_k}$ for every $x \in \{0, 1\}^n$. If we allow some of the x_{i_j} 's above to be replaced with \bar{x}_{i_j} , then f is a k -monomial. The function f is a monomial if it is a k -monomial for some $1 \leq k \leq n$.

A function f is an ℓ -term DNF function if it is a disjunction of ℓ monomials. If all monomials are monotone, then it is a monotone DNF function.

When the identity of the function f is clear from the context, we may use the following notation.

Definition 3 Define $F_0 \stackrel{\text{def}}{=} \{x | f(x) = 0\}$ and $F_1 \stackrel{\text{def}}{=} \{x | f(x) = 1\}$.

Definition 4 (Distance between functions) The distance according to the uniform distribution between two functions $f, g : \{0, 1\}^n \rightarrow \{0, 1\}$ is denoted by $\text{dist}(f, g)$, and is defined as follows: $\text{dist}(f, g) \stackrel{\text{def}}{=} \Pr_{x \in \{0, 1\}^n} [f(x) \neq g(x)]$.

The distance between a function f and a family of functions \mathcal{F} is $\text{dist}(f, \mathcal{F}) \stackrel{\text{def}}{=} \min_{g \in \mathcal{F}} \text{dist}(f, g)$. If $\text{dist}(f, \mathcal{F}) > \epsilon$ for some $0 < \epsilon < 1$, then we say that f is ϵ -far from \mathcal{F} . Otherwise, f is ϵ -close to \mathcal{F} .

Definition 5 (Testing Algorithms) A testing algorithm for a family of boolean functions \mathcal{F} over $\{0, 1\}^n$ is given a distance parameter ϵ , $0 < \epsilon < 1$, and is provided with query access to an arbitrary function $f : \{0, 1\}^n \rightarrow \{0, 1\}$.

If $f \in \mathcal{F}$ then the algorithm must output accept with probability at least $2/3$, and if f is ϵ -far from \mathcal{F} then it must output reject with probability at least $2/3$.

3 Testing Singletons

We start by presenting an algorithm for testing singletons. The testing algorithm for k -monomials will generalize this algorithm. More precisely, we present an algorithm for testing whether a function f is a *monotone* singleton. In order to test whether f is a singleton we can check whether either f or \bar{f} pass the monotone singleton test. For the sake of succinctness, in what follows we refer to monotone singletons simply as singletons.

The following characterization of monotone k -monomials motivates our tests. We later show that the requirement of monotonicity can be removed.

Claim 1 *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$. Then f is a monotone k -monomial if and only if the following two conditions hold:*

1. $\Pr[f = 1] = 1/2^k$;
2. $\forall x, y, f(x \wedge y) = f(x) \wedge f(y)$;

Proof: If f is a k -monomial then clearly the conditions hold. We turn to prove the other direction. We first observe that the two conditions imply that $f(x) = 0$ for all $|x| < k$, where $|x|$ denotes the number of ones in x . In order to verify this, assume in contradiction that there exists some x such that $|x| < k$ but $f(x) = 1$. But now consider any y such that $y_i = 1$ whenever $x_i = 1$. Then $x \wedge y = x$, and therefore $f(x \wedge y) = 1$. But by the second item, since $f(x) = 1$, it must also hold that $f(y) = 1$. However, since $|x| < k$, the number of such points y is strictly greater than 2^{n-k} , contradicting the first item.

Next let $y = \bigwedge_{x \in F_1} x$. Using the second item in the claim we get:

$$f(y) = f\left(\bigwedge_{x \in F_1} x\right) = \bigwedge_{x \in F_1} f(x) = 1.$$

However, we have just shown that $f(x) = 0$ for all $|x| < k$, and thus $|y| \geq k$. Hence, there exist k indices i_1, \dots, i_k such that $y_{i_j} = 1$ for all $1 \leq j \leq k$. But $y_{i_j} = \bigwedge_{x \in F_1} x_{i_j}$. Hence, $x_{i_1} = \dots = x_{i_k} = 1$ for every $x \in F_1$. The first item now implies that $f(x) = x_{i_1} \wedge \dots \wedge x_{i_k}$ for every $x \in \{0, 1\}^n$. ■

Definition 6 *We say that $x, y \in \{0, 1\}^n$ are a violating pair with respect to a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, if $f(x) \wedge f(y) \neq f(x \wedge y)$.*

Given the above definition, Claim 1 states that a basic property of monotone singletons, and more generally of monotone k -monomials, is that there are no violating pairs with respect to f . A natural candidate for a testing algorithm for singletons would take a sample of uniformly selected pairs x, y , and for each pair verify that it is not violating with respect to f . In addition, the test would check that $\Pr[f = 0]$ is roughly $1/2$ (or else any monotone k -monomial would pass the test).

As we discuss in Section 6, we were unable to give a complete proof for the correctness of this test. Somewhat counter-intuitively, the difficulty with the analysis lies in the case when the function f is *very far* from being a singleton. More precisely, the analysis is quite simple when the distance δ between f and the closest singleton is bounded away from $1/2$. However, the argument does not directly apply to δ arbitrarily close to $1/2$. We believe it would be interesting to prove that this simple test is in fact correct (or to come up with an example of a function f that is almost $1/2$ -far from any singleton, but passes the test).

In the algorithm described below we circumvent the above difficulty by “forcing more structure” on f . Specifically, we first perform another test that only accepts functions that have, or more precisely, that are close to having a certain structure. In particular, every singleton will pass the test. We then perform a slight variant of our original test. Provided that f passes the first test, it will be easy to show that f passes the second test with high probability only if it is close to a singleton function. Details follow.

The algorithm begins by testing whether the function f belongs to a larger family of functions that contains singletons as a sub-family. This is the family of *parity functions*.

Definition 7 *A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a parity function (a linear function over $\text{GF}(2)$) if there exists a subset $S \subseteq [n]$ such that $f(x) = \bigoplus_{i \in S} x_i$ for every $x \in \{0, 1\}^n$.*

The test for parity functions is a special case of the linearity test over general fields due to Blum, Luby and Rubinfeld [BLR93]. If the tested function f is a parity function, then the test always accepts, and if f is ϵ -far from any parity function then the test rejects with probability at least $9/10$. The query complexity of this test is $O(1/\epsilon)$. Specifically, the test uniformly picks $O(1/\epsilon)$ pairs $x, y \in \{0, 1\}^n$ and checks that $f(x) \oplus f(y) = f(x \oplus y)$.

Assuming this test passes, we still need to verify that f is actually close to a singleton function and not to some other parity function. If the parity test only accepted proper parity functions, then the following claim would suffice. It shows that if f is a non-singleton parity function, then a constant size sample of pairs x, y would, with high probability, contain a violating pair with respect to f .

Claim 2 *Let $g = \bigoplus_{i \in S} x_i$ for $S \subseteq [n]$. If $|S|$ is even then*

$$\Pr[g(x \wedge y) = g(x) \wedge g(y)] = \frac{1}{2} + \frac{1}{2^{|S|+1}}$$

and if $|S|$ is odd then

$$\Pr[g(x \wedge y) = g(x) \wedge g(y)] = \frac{1}{2} + \frac{1}{2^{|S|}}.$$

Proof: Let $s = |S|$, and let x, y be two strings such that (i) x has $0 \leq i \leq s$ ones in S , that is, $|\{\ell \in S : x_\ell = 1\}| = i$; (ii) $x \wedge y$ has $0 \leq k \leq i$ ones in S ; and (iii) y has a total of $j + k$ ones in S , where $0 \leq j \leq s - i$.

If $g(x \wedge y) = g(x) \wedge g(y)$, then either (1) i is even and k is even, or (2) i is odd and j is even. Let $Z_1 \subset \{0, 1\}^n \times \{0, 1\}^n$ be the subset of pairs x, y that obey the first constraint, and let $Z_2 \subset \{0, 1\}^n \times \{0, 1\}^n$ be the subset of pairs x, y that obey the second constraint. Since the two subsets are disjoint,

$$\Pr[g(x \wedge y) = g(x) \wedge g(y)] = 2^{-2n} \cdot (|Z_1| + |Z_2|) \tag{1}$$

It remains to compute the sizes of the two sets. Since the coordinates of x and y outside S do not determine whether the pair x, y belongs to one of these sets, we have

$$|Z_1| = 2^{n-s} \cdot 2^{n-s} \cdot \left(\sum_{i=0, i \text{ even}}^s \binom{s}{i} \sum_{k=0, k \text{ even}}^i \binom{i}{k} \sum_{j=0}^{s-i} \binom{s-i}{j} \right) \tag{2}$$

and

$$|Z_2| = 2^{n-s} \cdot 2^{n-s} \cdot \left(\sum_{i=0, i \text{ odd}}^s \binom{s}{i} \sum_{k=0}^i \binom{i}{k} \sum_{j=0, j \text{ even}}^{s-i} \binom{s-i}{j} \right) \tag{3}$$

The first expression equals

$$2^{2n-2s} \cdot (2^{2s-2} + 2^{s-1}) = 2^{2n-2} + 2^{2n-s-1} = 2^{2n} \cdot (2^{-2} + 2^{-(s+1)}).$$

The second sum equals $2^{2n} \cdot (2^{-2} + 2^{-(s+1)})$ if s is odd and 2^{2n-2} if s is even. The claim follows by combining Equations (2) and (3) with Equation (1). ■

Hence, if f is a parity function that is not a singleton, that is $|S| \geq 2$, then the probability that a uniformly selected pair x, y is violating with respect to f is at least $1/8$. In this case, a sample of 16 such pairs will contain a violating pair with probability at least $1 - (1 - 1/8)^{16} \geq 1 - e^{-2} > 2/3$.

However, what if f passes the parity test but is only close to being a parity function? Let g denote the parity function that is closest to f and let δ be the distance between them. (Note that g is unique, given that f is sufficiently close to a parity function). What we would like to do is check whether g is a singleton, by selecting a sample of pairs x, y and checking whether it contains a violating pair with respect to g . Observe that, since the distance between functions is measured with respect to the uniform distribution, then for a uniformly selected pair x, y , with probability at least $(1 - \delta)^2$, both $f(x) = g(x)$ and $f(y) = g(y)$. However, we cannot make a similar claim about $f(x \wedge y)$ and $g(x \wedge y)$, since $x \wedge y$ is *not* uniformly distributed. Thus it is not clear that we can replace the violation test for g with a violation test for f . In addition we would like to verify that g is not the all-0 function

The solution is to use a *self-corrector* for linear (parity) functions [BLR93]. Given query access to a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, which is strictly closer than $1/4$ to some parity function g , and an input $x \in \{0, 1\}^n$, the procedure $\text{Self-Correct}(f, x)$ returns the value of $g(x)$, with probability at least $9/10$. The query complexity of the procedure is constant.

The above discussion suggests the following testing algorithm.

Algorithm 1 Test for Singleton Functions

1. Apply the parity test to f with distance parameter $\min(1/5, \epsilon)$. If the parity test rejects then reject.
2. If $\text{Self-Correct}(f, \vec{1}) = 0$ then reject (where $\vec{1}$ is the all-1 vector).
3. Uniformly and independently select $m = 64$ pairs of points x, y .
 - For each such pair, let $b_x = \text{Self-Correct}(f, x)$, $b_y = \text{Self-Correct}(f, y)$ and $b_{x \wedge y} = \text{Self-Correct}(f, x \wedge y)$.
 - Check that $b_{x \wedge y} = b_x \wedge b_y$.
4. If one of the checks fails then reject. Otherwise, accept.

