Low Degree Test with Polynomially Small Error

Dana Moshkovitz *

January 31, 2016

Abstract

A long line of work in Theoretical Computer Science shows that a function is close to a low degree polynomial iff it is locally close to a low degree polynomial. This is known as low degree testing, and is the core of the algebraic approach to construction of PCP.

We obtain a low degree test whose error, i.e., the probability it accepts a function that does not correspond to a low degree polynomial, is polynomially smaller than existing low degree tests. A key tool in our analysis is an analysis of the sampling properties of the incidence graph of degree-k curves and k'-tuples of points in a finite space $\mathbb{F}^m$.

We show that the Sliding Scale Conjecture in PCP, namely the conjecture that there are PCP verifiers whose error is exponentially small in their randomness, would follow from a derandomization of our low degree test.

*dmoshkov@csail.mit.edu. Department of Electrical Engineering and Computer Science, MIT. This material is based upon work supported by the National Science Foundation under grants number 1218547 and 1452302. The paper originally appeared as "An Approach to the Sliding Scale Conjecture Via Parallel Repetition For Low Degree Testing" [24].

1 Introduction

1.1 Low Degree Test with Polynomially Small Error

A long line of work in Theoretical Computer Science shows that a function is close to a low degree polynomial iff it is *locally* close to a low degree polynomial (see, e.g., [19, 31, 4, 29, 20, 9, 27]). This is known as *low degree testing*, and is the core of the algebraic approach to construction of PCP, as well as of algebraic constructions of locally testable codes.

In this paper we consider the following formulation of low degree testing, which is conducive to the applications to PCP and codes. The parameters are a finite field $\mathbb{F}$, and natural numbers m and d. A verifier queries a constant number of non-interacting provers regarding a function $f : \mathbb{F}^m \to \mathbb{F}$ that is supposedly an m-variate polynomial of degree at most d. For instance, the verifier may ask the provers for the restriction of f to lines, planes, curves, etc. If f is indeed a polynomial of degree at most d, and the provers respond truthfully, then the verifier always accepts. In contrast, there exists $\varepsilon_0 > 0$, such that any prover strategy that the verifier accepts with non-negligible probability $\varepsilon \geq \varepsilon_0$ must correspond to a short list of polynomials of degree roughly d.

The most important parameter of a low degree test is its "error" ε_0. This parameter dominates the error probability of the PCP or the locally testable code. Additional parameters that are of interest are the randomness complexity of the verifier (i.e., the number of random bits that the verifier uses to generate the queries to the provers) and the answer size of the provers (i.e., the number of bits that comprise the answers of the provers). Previous work [4, 29, 27] showed that the error can be as low as $\varepsilon_0 = \text{poly}(m) \cdot (d/|\mathbb{F}|)^{\Omega(1)}$ with $O(m \log |\mathbb{F}|)$ randomness complexity and $\text{poly}(d, m, \log |\mathbb{F}|)$ answer size. For a sufficiently large field $\mathbb{F}$ (with respect to d and m), the error is $|\mathbb{F}|^{-\Omega(1)}$.

All low degree tests considered before inherently have error that is larger than $|\mathbb{F}|^{-1}$. A natural question is whether one can devise a test with lower error, e.g., $|\mathbb{F}|^{-k}$ for a parameter $k > 1$. This is the question that the current paper answers in the affirmative. We prove:

Theorem 1.1 (Low error low degree test; see Lemma 11.2)**.** *Assume that $\mathbb{F}$ is sufficiently large with respect to d, m and k (i.e., $|\mathbb{F}| = \text{poly}(m, d, k)$). Then, there is a low degree test whose error probability is $|\mathbb{F}|^{-\Omega(k)}$. The provers' answer size is $poly(d, k, \log |\mathbb{F}|)$, and the randomness of the verifier is $O(km \log |\mathbb{F}|)$.*

Moreover, since the tester of Theorem 1.1 queries the supposed restriction of the function on constant-dimensional surfaces, we get the following unusually strong property testing theorem, showing that one can probabilistically, to a very large degree of certainty, decide whether a function corresponds to a low degree polynomial, even in a very weak sense, by making a small number of queries to the function:

Theorem 1.2 (Property tester)**.** *Assume that $\mathbb{F}$ is sufficiently large with respect to d, m and k. Then, there is a randomized algorithm that, given $f : \mathbb{F}^m \to \mathbb{F}$, queries the evaluations of f on $\text{poly}(|\mathbb{F}|)$ points, and satisfies the following:*

- *If f is a polynomial of degree at most d, then the algorithm always accepts.*
- *There are $\delta = |\mathbb{F}|^{-\Theta(k)}$, $\varepsilon = |\mathbb{F}|^{-\Theta(1)}$, such that if f is ε-far from low degree (i.e., there is no polynomial of degree at most dk that agrees with f on at least ε fraction of the points $x \in \mathbb{F}^m$), then the probability that the algorithm rejects is at least $1 - \delta$.*

Note that the parameters ε and δ are extremely strong, since every function is $1/|\mathbb{F}|$-close to a constant function, and since δ is significantly lower than the probability $1/|\mathbb{F}|^{\Omega(1)}$ that was known before.

We prove Theorem 1.1 by developing a "direct product of low degree tests" to amplify the error of existing low degree tests (more on this in Section 1.2). Our analysis builds on a previous work by Imagliazzo, Kabanets and Wigderson [21]. We simplify and strengthen the IKW analysis and significantly improve the parameters it yields in our case. We stress that a direct application of IKW would have given a $2^{-\Omega(\sqrt{k})}$ error rather than $|\mathbb{F}|^{-\Omega(k)}$ error. At the original time of the paper's writing, it was not known how to improve IKW's work to error $2^{-\Omega(k)}$. A related theorem with error $2^{-\Omega(k)}$ was obtained by Dinur and Steurer [17] concurrently with our work. Moreover, it was not known how to have the base of the exponent as $|\mathbb{F}|^{-\Omega(1)}$ instead of $2^{-\Omega(1)}$. In the closely related setup of parallel repetition this was eventually obtained by several works [16, 25, 10]; the earliest of which was concurrent with our work. Our work combines the advantages of both lines of study.

As we explain in the sequel, our "direct product" is somewhat different than standard direct product. However, like the original direct product, the randomness complexity of our verifier is $O(k)$ *times* the randomness complexity of the original verifier. We show that if there were a way to implement our tester in a randomness-efficient way, so that the randomness complexity were $O(k \log |\mathbb{F}|)$ *plus* the randomness complexity of the original verifier (note that the error is exponentially small in $O(k \log |\mathbb{F}|)$), then one could have resolved the oldest open problem in PCP, known as "The Sliding Scale Conjecture" of Bellare, Goldwasser, Lund and Russell [7]. The conjecture is that in PCP with constant number of queries the error can be exponentially small in the randomness of the verifier.

Theorem 1.3 (From low degree test to Sliding Scale Conjecture)**.** *If there is a low degree test with randomness complexity $O(m \log |\mathbb{F}|)$, answer size $poly(d, m, \log |\mathbb{F}|)$, and error $|\mathbb{F}|^{-\Omega(m)}$ (see Conjecture 3.1), then the Sliding Scale Conjecture follows, i.e., there exists $c > 0$ such that*

$$NP \subseteq PCP_{1,1/n}[O(\log n), O(1)]_{[poly(n)]},$$

where $PCP_{c,s}[r, q]_{\Sigma}$ denotes the class of problems that have a PCP verifier with completeness c, soundness s, randomness r, number of queries q and where the proof is over an alphabet Σ.

For more details about the Sliding Scale Conjecture and its implications, see the previous version of this paper [24].

We note that a randomness-efficient low degree test obtained by amplifying a low degree test with higher error as in the current work must avoid known limitations on derandomized parallel repetition [18, 26]. We believe that this might be possible thanks to the algebraic structure of the low degree testing problem.