Theorem 1 *Algorithm 1 is a testing algorithm for monotone singletons. Furthermore, it has a one sided error. That is, if f is a monotone singleton, the algorithm always accepts. The query complexity of the algorithm is $O(1/\epsilon)$.*

Proof: Since the testing algorithm for parity functions has a one-sided error, if f is a singleton function then it always passes the test. In this case the self corrector always returns the value of f on every given input point. In particular, $\text{Self-Correct}(f, \vec{1}) = f(\vec{1}) = 1$, since every monotone

singleton has value 1 on the all-1 vector. Similarly, no violating pair can be found in Step 3. Hence, the test always accepts a singleton.

Assume, without loss of generality, that $\epsilon \leq 1/5$. Consider the case in which f is ϵ -far from any singleton. If it is also ϵ -far from any parity function, then it will be rejected with probability at least $9/10$ in the first step of the algorithm. Otherwise, there exists a unique parity function g such that f is ϵ -close to g . If g is the all-0 function, then f is rejected with probability at least $9/10$. Otherwise, g is a parity function of at least 2 variables. By Claim 2, the probability that a uniformly selected pair x, y is a violating pair with respect to g is at least $1/8$. Given such a pair, the probability that the self-corrector returns the value of g on all the three calls (that is, $b_x = g(x)$, $b_y = g(y)$, and $b_{x \wedge y} = g(x \wedge y)$), is at least $(1 - 1/10)^3 > 7/10$. The probability that Algorithm 1 obtains a violating pair with respect to g and all calls to the self-corrector return the correct value, is greater than $1/16$. Therefore, a sample of 64 pairs will ensure that a violation $b_{x \wedge y} \neq b_x \wedge b_y$ will be found with probability at least $9/10$. The total probability that f is accepted, despite being ϵ -far from any singleton, is hence at most $3 \cdot (1/10) < 1/3$.

The query complexity of the algorithm is dominated by the query complexity of the parity tester which is $O(1/\epsilon)$. The second stage takes a constant time. ■

4 Testing Monomials

In this section we describe an algorithm for testing *monotone* k -monomials, where k is provided to the algorithm. We discuss later how to extend this to testing monomials when k is not specified. As for the monotonicity requirement, the following observation and corollary show that this requirement can be easily removed, if desired.

Observation 3 *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and let $z \in \{0, 1\}^n$. Consider the function $f_z : \{0, 1\}^n \rightarrow \{0, 1\}$ that is defined by $f_z(x) = f(x \oplus z)$. Then the following are immediate:*

1. *The function f is a k -monomial if and only if f_z is a k -monomial.*
2. *Let $y \in F_1$. If f is a (not necessarily monotone) k -monomial, then $f_{\bar{y}}$ is a monotone k -monomial.*

Corollary 4 *If f is ϵ -far from every (not necessarily monotone) k -monomial, then for every $y \in F_1$, $f_{\bar{y}}$ is ϵ -far from every monotone k -monomial.*

We next observe that we can also assume without loss of generality that $\epsilon < 2^{-k+2}$, or else the testing problem is trivial.

Observation 5 *Suppose that $\epsilon \geq 2^{-k+2}$. Then:*

1. *If $\Pr[f = 1] \leq \frac{\epsilon}{2}$, then f is ϵ -close to every k -monomial and in particular to every monotone k -monomial.*
2. *If $\Pr[f = 1] > \frac{\epsilon}{4}$, then f is not a k -monomial.*

Proof: If $\Pr[f = 1] \leq \frac{\epsilon}{2}$ then for every k -monomial g ,

$$\text{dist}(f, g) = \Pr[f = 1 \wedge g = 0] + \Pr[f = 0 \wedge g = 1] \leq \Pr[f = 1] + \Pr[g = 1] \leq \frac{\epsilon}{2} + 2^{-k} \leq \epsilon.$$

Since $\epsilon \geq 2^{-k+2}$, if $\Pr[f = 1] > \frac{\epsilon}{4}$ then $\Pr[f = 1] > 2^{-k}$, while by the definition of a k -monomial, $\Pr[f = 1] = 2^{-k}$. ■

By Observation 5, if the algorithm receives parameters ϵ and k such that $\epsilon \geq 2^{-k+2}$, then it simply needs to obtain an estimate α for $p = \Pr[f = 1]$ such that the following holds with probability of at least $2/3$: if $p > \epsilon/2$ then $\alpha > 3\epsilon/8$, and if $p \leq \epsilon/4$ then $\alpha \leq 3\epsilon/8$. By a multiplicative Chernoff bound, such an estimate can be obtained using a sample of size $O(1/\epsilon)$. The algorithm accepts if $\alpha \leq 3\epsilon/8$ and rejects otherwise.

Hence, from this point on we can assume that $\epsilon < 2^{-k+2}$.

We now present the algorithm for testing monotone k -monomials. The first two steps of the algorithm are an attempt to generalize the parity test in Algorithm 1. Specifically, we test whether F_1 is an *affine subspace*.

Definition 8 (Affine Subspaces) *A subset $H \subseteq \{0, 1\}^n$ is an affine subspace of $\{0, 1\}^n$ if and only if there exist an $x \in \{0, 1\}^n$ and a linear subspace V of $\{0, 1\}^n$, such that $H = V \oplus x$. That is,*

$$H = \{y \mid y = v \oplus x, \text{ for some } v \in V\}.$$

The following is a well known alternative characterization of affine subspaces, which is a basis for our test. ⁴

Fact 6 *H is an affine subspace if and only if for every $y_1, y_2, y_3 \in H$ we have $y_1 \oplus y_2 \oplus y_3 \in H$.*

Note that the above fact implies that for every $y_1, y_2 \in H$ and $y_3 \notin H$ we have $y_1 \oplus y_2 \oplus y_3 \notin H$.

Algorithm 2 Test for monotone k -monomials

1. **Size Test:** *Uniformly and independently select a sample of $\Theta(2^k)$ strings in $\{0, 1\}^n$. For each x in the sample, obtain $f(x)$. Let α be the fraction of sample strings x such that $f(x) = 1$. If $|\alpha - 2^{-k}| > 2^{-(k+5)}$ then reject, otherwise continue.*
2. **Affinity Test:**
 - (a) *Set $\delta = 1/36$.*
 - (b) *Uniformly and independently select $m = 2^5/(\epsilon \cdot \delta)$ points $a_1, \dots, a_m \in \{0, 1\}^n$.*
 - (c) *Uniformly and independently select $m' = 4/\delta$ pairs of points $(x_1, y_1), \dots, (x_{m'}, y_{m'}) \in F_1 \times F_1$.*
 - (d) *If for some $1 \leq i \leq m$, $1 \leq j \leq m'$, the equality $f(a_i \oplus x_j \oplus y_j) = f(a_i)$ does not hold, then reject.*

As we show in our analysis, passing this step with sufficiently high probability ensures that f is close to some function g for which $g(x) \oplus g(y) \oplus g(z) = g(x \oplus y \oplus z)$ for all $x, y, z \in G_1 = \{x \mid g(x) = 1\}$. That is, G_1 is an affine subspace.

3. **Closure-Under-Intersection Test:**

(a) Uniformly and independently select 32 points $x \in F_1$.

⁴Here and in most of what follows we use y_i (similarly, x_i) to denote strings in $\{0, 1\}^n$, and not single bits. We find this notation easier to read than the alternative notation y^i . We use the latter only when necessary, that is when we need to refer to particular coordinates y_j^i of the string.

- (b) Uniformly and independently select 2^{k+3} points $y \in \{0, 1\}^n$.
- (c) If for some pair x, y selected, $\text{Self-Correct}(f, x \wedge y) \neq \text{Self-Correct}(f, y)$, then reject. Here Self-Correct is a procedure that given any input z and oracle access to f , asks a constant number of queries and returns with high constant probability, the value $g(z)$, where g is as described in Step 2.

4. If no step caused rejection, then accept.

In both the affinity test and the closure-under-intersection test, we need to select strings in F_1 uniformly. This is simply done by sampling from $\{0, 1\}^n$ and using only x 's for which $f(x) = 1$. Since in both tests the number of strings selected from F_1 is a constant, the total number of queries required is $O(2^k) = O(1/\epsilon)$ (recall that we can assume that $\epsilon < 2^{-k+2}$).

We now embark on proving the correctness of the algorithm.

Theorem 2 *Algorithm 2 is a testing algorithm for monotone k -monomials. The query complexity of the algorithm is $O(1/\epsilon)$.*

The proof of Theorem 2 is based on the following two lemmas whose proofs are provided in Subsections 4.1 and 4.2 respectively.

Lemma 7 *Let f be a function for which $|\Pr[f = 1] - 2^{-k}| < 2^{-k-3}$. If the probability that the affinity test accepts f is greater than $1/10$, then there exists a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ for which the following holds:*

1. $\text{dist}(f, g) \leq \epsilon/2^5$.
2. $G_1 \stackrel{\text{def}}{=} \{a : g(a) = 1\}$ is an affine subspace of dimension $n - k$.
3. There exists a procedure Self-Correct that given any input $a \in \{0, 1\}^n$ and oracle access to f , asks a constant number of queries and returns the value $g(a)$ with probability at least $1 - 1/40$.

Furthermore, if F_1 is an affine subspace then the affinity tests always accepts, $g = f$, and $\text{Self-Correct}(f, a) = f(a)$ with probability 1 for every $a \in \{0, 1\}^n$.

Lemma 8 *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a function for which $|\Pr[f = 1] - 2^{-k}| < 2^{-k-3}$. Suppose that there exists a function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ such that:*

1. $\text{dist}(f, g) \leq 2^{-k-3}$.
2. $G_1 \stackrel{\text{def}}{=} \{x : g(x) = 1\}$ is an affine subspace of dimension $n - k$.
3. There exists a procedure Self-Correct that given any input $a \in \{0, 1\}^n$ and oracle access to f returns the value $g(a)$ with probability at least $1 - 1/40$.

If g is not a monotone k -monomial, then the probability that the Closure-Under-Intersection Test rejects is at least $9/10$.

Proof of Theorem 2: If f is a monotone k -monomial, then $\Pr[f = 1] = 2^{-k}$. By a multiplicative Chernoff bound, for the appropriate constant in the $\Theta(\cdot)$ notation, the probability that it is rejected in the first step of Algorithm 2 is less than a $1/3$. By the definition of k -monomials, f always passes the affinity test and the closure-under-intersection test.

Suppose that f is ϵ -far from any monotone k -monomial. We show that it is rejected with probability greater than $2/3$.

1. If $|\Pr[f = 1] - 2^{-k}| \geq 2^{-k-3}$, then by a multiplicative Chernoff bound, f is rejected in the first step of the algorithm with probability at least $9/10$.
2. Otherwise, $|\Pr[f = 1] - 2^{-k}| < 2^{-k-3}$. If f is $(\epsilon/2^5)$ -far from every function g such that $G_1 = \{x : g(x) = 1\}$ is an affine subspace of dimension $n - k$, then by Lemma 7 it is rejected in the second step of the algorithm (the affinity test) with probability at least $9/10$.
3. Otherwise, both $|\Pr[f = 1] - 2^{-k}| < 2^{-k-3}$ and f is $(\epsilon/2^5)$ -close to a function g as described in the previous item. Since f is assumed to be ϵ -far from any monotone k -monomial, the function g cannot be a monotone k -monomial. Since $\epsilon \leq 2^{-k+2}$ then $\epsilon/2^5 \leq 2^{-k-3}$ and therefore f is 2^{-k-3} -close to g . Hence, by Lemma 8, f will be rejected with probability at least $9/10$ in the third step of the algorithm (the closure-under-intersection test).

Summing up, we get that the probability that f is accepted by the algorithm is less than a $1/3$, as required. ■

4.1 Analysis of the Affinity Test

In this subsection we prove Lemma 7. To this end we define the function g as follows:

$$g(a) \stackrel{\text{def}}{=} 1 \quad \text{if} \quad \Pr_{x,y \in F_1}[f(a \oplus x \oplus y) = 1] \geq 1/2 \quad \text{and} \quad g(a) \stackrel{\text{def}}{=} 0 \quad \text{otherwise.} \quad (4)$$

We shall prove two lemmas from which Lemma 7 follows.

Lemma 9 *If the probability that the affinity test accepts f is greater than $1/10$ then $\text{dist}(f, g) \leq \epsilon/2^5$.*

Lemma 10 *If the probability that the affinity test accepts f is greater than $1/10$, then for every $a, b, c \in G_1$ we have $(a \oplus b \oplus c) \in G_1$.*

Proof of Lemma 9: By definition,

$$\text{dist}(f, g) = 2^{-n} \cdot (|F_1 \setminus G_1| + |G_1 \setminus F_1|).$$

We observe that for any particular $a \in F_1 \setminus G_1$

$$\Pr_{x,y \in F_1}[f(a \oplus x \oplus y) \neq f(a)] = \Pr_{x,y \in F_1}[f(a \oplus x \oplus y) = 0] > 1/2$$

where the equality follows from $a \in F_1$, and the inequality from $a \notin G_1$. Similarly, for any particular $a \in G_1 \setminus F_1$,

$$\Pr_{x,y \in F_1}[f(a \oplus x \oplus y) \neq f(a)] = \Pr_{x,y \in F_1}[f(a \oplus x \oplus y) = 1] \geq 1/2.$$

Assume, contrary to the claim, that $\text{dist}(f, g) > \epsilon/2^5$. Then with probability at least $1 - (1 - \epsilon/2^5)^{2^5/(\epsilon \cdot \delta)} > 1 - e^{-1/\delta}$, one of the a_i 's selected in Step 2b of Algorithm 2 is in the symmetric

difference $(F_1 \setminus G_1) \cup (G_1 \setminus F_1)$. For that a_i , with probability at least $1 - (1 - 1/2)^{4/\delta} > 1 - e^{-2/\delta}$ Algorithm 2 selects, in Step 2c, a pair (x_j, y_j) such that $f(a_i) \neq f(a_i \oplus x_j \oplus y_j)$. Hence, the probability that f is rejected in this case is greater than $(1 - e^{-1/\delta})(1 - e^{-2/\delta}) > 9/10$, for $\delta \leq 1/5$. But this contradicts the premise of the lemma by which f is accepted with probability greater than $1/10$. ■

In order to prove Lemma 10 we introduce some notation and prove a few claims.

For any $a \in \{0, 1\}^n$ let

$$WP(a) \stackrel{\text{def}}{=} \{(x, y) \in F_1 \times F_1 : f(a \oplus x \oplus y) \neq f(a)\} \quad (5)$$

be the set of *witness pairs* (x, y) that together with a constitute evidence against the affinity of F_1 . Let

$$H = \{a \in \{0, 1\}^n : |WP(a)| > \delta \cdot |F_1|^2\} \quad (6)$$

be the set of *heavy points* $a \in \{0, 1\}^n$ for which there are relatively many pairs $(x, y) \in F_1 \times F_1$ that together with a constitute evidence against the affinity of F_1 .

The claim below follows from the definition of H and the test, similarly to Lemma 9, where we take into account that $\epsilon \leq 2^{-k+2}$ (see Observation 5 and the discussion following it).

Claim 11 *If the probability that the affinity test accepts f is greater than $1/10$, then $|H| \leq \delta \cdot 2^{n-k-1}$.*

In all that follows we assume that the affinity test accepts f with probability greater than $1/10$. Since we assume that $|F_1| \geq 2^{n-k-1}$, it directly follows from Claim 11 that $|H| \leq \delta \cdot |F_1|$.

By the definition of g we know that for every $a \in \{0, 1\}^n$, $\Pr_{x,y \in F_1}[g(a) = f(a \oplus x \oplus y)] \geq 1/2$. In the claim below we show that this agreement probability is actually higher.

Claim 12 *For every $a \in \{0, 1\}^n$, $\Pr_{x,y \in F_1}[g(a) = f(a \oplus x \oplus y)] \geq 1 - 4\delta$.*

Proof: We fix a and let $\gamma \stackrel{\text{def}}{=} \Pr_{x,y \in F_1}[g(a) = f(x \oplus y \oplus a)]$. By the definition of $g(\cdot)$, $\gamma \geq \frac{1}{2}$. We are interested in strengthening this bound. Consider the following equality and inequality for $\Pr_{x_1, x_2, y_1, y_2 \in F_1}[f(a \oplus x_1 \oplus y_1) = f(a \oplus x_2 \oplus y_2)]$:

$$\begin{aligned} & \Pr_{x_1, x_2, y_1, y_2 \in F_1}[f(a \oplus x_1 \oplus y_1) = f(a \oplus x_2 \oplus y_2)] \\ &= \Pr_{x_1, x_2, y_1, y_2 \in F_1}[(f(a \oplus x_1 \oplus y_1) = g(a)) \wedge (f(a \oplus x_2 \oplus y_2) = g(a))] \\ & \quad + \Pr_{x_1, x_2, y_1, y_2 \in F_1}[f(a \oplus x_1 \oplus y_1) \neq g(a) \wedge (f(a \oplus x_2 \oplus y_2) \neq g(a))] \\ &= \gamma^2 + (1 - \gamma)^2 \end{aligned} \quad (7)$$

and,

$$\begin{aligned} & \Pr_{x_1, x_2, y_1, y_2 \in F_1}[f(a \oplus x_1 \oplus y_1) = f(a \oplus x_2 \oplus y_2)] \\ & \geq \Pr_{x_1, x_2, y_1, y_2 \in F_1}[f(a \oplus x_1 \oplus y_1) = f(a \oplus x_1 \oplus x_2 \oplus y_1 \oplus y_2) \\ & \quad \wedge (f(a \oplus x_2 \oplus y_2) = f(a \oplus x_1 \oplus x_2 \oplus y_1 \oplus y_2))] \\ &= 1 - \Pr_{x_1, x_2, y_1, y_2 \in F_1}[(f(a \oplus x_1 \oplus y_1) \neq f(a \oplus x_1 \oplus x_2 \oplus y_1 \oplus y_2)) \\ & \quad \vee (f(a \oplus x_2 \oplus y_2) \neq f(a \oplus x_1 \oplus x_2 \oplus y_1 \oplus y_2))] \\ & \geq 1 - 2 \cdot \Pr_{x_1, x_2, y_1, y_2 \in F_1}[f(a \oplus x_1 \oplus y_1) \neq f(a \oplus x_1 \oplus x_2 \oplus y_1 \oplus y_2)]. \end{aligned} \quad (8)$$

Subclaim 12.1 $\Pr_{x_1, x_2, y_1, y_2 \in F_1} [(f(a \oplus x_1 \oplus y_1)) \neq f(a \oplus x_1 \oplus x_2 \oplus y_1 \oplus y_2)] \leq 2\delta$.

Proof: We bound $\Pr_{x_1, x_2, y_1, y_2 \in F_1} [(f(a \oplus x_1 \oplus y_1)) \neq f(a \oplus x_1 \oplus x_2 \oplus y_1 \oplus y_2)]$ by the sum of two terms:

$$\begin{aligned} & \Pr_{x_1, x_2, y_1, y_2 \in F_1} [f(a \oplus x_1 \oplus y_1) \neq f((a \oplus x_1 \oplus y_1) \oplus x_2 \oplus y_2)] \\ & \leq \Pr_{x_1, y_1 \in F_1} [(a \oplus x_1 \oplus y_1) \in H] \cdot \max_{a' \in H} \{ \Pr_{x_2, y_2 \in F_1} [f(a') \neq f(a' \oplus x_2 \oplus y_2)] \} \\ & \quad + \Pr_{x_1, y_1 \in F_1} [(a \oplus x_1 \oplus y_1) \notin H] \cdot \max_{a' \notin H} \{ \Pr_{x_2, y_2 \in F_1} [f(a') \neq f(a' \oplus x_2 \oplus y_2)] \}. \end{aligned} \quad (9)$$

By the definition of H and since $\Pr_{x_1, y_1 \in F_1} [(a \oplus x_1 \oplus y_1) \notin H] \leq 1$, the second term in the above sum is bounded by δ . It remains to bound the first term by δ as well. We shall use the trivial bound of 1 for $\max_{a' \in H} \{ \Pr_{x_2, y_2 \in F_1} [f(a') \neq f(a' \oplus x_2 \oplus y_2)] \}$, and show that $\Pr_{x_1, y_1 \in F_1} [(a \oplus x_1 \oplus y_1) \in H] \leq \delta$.

For each choice of $x_1 \in F_1$, consider the set $B(a, x_1) \stackrel{\text{def}}{=} \{y_1 \in F_1 : a \oplus x_1 \oplus y_1 \in H\}$. Since for each pair $y_1, y'_1 \in F_1$ such that $y_1 \neq y'_1$ we have $a \oplus x_1 \oplus y_1 \neq a \oplus x_1 \oplus y'_1$, then there is a one-to-one mapping from $B(a, x_1)$ into H . Therefore, $|B(a, x_1)| \leq |H|$ for every $x_1 \in F_1$. It follows that

$$\begin{aligned} \Pr_{x_1, y_1 \in F_1} [(a \oplus x_1 \oplus y_1) \in H] &= \frac{1}{|F_1|^2} \sum_{x_1 \in F_1} |B(a, x_1)| \\ &\leq \frac{|H|}{|F_1|} \leq \frac{\delta \cdot |F_1|}{|F_1|} = \delta \end{aligned} \quad (10)$$

where we have used Claim 11 in the last inequality. \square

Subclaim 12.2 *Let $\frac{1}{2} \leq \gamma \leq 1$ and suppose that $\gamma^2 + (1 - \gamma)^2 \geq 1 - \beta$, for some $0 < \beta < 1$. Then $\gamma \geq 1 - \beta$.*

Proof: If $\gamma^2 + (1 - \gamma)^2 \geq 1 - \beta$, then $2\gamma(1 - \gamma) \leq \beta$. Since $\gamma \geq 1/2$, this implies that $1 - \gamma \leq \beta$, or equivalently that $\gamma \geq 1 - \beta$. \square

By combining Equations (7) and (8) with Subclaim 12.1, we obtain that $\gamma^2 + (1 - \gamma)^2 \geq 1 - 4\delta$. Claim 12 follows by applying Subclaim 12.2. \blacksquare

Recall that in order to prove Lemma 10 we need to show that for every $a, b, c \in G_1$, $g(a \oplus b \oplus c) = 1$. That is, by the definition of g , we need to show that $\Pr_{x, y \in F_1} [f((a \oplus b \oplus c) \oplus x \oplus y) = 1] \geq 1/2$ for every $a, b, c \in G_1$. To this end we first prove the following related claim.