1.2 Direct Product of Low Degree Tests

In direct product testing a randomized verifier queries a constant number of provers regarding a function $s : [n] \to \Sigma$ they were supposed to agree on. The verifier's questions are about the restriction of s to different sets $A \subset [n]$ (i.e., $(s(i) : i \in A)$). The verifier tests probabilistically whether the provers are consistent by comparing the answers of the provers on intersecting sets. If the provers' responses are consistent with some function $s : [n] \to \Sigma$, the verifier always accepts. Moreover, the probability that the verifier accepts when the provers' responses do not

match one of a few functions $s_1, \ldots, s_l : [n] \to \Sigma$, is small. Direct product testing was introduced by Goldreich and Safra [20] and studied in many works since [15, 13, 21, 17].

In this work, we define an algebraic analogue of direct product testing for low degree testing. Instead of a function $s : [n] \to \Sigma$, we have a polynomial p of degree at most d over $\mathbb{F}$ that the provers are supposed to agree on. The verifier queries the provers about the supposed restriction of p to algebraic surfaces of degree k rather than to general sets. The algebraic structure allows the verifier to test for low degree. The actual test is a variation of the IKW test. The verifier picks three curves, where each curve intersects the previous curve in $k' < k/2$, $k' = \Theta(k)$ points. The verifier sends different curves and k'-tuples of points to different provers, who are supposed to respond with the restriction of the polynomial to their sets. The verifier checks consistency on the intersections. The large intersections enable the error to be exponentially small in k'. Crucially, the third curve is independent of the first curve, which forces the provers to be consistent with a global polynomial over $\mathbb{F}^m$.

Curves Test

1. Pick uniformly and independently at random:
   - a degree-k curve c_1;
   - two k'-tuples of points S_0, S_1 on c_1 such that $S_0 \cap S_1 = \phi$;
   - a degree-k curve c_2 that passes through S_1;
   - a k'-tuple of points S_2 contained in c_2 such that $S_1 \cap S_2 = \phi$;
   - a degree-k curve c_3 that passes through S_2;
   - a k'-tuple of point S_3 contained in c_3 such that $S_2 \cap S_3 = \phi$.
2. Send each of $S_0, S_1, S_2, S_3, c_1, c_2, c_3$ to a different prover.
   - For each curve c_i the prover sends a univariate polynomial p_i of degree at most dk that is supposed to be the restriction of the polynomial p to c_i.
   - For each tuple S_j the prover sends assignments a_j over $\mathbb{F}$ to the points in S_j. The assignments are supposed to be the evaluations of p on the points.
3. Check that $p_i(x) = a_j(x)$ for every $S_j \subseteq c_i$ and $x \in S_j$.

To prove Theorem 1.1 we end up analyzing a similar test that considers constant-dimensional surfacess that span whole lines instead of curves. This way we can directly use earlier works that analyze tests based on lines [4].

1.3 The Sampling Properties of Curves and Tuples

A key lemma in the proof of Theorem 1.1, and the main source for the quantitative improvement over IKW, is an analysis of the sampling properties of the incidence graph of "degree-k curves vs. k'-tuples" (similarly, "degree-k surfaces vs. k'-tuples"). This graph is the bipartite graph that has on one side all degree-k curves in a space $\mathbb{F}^m$, and on the other side all k'-tuples of points in $\mathbb{F}^m$. A curve is connected to a tuple if it contains it. Let S be a subset of μ fraction of the k'-tuples. We say that a curve ε-samples S if $\mu \pm \varepsilon$ fraction of the tuples on the curve are in S. We say that the curves vs. tuples graph is a (δ, ε)-sampler if for any subset S as above at least $1 - \delta$ fraction of the curves ε-sample S. We call δ the *sampling error*, and we call ε the *deviation error*. We show:

Lemma 1.1 (See Corollary 4.5). *Let $0 < \varepsilon < 1$. Then, for all k', the incidence graph "degree-k curves vs. k'-tuples" is a $(2\delta(k')/\varepsilon, 2k'\varepsilon)$-sampler for*

$$\delta(k') = \frac{k^k(k+1)}{\varepsilon^k(|\mathbb{F}| - k' + 1)^{(k-k'+1)/2}}.$$

For sufficiently large field $\mathbb{F}$ (polynomial size) with respect to k, and sufficiently large k (linear size) with respect to k', one gets deviation error $\varepsilon = |\mathbb{F}|^{-\Omega(1)}$ with sampling error $\delta = |\mathbb{F}|^{-\Omega(k)}$.