Claim 13 *For every $a, b, c \in G_1$*

$$\Pr_{x_1, y_1, x_2, y_2, x_3, y_3 \in F_1} [f((a \oplus x_1 \oplus y_1) \oplus (b \oplus x_2 \oplus y_2) \oplus (c \oplus x_3 \oplus y_3)) = 1] \geq 1 - 14\delta.$$

Proof: We first observe that by the definition of $WP(\cdot)$ (see Equation (5)) we have:

$$\begin{aligned} & \Pr_{x_1, y_1, x_2, y_2, x_3, y_3 \in F_1} [f((a \oplus x_1 \oplus y_1) \oplus (b \oplus x_2 \oplus y_2) \oplus (c \oplus x_3 \oplus y_3)) = 1] \\ & \geq \Pr_{x_1, y_1, x_2, y_2, x_3, y_3 \in F_1} [(a \oplus x_1 \oplus y_1) \in F_1 \setminus H \text{ and} \\ & \quad ((b \oplus x_2 \oplus y_2), (c \oplus x_3 \oplus y_3)) \in (F_1 \times F_1) \setminus WP(a \oplus x_1 \oplus y_1)] \\ & \geq \Pr_{x_1, y_1 \in F_1} [(a \oplus x_1 \oplus y_1) \in F_1 \setminus H] \\ & \quad \times \min_{a' \in F_1 \setminus H} \Pr_{x_2, y_2, x_3, y_3 \in F_1} [((b \oplus x_2 \oplus y_2), (c \oplus x_3 \oplus y_3)) \in (F_1 \times F_1) \setminus WP(a')]. \end{aligned} \quad (11)$$

By Claim 12 and since $a \in G_1$, we know that $\Pr_{x_1, y_1 \in F_1}[(a \oplus x_1 \oplus y_1) \in F_1] \geq 1 - 4\delta$. By Equation (10) we know that $\Pr_{x_1, y_1 \in F_1}[(a \oplus x_1 \oplus y_1) \in H] \leq \delta$. Hence,

$$\Pr_{x_1, y_1 \in F_1}[(a \oplus x_1 \oplus y_1) \in F_1 \setminus H] \geq 1 - 5\delta. \quad (12)$$

It remains to bound the second term in the product above in Equation (11). By Claim 12, and since $b, c \in G_1$, we know that $\Pr_{x_2, y_2 \in F_1}[(b \oplus x_2 \oplus y_2) \in F_1] \geq 1 - 4\delta$ and similarly that $\Pr_{x_3, y_3 \in F_1}[(c \oplus x_3 \oplus y_3) \in F_1] \geq 1 - 4\delta$. Hence

$$\begin{aligned} \Pr_{x_2, y_2, x_3, y_3 \in F_1}[(b \oplus x_2 \oplus y_2), (c \oplus x_3 \oplus y_3) \in (F_1 \times F_1)] &\geq (1 - 4\delta)^2 \\ &\geq 1 - 8\delta. \end{aligned} \quad (13)$$

Let us fix some $a' \in F_1 \setminus H$. By the definition of H we know that $|WP(a')| \leq \delta \cdot |F_1|^2$. For any fixed $x_2, x_3 \in F_1$ let

$$B(a', b, c, x_2, x_3) \stackrel{\text{def}}{=} \{(y_2, y_3) \in F_1 \times F_1 : ((b \oplus x_2 \oplus y_2), (c \oplus x_3 \oplus y_3)) \in WP(a')\}.$$

Observe that for every two different pairs $(y_2, y_3), (y'_2, y'_3) \in F_1 \times F_1$ (i.e., such that either $y_2 \neq y'_2$ or $y_3 \neq y'_3$), the pair $((b \oplus x_2 \oplus y_2), (c \oplus x_3 \oplus y_3))$ differs from $((b \oplus x_2 \oplus y'_2), (c \oplus x_3 \oplus y'_3))$. That is, there is a one-to-one mapping from $B(a', b, c, x_2, x_3)$ into $WP(a')$. It follows that for every $x_2, x_3 \in F_1$,

$$|B(a', b, c, x_2, x_3)| \leq |WP(a')| \leq \delta \cdot |F_1|^2$$

and so

$$\begin{aligned} \Pr_{x_2, y_2, x_3, y_3 \in F_1}[(b \oplus x_2 \oplus y_2), (c \oplus x_3 \oplus y_3) \in WP(a')] &= \frac{1}{|F_1|^4} \sum_{x_2, x_3 \in F_1} |B(a', b, c, x_2, x_3)| \\ &\leq \frac{\delta \cdot |F_1|^2}{|F_1|^2} = \delta. \end{aligned} \quad (14)$$

The claim follows by combining Equations (11)–(14). \blacksquare

We are now ready to prove Lemma 10.

Proof of Lemma 10: Consider any fixed choice of $a, b, c \in G_1$. If $a = b$ then $(a \oplus b \oplus c) = c$, so that $(a \oplus b \oplus c) \in G_1$ and the claim holds by definition. Similarly for the case $a = c$ or $b = c$. Hence we may assume from now on the all three points a, b, c are different.

Assume contrary to the claim that there exist $a, b, c \in G_1$ such that $g(a \oplus b \oplus c) \neq 1$. By the definition of g this means that $\Pr_{x, y \in F_1}[f((a \oplus b \oplus c) \oplus x \oplus y) = 1] < 1/2$. In other words, if we let

$$O(a, b, c) \stackrel{\text{def}}{=} \{(x, y) \in F_1 \times F_1 : f((a \oplus b \oplus c) \oplus x \oplus y) = 1\}$$

then $|O(a, b, c)| < |F_1|^2/2$. We shall show that this contradicts Claim 13.

For every fixed choice of $x_1, x_2, y_1, y_2 \in F_1$, let

$$O(x_1, x_2, y_1, y_2) \stackrel{\text{def}}{=} \{(x_3, y_3) \in F_1 \times F_1 : ((x_1 \oplus x_2 \oplus x_3), (y_1 \oplus y_2 \oplus y_3)) \in O(a, b, c)\}.$$

(In order to be consistent with previous notation, we should have let $O(x_1, x_2, y_1, y_2)$ be denoted by $O(a, b, c, x_1, x_2, y_1, y_2)$, but have chosen the above notation in consideration to the reader.) Then

similarly to what we have argued before for similar subsets, $|O(x_1, x_2, y_1, y_2)| \leq |O(a, b, c)|$. Hence,

$$\begin{aligned}
& \Pr_{x_1, x_2, x_3, y_1, y_2, y_3 \in F_1} [((x_1 \oplus x_2 \oplus x_3), (y_1 \oplus y_2 \oplus y_3)) \in O(a, b, c)] \\
&= \frac{1}{|F_1|^6} \sum_{x_1, x_2, y_1, y_2 \in F_1} |O(x_1, x_2, y_1, y_2)| \\
&\leq \frac{|O(a, b, c)|}{|F_1|^2} < 1/2.
\end{aligned} \tag{15}$$

Recalling that $O(a, b, c) \subset F_1 \times F_1$, this implies that

$$\begin{aligned}
& \Pr_{x_1, x_2, x_3, y_1, y_2, y_3 \in F_1} [f((a \oplus x_1 \oplus y_1) \oplus (b \oplus x_2 \oplus y_2) \oplus (c \oplus x_3 \oplus y_3)) = 1] \\
&\leq \Pr_{x_1, x_2, x_3, y_1, y_2, y_3 \in F_1} [((x_1 \oplus x_2 \oplus x_3), (y_1 \oplus y_2 \oplus y_3)) \in O(a, b, c)] \\
&\quad + \Pr_{x_1, x_2, x_3, y_1, y_2, y_3 \in F_1} [((x_1 \oplus x_2 \oplus x_3), (y_1 \oplus y_2 \oplus y_3)) \notin F_1 \times F_1].
\end{aligned} \tag{16}$$

By Equation (15), the first term in the sum above is less than 1/2. We turn to bound the second term.

$$\begin{aligned}
& \Pr_{x_1, x_2, x_3, y_1, y_2, y_3 \in F_1} [((x_1 \oplus x_2 \oplus x_3), (y_1 \oplus y_2 \oplus y_3)) \notin F_1 \times F_1] \\
&\leq 2 \cdot \Pr_{x_1, x_2, x_3 \in F_1} [(x_1 \oplus x_2 \oplus x_3) \notin F_1] \\
&\leq 2 \cdot \Pr_{x_1 \in F_1} [x_1 \in H] \\
&\quad + 2 \cdot \Pr_{x_1 \in F_1} [x_1 \notin H] \cdot \max_{x'_1 \notin H} \Pr_{x_2, x_3 \in F_1} [(x'_1 \oplus x_2 \oplus x_3) \notin F_1] \\
&\leq 2 \cdot (\delta + \delta)
\end{aligned} \tag{17}$$

where in the last inequality we have applied Claim 11 and the definition of H . Thus we get that

$$\Pr_{x_1, x_2, x_3, y_1, y_2, y_3 \in F_1} [f((a \oplus x_1 \oplus y_1) \oplus (b \oplus x_2 \oplus y_2) \oplus (c \oplus x_3 \oplus y_3)) = 1] < (1/2) + 4\delta.$$

But for $\delta \leq 1/36$ we get a contradiction to Claim 13. ■

Proof of Lemma 7: Suppose that the affinity test accepts with probability greater than 1/10, and let g be as defined in Equation (4). By Lemma 9 we have that $\text{dist}(f, g) \leq \epsilon/2^5$ as required. By applying Lemma 10 we get that G_1 is an affine subspace. Furthermore, its dimension must be $n - k$ given the premise of the lemma concerning the size of F_1 . The last item in the lemma follows from the definition of g : Given any input a , the procedure Self-Correct simply selects a sufficiently large (but constant) number of pairs $x, y \in F_1$, and returns the majority value obtained for $f(a \oplus x \oplus y)$. By Claim 12 it returns the correct value $g(a)$ with high probability.

If F_1 is an affine subspace then by Fact 6 the test always accepts f , and by definition of g in Equation (4), $g = f$. Finally, the procedure Self-Correct as defined above always returns $f(a)$ for every $a \in \{0, 1\}^n$. ■

4.2 Analysis of the Closure-Under-Intersection Test

In this Subsection we prove Lemma 8.

4.2.1 Properties of Affine Subspaces

We first recall several simple properties of affine subspaces.

Claim 14 *Let H be an affine subspace such that $H = V \oplus x$, where $x \in \{0, 1\}^n$ and $V \subseteq \{0, 1\}^n$ is a linear subspace. Then,*

1. $x \in H$.
2. For every $z \in H$ we have that $H = V \oplus z$. By the definition of the \oplus operator, we thus also have that $V = H \oplus z$, for every $z \in H$.
3. $|H| = |V| = 2^{\dim V}$.

Claim 15 *Let H, H' be two affine subspaces of $\{0, 1\}^n$, such that $H \not\subseteq H'$. Then:*

$$\frac{|H \cap H'|}{|H|} \leq \frac{1}{2}.$$

Proof: The claim follows from the corresponding property of linear subspaces, namely $V \not\subseteq V'$ implies that $|V \cap V'|/|V| \leq 1/2$. ■

The following corollary is immediate:

Corollary 16 *Let H, H' be two affine subspaces of $\{0, 1\}^n$ such that $H' \subseteq H$. Then either $H' = H$, or $|H'| \leq |H|/2$.*

Claim 17 *Let H, H' be two affine subspaces of $\{0, 1\}^n$ such that $H' \subseteq H$ and let $y \in H'$. Denote by V' the linear subspace such that $H' = V' \oplus y$, and by V the linear subspace such that $H = V \oplus y$. Then:*

1. $V' \subseteq V$.
2. For any $x \in V$ we have $(H' \oplus x) \subseteq H$, and for any $x \notin V$ we have $(H' \oplus x) \cap H = \emptyset$.

Proof: By definition, $V' = H' \oplus y \subseteq H \oplus y = V$. This proves the first part of the lemma.

Now let $x \in V$. Since V and V' are linear subspaces and $V' \subseteq V$, then $(V' \oplus x) \subseteq V$. Thus, $H' \oplus x = (V' \oplus y) \oplus x = (V' \oplus x) \oplus y \subseteq V \oplus y = H$. On the other hand, let $x \notin V$. Observe that $(V' \oplus x) \cap V = \emptyset$. Since $H' \oplus x = (V' \oplus y) \oplus x = (V' \oplus x) \oplus y$, we get that $(H' \oplus x) \cap H = (V' \oplus x) \oplus y \cap (V \oplus y) = \emptyset$. This concludes the proof of the claim. ■

4.2.2 Auxiliary Claims

In order to prove Lemma 8 we will need several auxiliary claims. The first claim relates affine spaces that correspond to k -monomials and monotonicity.

Claim 18 *Let H be an affine subspace of $\{0, 1\}^n$ of size 2^{n-k} . Assume also that H is monotone. Namely, if $x \in H$ and $y \succeq x$, then $y \in H$. Then $H = \{x : x_{i_1} = 1, \dots, x_{i_k} = 1\}$, for some subset i_1, \dots, i_k of coordinates.*

Proof: Let V be an $n - k$ dimensional linear subspace and let $y \in \{0, 1\}^n$ be such that $H = V \oplus y$. Let v_1, \dots, v_{n-k} be a basis of V . Consider an $(n - k) \times n$ matrix with rows v_1, \dots, v_{n-k} . Its rank is $n - k$, and therefore it has $n - k$ linearly independent columns. Without loss of generality, these are the first $n - k$ columns. Therefore the restriction of the rows to the first $n - k$ coordinates is a basis of $\{0, 1\}^{n-k}$, and thus it spans all the vectors in $\{0, 1\}^{n-k}$ and in particular the first $n - k$ coordinates of y . It follows that there is a vector $v \in V$, namely a linear combination of the rows, that is equal to y on the first $n - k$ coordinates. Therefore, $z = (v \oplus y) \in H$ is 0 on the first $n - k$ coordinates.

Since H is monotone, if $|z| < k$, or there exists a $z' \not\preceq z$ such that $z' \in H$, then $|H| > 2^{n-k}$, contradicting our assumption on H . Hence $H = \{x : x_{i_1} = 1, \dots, x_{i_k} = 1\}$ where i_1, \dots, i_k are the coordinates on which z is 1. ■

Recall that by the premise of Lemma 8, there exists a function g such that $\text{dist}(f, g) \leq 2^{-k-3}$, and $G_1 \stackrel{\text{def}}{=} \{x : g(x) = 1\}$ is an affine subspace of dimension $n - k$. Claim 18 implies that if g is not a k -monomial, then the affine subspace G_1 cannot be monotone. We shall use this, together with the fact that f and g are close, to prove that there are many pairs $x \in F_1$, $y \in \{0, 1\}^n$ such that $f(y) \neq f(x \wedge y)$. To this end we define the following subsets.

Definition 9 Let $x \in \{0, 1\}^n$ and $z \in 2^x$. Define $G(x, z) \stackrel{\text{def}}{=} \{y \mid x \wedge y = z\}$.

We shall show that for many pairs (x, z) , with $x \in G_1$ and $z \in 2^x$, the function g is far from constant on $G(x, z)$. Since the functions f and g are close to each other, this will imply the existence of many violating pairs, as desired. First, we prove some properties of the subsets $G(x, z)$.

Claim 19 For every $x \in \{0, 1\}^n$ and $z \in 2^x$, $G(x, z)$ is an affine subspace of $\{0, 1\}^n$ of size $2^{n-|x|}$. Furthermore, for every $x \in \{0, 1\}^n$, the affine subspaces $\{G(x, z)\}_{z \in 2^x}$ partition $\{0, 1\}^n$.

Proof: These facts about $G(x, z)$ follow easily from the following observation: for a fixed x , the map $m_x : y \rightarrow x \wedge y$ is a linear map from $\{0, 1\}^n$ to 2^x , and $G(x, z) = m_x^{-1}(z)$. ■

Claim 20 Let $x \in G_1$ be such that there exists $z \in 2^x$ for which $G(x, z) \subseteq G_1$. Then, $G(x, x) \subseteq G_1$.

Proof: We first show that $G(x, z) \oplus x \oplus z \subseteq G_1$. Since G_1 is an affine subspace, by Fact 6, is it enough to show that x and z lie in G_1 , and that $G(x, z)$ is a subset of G_1 . Taking into account the assumptions of the Claim, we only need to show that $z \in G_1$. Since $z \preceq x$, we have $z \wedge x = z$. Hence, $z \in G(x, z) \subseteq G_1$.

Next, we show that $G(x, x) \subseteq G(x, z) \oplus x \oplus z$. Take $y \in G(x, x)$. Now, define y' as follows. If $z_i = 1$, then $y'_i = 1$ (in this case always $x_i = 1$). If $z_i = 0$ and $x_i = 1$, then $y'_i = 0$, and if $z_i = 0$ and $x_i = 0$, then $y'_i = y_i$. Thus, $y' \wedge x = z$ and so $y' \in G(x, z)$. It is also easy to verify that $y' \oplus x \oplus z = y$. (Note that $y \succeq x$, and therefore $x_i = 1$ implies that $y_i = 1$). Hence, $y \in G(x, z) \oplus x \oplus z$. Since we have shown that $G(x, z) \oplus x \oplus z \subseteq G_1$, the claim follows. ■

We shall be interested in the following set:

$$\mathcal{X} \stackrel{\text{def}}{=} \{x \in G_1 : G(x, x) \subseteq G_1\}. \quad (18)$$

Thus \mathcal{X} consists of those $x \in G_1$ for which every $y \succeq x$ is in G_1 . Since, by Claim 18, the set G_1 is not monotone, then necessarily $\mathcal{X} \neq G_1$. As we show momentarily, \mathcal{X} is actually significantly smaller than G_1 , and we shall exploit this in our proof.

Claim 21 *The set \mathcal{X} is an affine subspace of G_1 . Furthermore, if g is not a k -monomial then $|\mathcal{X}| \leq \frac{1}{2}|G_1|$.*

Proof: By Fact 6, in order to prove the first part of the lemma it suffices to show that for every $x^1, x^2, x^3 \in \mathcal{X}$, we have $x^1 \oplus x^2 \oplus x^3 \in \mathcal{X}$. Let us fix $x^1, x^2, x^3 \in \mathcal{X}$, and let $x = x^1 \oplus x^2 \oplus x^3$. To show that $x \in \mathcal{X}$ we have to show that $G(x, x) \subseteq G_1$. Namely, that for every $y \succeq x$, we have $y \in G_1$. Let $y \succeq x$. Then there exist y^1, y^2, y^3 such that $y = y^1 \oplus y^2 \oplus y^3$, where $y^j \succeq x^j$ for $j = 1, \dots, 3$. (To verify this, choose a coordinate i : (1) If $y_i = x_i$: set $y_i^j = x_i^j$ for all j . (2) If $y_i = 1$ and $x_i = 0$: Set $y_i^j = 1$ for all j .) That is, $y^j \in G(x^j, x^j) \subseteq G_1$. Therefore $y^j \in G_1$ for all j , and so $y = y^1 \oplus y^2 \oplus y^3 \in G_1$.

By Corollary 16, since \mathcal{X} is an affine subspace of G_1 , either $\mathcal{X} = G_1$, or $|\mathcal{X}| \leq \frac{1}{2}|G_1|$. If $\mathcal{X} = G_1$, then for any $x \in G_1$ we have $G(x, x) = \{y : y \succeq x\} \subseteq G_1$, namely G_1 is monotone. By Claim 18, g is a k -monomial, which contradicts our assumptions. Therefore, $|\mathcal{X}| \leq \frac{1}{2}|G_1|$. ■

In the next claim we show that for every $x \in G_1 \setminus \mathcal{X}$, the function g is far from constant on $G(x, z)$, for many $z \in 2^x$. Observe that this is trivially true if g is a monotone monomial, since in this case the set $G_1 \setminus \mathcal{X}$ is empty.

Claim 22 *For every $x \in G_1 \setminus \mathcal{X}$, and for any fixed function $h : \{0, 1\}^n \rightarrow \{0, 1\}$,*

$$\frac{1}{2^{|x|}} \cdot \sum_{z \in 2^x} \Pr_{y \in G(x, z)}[g(y) \neq h(z)] \geq 2^{-k}.$$

Proof: Let us fix $x \in G_1 \setminus \mathcal{X}$ and a function h . For every $z \in 2^x$, if $h(z) = 0$ then

$$\Pr_{y \in G(x, z)}[g(y) \neq h(z)] = \frac{|G(x, z) \cap G_1|}{|G(x, z)|}$$

and if $h(z) = 1$ then

$$\Pr_{y \in G(x, z)}[g(y) \neq h(z)] = \frac{|G(x, z) \setminus G_1|}{|G(x, z)|} = 1 - \frac{|G(x, z) \cap G_1|}{|G(x, z)|}.$$

Hence,

$$\Pr_{y \in G(x, z)}[g(y) \neq h(z)] \geq \min \left\{ \frac{|G(x, z) \cap G_1|}{|G(x, z)|}, 1 - \frac{|G(x, z) \cap G_1|}{|G(x, z)|} \right\}. \quad (19)$$

But, for all $z \in 2^x$, $G(x, z) \not\subseteq G_1$ (otherwise, by Claim 20, we would have $G(x, x) \subseteq G_1$, and so $x \in \mathcal{X}$). Thus, by Claim 15, $\frac{|G(x, z) \cap G_1|}{|G(x, z)|} \leq \frac{1}{2}$. Combining this with Equation (19),

$$\begin{aligned} \frac{1}{2^{|x|}} \cdot \sum_{z \in 2^x} \Pr_{y \in G(x, z)}[g(y) \neq h(z)] &\geq \frac{1}{2^{|x|}} \cdot \sum_{z \in 2^x} \frac{|G(x, z) \cap G_1|}{|G(x, z)|} \\ &= 2^{-n} \cdot \sum_{z \in 2^x} |G(x, z) \cap G_1| = 2^{-n} \cdot |G_1| = 2^{-k} \end{aligned} \quad (20)$$

In the last sequence of steps we have used the following: (1) $|G(x, z)| = 2^{n-|x|}$ for every z (Claim 19); (2) For every x , the subsets $G(x, z)$ form a partition of $\{0, 1\}^n$ (Claim 19); (3) G_1 is of size 2^{n-k} . ■

4.2.3 Proof of Lemma 8

Let g be the function ensured by the premise of Lemma 8. If g is not a k -monomial then by Claim 21 we know that $|\mathcal{X}| \leq |G_1|/2$. In other words,

$$|G_1 \setminus \mathcal{X}| \geq \frac{1}{2}|G_1| \geq 2^{n-k-1}$$

where the second inequality follows from the fact that G_1 is an affine subspace of dimension $n - k$. Since $\text{dist}(f, g) \leq 2^{-k-3}$ where $||F_1| - 2^{n-k}| \leq 2^{n-k-3}$, we know that

$$|(F_1 \cap G_1) \setminus \mathcal{X}| \geq 2^{n-k-2}.$$

Hence, with probability of at least $1 - (1 - 1/4)^{32} > 1 - e^{-8}$, Algorithm 2 obtains in Step 3a, a point $x \in (F_1 \cap G_1) \setminus \mathcal{X}$. Let us denote this point by x^1 .

Since $x^1 \in G_1 \setminus \mathcal{X}$ we know by Claim 22 that

$$\frac{1}{2^{|x^1|}} \cdot \sum_{z \in 2^{x^1}} \Pr_{y \in G(x^1, z)}[g(y) \neq g(z)] \geq 2^{-k}. \quad (21)$$

By Claim 19, $|G(x^1, z)| = 2^{n-|x^1|}$ and so we have that

$$\begin{aligned} \frac{1}{2^{|x^1|}} \cdot \sum_{z \in 2^{x^1}} \Pr_{y \in G(x^1, z)}[g(y) \neq g(z)] &= \frac{1}{2^{|x^1|}} \sum_{z \in 2^{x^1}} \frac{1}{2^{n-|x^1|}} |\{y \in G(x^1, z) : g(y) \neq g(z)\}| \\ &= \frac{1}{2^n} |\{y \in \{0, 1\}^n : g(y) \neq g(x^1 \wedge y)\}| \\ &= \Pr_{y \in \{0, 1\}^n} [g(y) \neq g(x^1 \wedge y)] \end{aligned} \quad (22)$$

By combining Equations (21) and (22) we have that

$$\Pr_{y \in \{0, 1\}^n} [g(y) \neq g(x^1 \wedge y)] \geq 2^{-k}. \quad (23)$$

Hence, with probability of at least $1 - (1 - 2^{-k})^{2^{k+3}} > 1 - e^{-8}$, Algorithm 2 will select, in Step 3b, such a point $y^1 \in \{0, 1\}^n$ for which $g(y^1) \neq g(x^1 \wedge y^1)$. Finally, if both calls to Self-Correct, in Step 3c, return correct values, which occurs with probability of at least $1 - 1/20$, then the algorithm will reject as desired. The lemma follows by combining all error probabilities. ■

4.3 Testing Monomials when k is Unspecified

Suppose that we want to test whether a function f is a monomial without the size of the monomial, k , being specified. In this case we start by finding k . We obtain an estimate α to $\Pr[f = 1]$, by taking a sample of size $\Theta(1/\epsilon)$. By a multiplicative Chernoff bound, such a sample ensures that, with high probability, if $\Pr[f = 1] \geq \epsilon/2$ then $\alpha \geq \epsilon/4$, while if $\Pr[f = 1] < \epsilon/8$, then $\alpha < \epsilon/4$. Hence, if $\alpha < \epsilon/4$ then we can immediately accept. This is true, since we may assume that $\Pr[f = 1] < \epsilon/2$, and so f is close to every monomial that contains at least $\log(2/\epsilon)$ literals.

Otherwise, we may assume that $\Pr[f = 1] \geq \epsilon/8$, and a multiplicative Chernoff bound implies that, with high probability, $(1 - 1/4) \cdot \Pr[f = 1] < \alpha < (1 + 1/4) \cdot \Pr[f = 1]$. Now, we look for an integer k for which $4/5\alpha \leq 2^{-k} \leq 4/3\alpha$. If there is no such integer, we reject. If there is, there is at most one, and we choose it as our estimate for k . If f is in fact a monomial, then this estimate of k is correct with high probability. Given this k , we proceed as before.

5 Testing Monotone DNF Formulae

In this section we describe an algorithm for testing whether a function f is a monotone DNF formula with at most ℓ terms, for a given integer ℓ .

In other words, we test whether $f = T_1 \vee T_2 \vee \dots \vee T_{\ell'}$, where $\ell' \leq \ell$, and each term T_i is a monotone monomial. Note that we allow the size of the terms to vary. We assume, without loss of generality, that no term contains the set of variables of any other term (or else we can ignore the more specific term), though the same variable can of course appear in several terms.

The basic idea underlying the algorithm is to test whether the set $F_1 \stackrel{\text{def}}{=} \{x : f(x) = 1\}$ can be “approximately covered” by at most ℓ terms (monomials). To this end, the algorithm finds strings $x_i \in \{0, 1\}^n$ and uses them to define functions f_i that are tested for being monomials. If the original function f is in fact an ℓ -term DNF, then, with high probability, each such function f_i corresponds to one of the terms of f .

The following notation will be useful. Let f be a monotone ℓ -term DNF, and let its terms be T_1, \dots, T_ℓ . Then, for any $x \in \{0, 1\}^n$, we let $S(x) \subseteq \{1, \dots, \ell\}$ denote the subset of indices of the terms satisfied by x . That is:

$$S(x) \stackrel{\text{def}}{=} \{i : T_i(x) = 1\} .$$

In particular, if $f(x) = 0$ then $S(x) = \emptyset$. This notion extends to a set $R \subseteq F_1$, where $S(R) \stackrel{\text{def}}{=} \bigcup_{x \in R} S(x)$. We observe that if f is a monotone ℓ -term DNF, then for every $x, y \in \{0, 1\}^n$

$$S(x \wedge y) = S(x) \cap S(y) .$$

We shall also need the following definitions.

Definition 10 (Single-Term Representatives) *Let f be a monotone ℓ -term DNF. We say that $x \in F_1$ is a single-term representative for f if $|S(x)| = 1$. That is, x satisfies only a single term in f .*

Definition 11 (Neighbors) *Let $x \in F_1$. The set of neighbors of x , denoted by $N(x)$, is defined as follows:*

$$N(x) \stackrel{\text{def}}{=} \{y \mid f(y) = 1 \text{ and } f(x \wedge y) = 1\} .$$

The notion of neighbors extends to a set $R \subseteq F_1$, where $N(R) \stackrel{\text{def}}{=} \bigcup_{x \in R} N(x)$.

Note that the above definition of neighbors is very different from the standard notion (that is, strings at Hamming distance 1), and in particular depends on the function f .

Consider the case in which x is a single-term representative of f , and $S(x) = \{i\}$. Then, for every neighbor $y \in N(x)$, we must have $i \in S(y)$ (or else $S(x \wedge y)$ would be empty, implying that $f(x \wedge y) = 0$). Notice that the converse statement holds as well, that is, $i \in S(y)$ implies that x and y are neighbors. Therefore, the set of neighbors of x is exactly the set of all strings satisfying the term T_i . The goal of the algorithm will be to find at most ℓ such single-term representatives $x \in \{0, 1\}^n$, and for each such x to test that its set of neighbors $N(x)$ satisfies some common term. We shall show that if f is in fact a monotone ℓ -term DNF, then all these tests pass with high probability. On the other hand, if all the tests pass with high probability, then f is close to some monotone ℓ -term DNF.

We start with a high-level description of the algorithm, and then show how to implement its main step of finding single-term representatives.

Algorithm 3 Test for Monotone ℓ -term DNF

1. $R \leftarrow \emptyset$. R is designated to be a set of single-term representatives for f .
2. For $i = 1$ to $\ell + 1$ (Try to add ℓ single-term representatives to R):
 - (a) Take a uniform sample U_i of size $m_1 = \Theta(\ell \log \ell / \epsilon)$ strings. Let $W_i = (U_i \cap F_1) \setminus N(R)$. That is, W_i consists of strings x in the sample such that $f(x) = 1$, and x is not a neighbor of any string already in R .
Observe that if the strings in R are in fact single-term representatives, then every $x \in W_i$ satisfies only terms not satisfied by the representatives in R .
 - (b) If $i = \ell + 1$ and $W_i \neq \emptyset$, then reject.
If there are more than ℓ single term representatives for f then necessarily f is not an ℓ -term DNF.
 - (c) Else, if $\frac{|W_i|}{m_1} < \frac{\epsilon}{4}$ then go to Step 3.
The current set of representatives already “covers” almost all of F_1 .
 - (d) Else $\left(\frac{|W_i|}{m_1} \geq \frac{\epsilon}{4} \text{ and } i \leq \ell\right)$, use W_i in order to find a string x_i that is designated to be a single-term representative of a term not yet represented in R . This step will be described subsequently.
3. For each string $x_i \in R$, let the function $f_i : \{0, 1\}^n \mapsto \{0, 1\}$ be defined as follows: $f_i(y) = 1$ if and only if $y \in N(x_i)$.
As observed previously, if x_i is in fact a single-term representative, then f_i is a monomial.
4. For each f_i , test that it is monomial, using distance parameter $\epsilon' = \frac{\epsilon}{2\ell}$ and confidence $1 - \frac{1}{6\ell}$ (instead of $\frac{2}{3}$ — this can simply be done by $O(\log \ell)$ repeated applications of each test).
Note that we do not specify the size of the monomial, and so we need to apply the appropriate variant of our test, as described in Subsection 4.3.
5. If any of the tests fail then reject, otherwise accept.

The heart of the algorithm lies in finding a new representative in each iteration of Step 2. This procedure will be described and analyzed shortly. In particular, we shall prove the following lemma.

Lemma 23 Suppose that f is a monotone ℓ -term DNF, and let $R \subset \{0, 1\}^n$ be a subset of single-term representatives for f such that $\Pr[x \in (F_1 \setminus N(R))] \geq \epsilon/8$. Let U_i be a uniformly selected sample of $m_1 = \Theta(\ell \log \ell / \epsilon)$ strings, and let $W_i = (U_i \cap F_1) \setminus N(R)$. Then there exists a procedure that receives W_i as input, for which the following holds:

1. With probability at least $1 - \frac{1}{6\ell}$, taken over the choice of U_i and the internal coin flips of the procedure, the procedure returns a string x_i that is a single term representative for f of a term not yet represented in R . That is, $|S(x_i)| = 1$ and $S(x_i) \cap S(R) = \emptyset$.
2. The query complexity of the procedure is $O(\ell \log^2 \ell / \epsilon)$.

Conditioned on the above lemma we can prove the following theorem.

Theorem 3 Algorithm 3 is a testing algorithm for monotone ℓ -term DNF. The query complexity of the algorithm is $\tilde{O}(\ell^2 / \epsilon)$.

Proof: We shall use the following notation: for any set $R \subset \{0, 1\}^n$, let

$$\bar{p}(R) \stackrel{\text{def}}{=} \Pr[x \in (F_1 \setminus N(R))].$$

Suppose that f is a monotone ℓ -term DNF, and consider each iteration of Step 2. By Lemma 23, if all strings in R are single-term representatives for f and $\bar{p}(R) \geq \epsilon/8$, then with probability of at least $1 - \frac{1}{6\ell}$, the procedure for finding a single term representative in fact returns a new representative (of a term not yet represented in R). Hence, the probability that, for some iteration i , the string x_i returned by the procedure is not a single-term representative, is at most $1/6$. Conditioned on such an event not occurring, Algorithm 3 completes Step 2 with a set R that contains at most ℓ single-term representatives for f .

In such a case, by the definition of single-term representatives, each f_i defined in Step 3 of Algorithm 3 is a monotone monomial. For each fixed f_i , the probability that it fails the monomial test is at most $\frac{1}{6\ell}$. By applying a union bound, the probability that any one of the f_i 's fail, is at most $\frac{1}{6}$. Adding up the error probabilities, we obtain that f is accepted with probability of at least $2/3$.

We now turn to the case in which f is ϵ -far from being a monotone ℓ -term DNF. Consider the value of $\bar{p}(R)$ at the start of each iteration i of Step 2. Observe that $\bar{p}(R)$ does not increase with i . If $\bar{p}(R) > \epsilon/2$, then, by a multiplicative Chernoff bound, the probability that $\frac{|W_i|}{m_1} \leq \epsilon/4$ (causing Algorithm 3 to exit Step 2) is smaller than $\frac{1}{6\ell}$. Hence, the probability that Algorithm 3 completes Step 2 without rejecting and with a set R for which $\bar{p}(R) > \epsilon/2$, is at most $1/6$.

Conditioned on such an event not occurring, consider the functions f_i defined in Step 3 of Algorithm 3. We claim that at least one of these functions is $\frac{\epsilon}{2\ell}$ -far from being a monomial. To verify this, assume in contradiction that all these $|R| \leq \ell$ functions are $\frac{\epsilon}{2\ell}$ -close to being monomials. For each such function f_i , let g_i be a closest monomial to f_i , and let $g = g_1 \vee g_2 \vee \dots \vee g_{|R|}$. Then $\text{dist}(f, g) \leq |R| \cdot \frac{\epsilon}{2\ell} + \bar{p}(R) \leq \epsilon$, contradicting the fact that f is ϵ -far from any ℓ -term DNF. Thus, let f_t be one of the f_i 's that is $\frac{\epsilon}{2\ell}$ -far from being a monomial. The probability that the monomial test does not reject f_t is at most $\frac{1}{6\ell}$. Adding up the error probabilities, f is rejected with probability of at least $2/3$.

Finally, we bound the query complexity of Algorithm 3. There are at most $\ell + 1$ iterations in Step 2 of the algorithm. In each iteration, $m_1 = O(\ell \log \ell / \epsilon)$ strings are queried in Step 2a. By Lemma 23, $O(\ell \log^2 \ell / \epsilon)$ strings are queried by the procedure for finding a new representative that is called in Step 2d. By Theorem 2, testing that each of the at most ℓ functions f_i is a monomial, requires a total of $\ell \cdot O(1/\epsilon') \cdot O(\log \ell) = \tilde{O}(\ell^2/\epsilon)$ queries. Therefore, the total number of queries performed by Algorithm 3 is $\tilde{O}(\ell^2/\epsilon)$. ■

5.1 Finding New Representatives

Suppose that f is a monotone ℓ -term DNF with terms T_1, \dots, T_ℓ , and consider an arbitrary iteration i in Step 2 of Algorithm 3. Assume that $R \subset \{0, 1\}^n$ is a subset of single-term representatives for f , such that $\Pr[x \in (F_1 \setminus N(R))] \geq \epsilon/8$. Let $\bar{N}(R) \stackrel{\text{def}}{=} F_1 \setminus N(R)$ be the set of all the strings that are not neighbors of any string in R , and let $\bar{S}(R) \stackrel{\text{def}}{=} \{1, \dots, \ell\} \setminus S(R)$ be the set of indices of terms not yet represented in R . By definition, $W_i \subseteq \bar{N}(R)$, and for every $x \in W_i$ we have $S(x) \subseteq \bar{S}(R)$.

Given a string $x_0 \in W_i$, we shall try to “remove” terms from $S(x_0)$, until we are left with a single term. More precisely, we produce a sequence of strings x_0, \dots, x_r , where $x_0 \in W_i$, such that $\emptyset \neq S(x_{j+1}) \subseteq S(x_j)$, and in particular $|S(x_r)| = 1$. The aim is to decrease the size of $S(x_j)$

by a constant factor for most j 's. This will ensure that for $r = \Theta(\log \ell)$, the final string x_r is a single-term representative as desired.

How is such a sequence obtained? Given a string $y_j \in N(x_j)$, define $x_{j+1} = x_j \wedge y_j$. Then $f(x_{j+1}) = 1$ (i.e., $S(x_{j+1}) \neq \emptyset$), and $S(x_{j+1}) = S(x_j) \cap S(y_j) \subseteq S(x_j)$. The string y_j is acquired by uniformly selecting a sufficiently large sample from $\{0, 1\}^n$, and picking the first string in the sample that belongs to $N(x_j)$, if such exists. The exact procedure follows.

Procedure for finding a new representative, given $W_i \subseteq \overline{N}(R)$:

1. Let the strings in W_i be denoted $w_1, \dots, w_{|W_i|}$.
2. Uniformly and independently select $r = \Theta(\log \ell)$ samples, Y_0, \dots, Y_{r-1} , each consisting of $m_2 = O(\ell \log \ell / \epsilon)$ strings from $\{0, 1\}^n$.
3. $found \leftarrow FALSE$; $t \leftarrow 0$;
4. While $found \neq TRUE$ and $t < |W_i|$ do:
 - (a) $t \leftarrow t + 1$; $x_0 \leftarrow w_t$.
 - (b) For $j = 1$ to r
 - i. If $Y_{j-1} \cap N(x_{j-1}) = \emptyset$ then exit the "for" loop and go to (4a).
 - ii. Otherwise, pick the first string $y_{j-1} \in Y_{j-1} \cap N(x_{j-1})$, and let $x_j = x_{j-1} \wedge y_{j-1}$.
 - (c) If $j = r$ then $found \leftarrow TRUE$.
5. If $found = TRUE$ then return x_r . Otherwise, return an arbitrary string.

We first prove that if Y_j intersects $N(x_j)$, then the probability that the size of $S(x_{j+1})$ is significantly smaller than that of $S(x_j)$ is at least a $1/3$. Observe that since the sample Y_j is uniformly distributed in $\{0, 1\}^n$, then $Y_j \cap N(x_j)$ is uniformly distributed in $N(x_j)$.

Claim 24 *Let x_j be a fixed string. With probability of at least a $1/3$ over the uniform choice of a string $y_j \in N(x_j)$, $|S(x_j \wedge y_j)| \leq 1 + \frac{3}{4} \cdot (|S(x_j)| - 1)$.*

Proof: Without loss of generality, let $S(x_j) = \{1, \dots, t\}$. We partition the set of neighbors $N(x_j)$ into disjoint subsets $N_i(x_j)$, for $1 \leq i \leq t$, where

$$N_i(x_j) = \{y \in N(x_j) : i \in S(y) \text{ and for every } i' < i, i' \notin S(y)\}.$$

Since y_j is uniformly distributed in $N(x_j)$, we can view it as being selected by first choosing i with probability $\frac{|N_i(x_j)|}{|N(x_j)|}$, and then selecting y uniformly in $N_i(x_j)$.

Consider the case $y_j \in N_1(x_j)$. In order to select a string uniformly in $N_1(x_j)$, we first set to 1 all the bits corresponding to the variables in T_1 , and then set the remaining bits to 0 or 1 with equal probability. Since for every $i \neq 1$ there is at least one variable that appears in T_i and not in T_1 , we have that

$$\Pr[T_i(y_j) = 0 \mid y_j \in N_1(x_j)] \geq \frac{1}{2}.$$

It follows that the expected number of indices $i \in S(x_j)$, $i \neq 1$, for which $T_i(y_j) = 1$ is at most $(t-1)/2$. By Markov's inequality, the probability that there are more than $(1-\alpha)(t-1)$ terms

T_i , $i \neq 1$, satisfied by a uniformly selected $y_j \in N_1(x_j)$, is at most $\frac{1}{2(1-\alpha)}$. Setting $\alpha = 1/4$, we get that, with probability of at least a $\frac{1}{3}$, over the choice of a uniformly selected $y_j \in N_1(x_j)$, we have that $|S(x_{j+1})| \leq 1 + \frac{3}{4} \cdot (|S(x_j)| - 1)$. It is easy to see that for any $N_i(x_j)$, $i > 1$, this probability is at least as large. In particular, note that for $i = t$, for any $y_j \in N_t(x_j)$, $|S(x_{j+1})| = 1$. ■

The next corollary follows directly from Claim 24 and the fact that $|S(x_0)| \leq \ell$.

Corollary 25 *Let $r = c \cdot \log \ell$, where c is a sufficiently large constant, and let x_0 be a fixed string in W_i . Consider the following process, consisting of r steps, where in the j 's step we uniformly and independently select a string $y_{j-1} \in N(x_{j-1})$ and set $x_j = x_{j-1} \wedge y_{j-1}$. Then, with probability of at least $1 - \frac{1}{18\ell}$ over the choice of y_0, \dots, y_{r-1} , we obtain $|S(x_r)| = 1$.*

Finally, we bound the size of a sample Y_j sufficient for acquiring a string $y_j \in N(x_j)$ with high probability. We first define a “good initial string” x_0 . This is a string that satisfies only relatively “large” terms.

Definition 12 *A string x_0 will be called a good initial string if for every $i \in S(x_0)$, $\Pr[T_i = 1] \geq \frac{\epsilon}{16\ell}$.*

Recall that $\overline{N}(R) = F_1 \setminus N(R)$ and define the set:

$$\text{Good} \stackrel{\text{def}}{=} \left\{ x \in \overline{N}(R) \mid x \text{ is a good initial string} \right\}.$$

Claim 26 *Suppose that $\Pr[x \in \overline{N}(R)] \geq \frac{\epsilon}{8}$. Then the probability, taken over the choices of U_i , that W_i does not contain any good initial strings, is at most $\frac{1}{18\ell}$.*

Proof: Recall that $\bar{p}(R) \stackrel{\text{def}}{=} \Pr[x \in \overline{N}(R)]$, and that $\overline{S}(R) \stackrel{\text{def}}{=} \{1, \dots, \ell\} \setminus S(R)$. For any $i \in \overline{S}(R)$, consider the event

$$E_i \stackrel{\text{def}}{=} \left\{ x \in \overline{N}(R) \text{ and } T_i(x) = 1 \right\}.$$

By definition, $\bar{p}(R) = \Pr \left[\bigcup_{i \in \overline{S}(R)} E_i \right]$. Let

$$\overline{S}_{\text{small}}(R) = \left\{ i \in \overline{S}(R) \text{ and } \Pr[E_i] \leq \frac{\bar{p}(R)}{2\ell} \right\}.$$

Clearly, for any term i , $\Pr[T_i = 1] \geq \Pr[E_i]$. Therefore, if $x \in \left(\bigcup_{i \in \overline{S}(R)} E_i \right) \setminus \left(\bigcup_{i \in \overline{S}_{\text{small}}(R)} E_i \right)$ then $S(x) \subseteq \overline{S}(R) \setminus \overline{S}_{\text{small}}(R)$, and therefore for all $i \in S(x)$ we have that $\Pr[T_i(x) = 1] \geq \Pr[E_i] \geq \frac{\bar{p}(R)}{2\ell} \geq \frac{\epsilon}{16\ell}$. Thus, $x \in \text{Good}$. Therefore,

$$\begin{aligned} \Pr[\text{Good}] &\geq \Pr \left[\left(\bigcup_{i \in \overline{S}(R)} E_i \right) \setminus \left(\bigcup_{i \in \overline{S}_{\text{small}}(R)} E_i \right) \right] \\ &\geq \Pr \left[\bigcup_{i \in \overline{S}(R)} E_i \right] - \Pr \left[\bigcup_{i \in \overline{S}_{\text{small}}(R)} E_i \right] \\ &\geq \bar{p}(R) - \ell \cdot \frac{\bar{p}(R)}{2\ell} = \frac{\bar{p}(R)}{2}. \end{aligned}$$

Since $\bar{p}(R) \geq \frac{\epsilon}{8}$ and the size of the sample U_i is $\Theta(\ell \log \ell / \epsilon)$, then the probability that W_i does not contain *any* good initial string is, for a sufficiently large constant in the $\Theta(\cdot)$ notation, smaller than $\frac{1}{18\ell}$. ■

The next claim follows from the definition of a good initial string.

Claim 27 *Let $m_2 = c_1 \cdot \ell \log \ell / \epsilon$ and $r = c_2 \cdot \log \ell$, where c_1, c_2 are sufficiently large constants. Let Y_1, \dots, Y_r be samples of m_2 strings each, and suppose that x_0 is a good initial string. Then, for each $1 \leq j \leq r$, the probability that $Y_j \cap N(x_j) \neq \emptyset$ is at least $1 - \frac{1}{18\ell^2}$.*

We can now complete the proof of Lemma 23, and hence the correctness of Theorem 3.

Proof of Lemma 23: By the premise of the lemma, $\Pr[x \in \overline{N}(R)] \geq \epsilon/8$. By Claim 26, the set W_i contains a good initial string with probability at least $1 - \frac{1}{18\ell}$. Conditioned on this event, let us fix such a string x_0 , and consider the execution of Step 4b in the procedure. By Claim 27, the probability that there exists a $j \leq r$ for which the sample Y_j does not contain a string in $N(x_j)$ is at most $\frac{1}{18\ell}$. Since the strings in Y_j are uniformly selected from $\{0, 1\}^n$, the strings in $Y_j \cap N(x_j)$ are uniformly distributed in $N(x_j)$. Hence, conditioned on each Y_j containing a string from $N(x_j)$, we can apply Corollary 25 and get that with probability of at least $1 - \frac{1}{18\ell}$, $|S(x_r)| = 1$. Since $x_0 \in \overline{N}(R)$, necessarily $x_r \in \overline{N}(R)$. Therefore, with probability of at least $1 - 3 \cdot \frac{1}{18\ell} = 1 - \frac{1}{6\ell}$, taken over the choices of U_i and the samples Y_j , the procedure returns a string x_r that is a single-term representative for f of a term not yet represented in f .

The number of queries performed by the procedure is $r \cdot m_2 = O(\ell \log^2 \ell / \epsilon)$ as promised. ■

6 Testing Singletons Without Testing Linearity

Recall that by Claim 1 an alternative characterization of singletons is that $\Pr[f = 1] = 1/2$, and furthermore that there are no violating pairs $x, y \in \{0, 1\}^n$. That is, there are no x, y such that $f(x \wedge y) \neq f(x) \wedge f(y)$. We show that the following simple algorithm that checks these properties, is a testing algorithm for singletons if f is not too far from a singleton function. Let $\mathcal{F}_{\text{SING}}$ denote the class of singletons. The algorithm will receive a value γ_0 such that $\min_{g \in \mathcal{F}_{\text{SING}}} \text{dist}(f, g) \leq \frac{1}{2} - \gamma_0$. That is, γ_0 is a lower bound on the difference between $1/2$ and the distance of f to the closest singleton. We shall think of γ_0 as a constant.

Algorithm 4 Test for Singletons with lower bound γ_0

1. **Size Test:** *Uniformly select a sample of $m = \Theta(1/\epsilon^2)$ strings in $\{0, 1\}^n$. For each x in the sample, obtain $f(x)$. Let α be the fraction of sample strings x such that $f(x) = 1$. If $|\alpha - 1/2| > \epsilon/4$ then reject, otherwise continue.*
2. **Closure-Under-Intersection Test:** *Repeat the following $\Theta(\epsilon^{-1}\gamma_0^{-1})$ times: Uniformly select $x, y \in \{0, 1\}^n$. If x and y are a violating pair, then reject.*
3. *If no step caused rejection, then accept.*

Theorem 4 *If f is a singleton, then Algorithm 4 accepts with probability of at least $2/3$. If f is ϵ -far from any singleton where ϵ is bounded away from $1/2$, then the algorithm rejects with probability of at least $2/3$. The query complexity of the algorithm is $O(1/\epsilon^2)$.*

Proof: If f is a singleton then $\Pr[f = 1] = 1/2$. By an additive Chernoff bound, and for the appropriate constant in the $\Theta(\cdot)$ notation, the probability that it is rejected in the first step of Algorithm 4 is less than a $1/3$. By the definition of singletons, f always passes the closure-under-intersection test.

Suppose that f is ϵ -far from any singleton and let δ be its distance to the closest singleton. Thus $\epsilon < \delta \leq 1/2 - \gamma_0$. We show that f is rejected with probability greater than $2/3$.

1. If $|\Pr[f = 1] - 1/2| > \frac{\epsilon}{2}$, then f is rejected in the first step of Algorithm 4 with probability of at least $5/6$.
2. Otherwise, $|\Pr[f = 1] - 1/2| \leq \frac{\epsilon}{2} < \frac{\delta}{2}$. In this case, as we show shortly in Lemma 28, the probability of obtaining a violating pair is at least $\frac{\delta}{4}(\frac{1}{2} - \delta) \geq \frac{\epsilon}{4} \cdot \gamma_0$. Therefore, f will be rejected with probability of at least $5/6$ in the second step of the algorithm (the closure-under-intersection test).

Thus, the probability that f is accepted by the algorithm is at most a $1/3$, as required. ■

Lemma 28 *Let δ be the distance of f to the closest singleton. If $\Pr[f(x) = 1] \geq \frac{1}{2} - \frac{\delta}{2}$, then the probability of obtaining a violating pair is at least $\frac{\delta}{4}(\frac{1}{2} - \delta)$.*

Proof: Let x_i be the closest singleton to f , so that $\Pr[f(x) \neq x_i] = \delta$. Define

$$\begin{aligned} G_1 &= \{x | f(x) = 1, x_i = 1\}, & B_1 &= F_1 \setminus G_1, \\ G_0 &= \{x | f(x) = 0, x_i = 0\}, & B_0 &= F_0 \setminus G_0. \end{aligned}$$

A simple counting argument shows that there are $(\frac{1}{2} - \delta)2^n$ disjoint pairs x, x' , such that: (1) $x \in G_1, x' \in G_0$; (2) x and x' differ only on the i 'th bit. To see why this is true, simply match each $x \in G_1$ to a point x' , which differs with x only on the i 'th bit. Thus, there are at least $|G_1| - |B_1|$ points $x \in G_1$ that must be matched to points $x' \in G_0$. But $|G_1| + |B_0| = 2^{n-1}$, and $|B_1| + |B_0| = \delta 2^n$ and therefore $|G_1| - |B_1| = (\frac{1}{2} - \delta)2^n$.

Now consider any point $y \in B_1$, and let $x \in G_1, x' \in G_0$ be a matched pair as defined above. Then $x \wedge y = x' \wedge y$, but $f(x) \wedge f(y) = 1$ while $f(x') \wedge f(y) = 0$. Therefore, either $f(x \wedge y) \neq f(x) \wedge f(y)$ or $f(x' \wedge y) \neq f(x') \wedge f(y)$, and so either y and x are a violating pair, or y and x' are a violating pair.

Since $\Pr[f(x) = 1] \geq \frac{1}{2} - \frac{\delta}{2}$, then $|G_1| + |B_1| \geq 2^n(\frac{1}{2} - \frac{\delta}{2})$. Using again the fact that $|G_1| - |B_1| = (\frac{1}{2} - \delta)2^n$, we get that $|B_1| \geq \delta 2^{n-2}$. It follows that the probability of obtaining a violating pair, is at least $\frac{\delta}{4}(\frac{1}{2} - \delta)$. ■

The above analysis breaks when f is actually almost $1/2 - \text{far}$ from every singleton, since in this case δ is close to $1/2$, and the probability $\frac{\delta}{4}(\frac{1}{2} - \delta)$ of obtaining a violating pair is not bounded from below. Another disadvantage of Algorithm 4 is the two sided error probability for testing singletons, as opposed to the one sided error we achieved in Algorithm 1 when we added the parity test.

Algorithm 4 can be generalized to testing k -monomials, with a query complexity of $O(1/\epsilon^2)$. The probability of choosing a violating pair can be shown to be at least $\frac{\delta}{4}(2^{-k} - \delta)$. Thus the requirement here is that δ will be strictly smaller than 2^{-k} . Notice that requiring that $\delta = O(1/2^k)$ is not really a restriction: every function f for which $\Pr[f(x) = 1]$ is approximately 2^{-k} is $O(2^{-k})$ -close to being a monomial, and our algorithm first verifies that in fact $\Pr[f(x) = 1]$ is approximately 2^{-k} . The restriction is in requiring that it be *strictly* smaller than 2^{-k} .

Another alternative test for singletons is to replace the relatively expensive test of checking whether $\Pr[f(x) = 1]$ is approximately $1/2$, by extending the notion of a violating pair. We will say that $x, y \in \{0, 1\}^n$ are a violating pair if $f(x \wedge y) \neq f(x) \wedge f(y)$ or if $f(x \vee y) \neq f(x) \vee f(y)$. Then in a similar way to the proof of Lemma 28, it can be shown that the probability of obtaining a violating pair is at least $\frac{\delta}{2}(\frac{1}{2} - \delta)$ (In this case the size of either B_0 or B_1 is at least $\delta 2^{n-1}$. Therefore choosing $y \in B_1$ or $y \in B_0$ and x, x' as before, will result in a violating pair either to the \wedge test or to the \vee test). The query complexity of this algorithm will be only $O(1/\epsilon)$, and it will have a one-sided error. Unfortunately this algorithm does not extend to testing monomials.

Further Research

Our results raise several questions that we believe may be interesting to study.

- Our algorithms for testing singletons and, more generally, monomials, apply two tests. The role of the first test is essentially to facilitate the analysis of the second, natural test (the closure under intersection test). The question is whether the first test is necessary.
- The query complexity of our algorithm for testing ℓ -term monotone DNF has a quadratic dependence in ℓ . While some dependence on ℓ seems necessary, we conjecture that a lower dependence is achievable. In particular, suppose we slightly relax the requirements of the testing algorithm and only ask that it rejects functions that are ϵ -far from any monotone DNF with at most $c \cdot \ell$ (or possibly ℓ^c) terms, for some constant c . Is it possible, under this relaxation, to devise an algorithm that has only polylogarithmic dependence on ℓ ?
- Finally, can our algorithm for testing monotone DNF functions be extended to testing general DNF functions?

Acknowledgments

We would like to thank Luca Trevisan for bringing to our attention the relation between testing the Long-Code and testing singletons. We would also like to thank Uri Feige and Adi Shamir for useful comments which led to an improvement in the complexity of the monomial test.

References

- [AFKS99] N. Alon, E. Fischer, M. Krivelevich, and M. Szegedy. Efficient testing of large graphs. In *Proceedings of FOCS*, pages 645–655, 1999.
- [AHR99] Y. Aumann, J. Håstad, M. Rabin, and M. Sudan. Linear consistency testing. In *Proceedings of RANDOM*, pages 109–120, 1999.
- [Ang88] D. Angluin. Queries and concept learning. *Machine Learning*, 2:319–342, 1988.
- [BCH⁺95] M. Bellare, D. Coppersmith, J. Håstad, M. Kiwi, and M. Sudan. Linearity testing in characteristic two. In *Proceedings of FOCS*, pages 432–441, 1995.
- [BEHW87] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Occam’s razor. *Information Processing Letters*, 24(6):377–380, April 1987.

- [BGS98] M. Bellare, O. Goldreich, and M. Sudan. Free bits, PCPs and non-approximability – towards tight results. *SIAM Journal on Computing*, 27(3):804–915, 1998.
- [BJT99] N. Bshouty, J. Jackson, and C. Tamon. More efficient PAC-learning of DNF with membership queries under the uniform distribution. In *Proceedings of COLT*, pages 286–295, 1999.
- [BLR93] M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. *JACM*, 47:549–595, 1993.
- [DGL⁺99] Y. Dodis, O. Goldreich, E. Lehman, S. Raskhodnikova, D. Ron, and A. Samorodnitsky. Improved testing algorithms for monotonicity. In *Proceedings of RANDOM*, pages 97–108, 1999.
- [Fis01] E. Fischer. The art of uninformed decisions: A primer to property testing. *The Bulletin of the European Association for Theoretical Computer Science*, 75:97–126, 2001.
- [FKR⁺02] E. Fischer, G. Kindler, D. Ron, S. Safra, and A. Smorodnitsky. Testing juntas. Unpublished manuscript, 2002.
- [GGL⁺00] O. Goldreich, S. Goldwasser, E. Lehman, D. Ron, and A. Samorodnitsky. Testing monotonicity. *Combinatorica*, 20(3):301–337, 2000.
- [GGR98] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *JACM*, 45(4):653–750, 1998.
- [GR97] O. Goldreich and D. Ron. Property testing in bounded degree graphs. In *Proceedings of STOC*, pages 406–415, 1997. To appear in *Algorithmica*.
- [Jac97] J. Jackson. An efficient membership-query algorithm for learning DNF with respect to the uniform distribution. *JCSS*, 55:414–440, 1997.
- [KS99] A. Klivans and R. Servedio. Boosting and hard-core sets. In *Proceedings of FOCS*, pages 624–633, 1999.
- [PRS01a] M. Parnas, D. Ron, and A. Samorodnitsky. Proclaiming dictators and juntas or testing boolean formulae. In *Proceedings of RANDOM01*, pages 273–284, 2001.
- [PRS01b] M. Parnas, D. Ron, and A. Samorodnitsky. Testing boolean formulae. Available from *ECCC*, <http://www.eccc.uni-trier.de/eccc/>, TR01-063, 2001.
- [Ron01] D. Ron. Property testing. In *Handbook of Randomized Computing, Volume II*, 2001.
- [RS96] R. Rubinfeld and M. Sudan. Robust characterization of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.
- [Rub99] R. Rubinfeld. Robust functional equations and their applications to program testing. *SIAM Journal on Computing*, 28(6):1972–1997, 1999.