For $k' = 1$, it is well known that the "degree-k curves vs. points" graph has sampling error $\delta = |\mathbb{F}|^{-\Omega(k)}$ and deviation error $\varepsilon = |\mathbb{F}|^{-\Omega(1)}$. This follows from the $(k+1)$-wise independence of degree-k curves. Extending this argument to "degree-k curves vs. k'-tuples" for larger k' results in a large sampling error $\delta = |\mathbb{F}|^{-\Omega(k/k')}$. Similarly, it is shown in [21] that the graph "k-tuples vs. k'-tuples" has sampling error $\delta = exp(-k/k')$ with a small constant deviation error ε. The reason for the error $exp(-\sqrt{k})$ in [21] is taking $k' = \sqrt{k}$ as to balance $exp(-k/k')$ and $exp(-k')$. On the other hand, we show that the "degree-k curves vs. k'-tuples" incidence graph has sampling error $|\mathbb{F}|^{-\Omega(k-k')}$ while maintaining $\varepsilon = |\mathbb{F}|^{-\Omega(1)}$ deviation error (for sufficiently large field $\mathbb{F}$). This allows us to take $k' = \Theta(k)$, and achieve error $|\mathbb{F}|^{-\Omega(k)}$ for the repeated test.

Our approach to analyzing the sampling properties of "degree-k curves vs. k'-tuples" is to view the incidence graph of "degree-k curves vs. k'-tuples" as a k'-fold product of the incidence graph of "degree-k curves vs. points". With the appropriate choice of product, the sampling error of the product graph is k' times the sampling error of the initial graph. Hence, the sampling error of "degree-k curves vs. k'-tuples" is similar to that of "degree-k curves vs. points".

The advantage of this abstract view is that one can use our technique to analyze the sampling properties of general incidence graphs. Interestingly, this approach does not apply to the "k-tuples vs. k'-tuples" incidence graph relevant for [21]. The reason is that the deviation error accumulated in the k' applications of the product builds up, and – unlike in the curves graph – the initial deviation error is not sufficiently low to withstand that.

1.4 The Wide Agreement Lemma

Even with the improved sampling parameters, it is not clear that the analysis of IKW can be lifted to low degree testing in a way that proves Theorem 1.1. Briefly, the reason is that IKW relies on "majority decoding", while we have to deal with the "list decoding" regime. One of our main contributions is replacing the heart of IKW's analysis – the same part that uses samplers – with a new, stronger, lemma that we call "The Wide Agreement Lemma".

If our test passes with significant probability, then curves that intersect on k'-tuples of points often agree on their intersection. In contrast, to prove Theorem 1.1, we need to show that curves that intersect on a *single point* often agree on their intersection. In other words, we require a much wider agreement – not just between curves that have large intersection, but even among curves that have a small intersection. A-priori it could be that curves that disagree on an assignment to a point x typically do not share k'-tuples. However, since the incidence graph that has on one side curves that contain x and on the other side k'-tuples that contain x is a good sampler for every $x \in \mathbb{F}^m$, this cannot happen. The Wide Agreement Lemma appears as Lemma 8.1 in the sequel, and led to the later work [25].

1.5 Abstraction of Arora-Safra Composition

As a side-benefit of our proof of Theorem 1.3, we provide an abstract version of the Arora-Safra composition as described next.

Arora and Safra [3] were the first to suggest the technique of composition to decrease the number of queries (or alphabet) of a PCP verifier, leading to the first PCP with constant number of queries [2]. Since then, every proposed PCP construction (including the current one) used composition. Alas, the Arora-Safra composition was tailored to low degree extensions, and led to somewhat cumbersome and restricted usage.

In recent years there has been an attempt to formulate abstract composition lemmas that are widely applicable and lead to modular, easier to understand, constructions. One combinatorial method of composition was formulated by Szegedy [32], Dinur and Reingold [15] and Ben-Sasson et al [8]. Their works revealed the advantage of a "robust" PCP construction for composition. Robustness means that in the soundness case, not only that – with significant probability – the verifier rejects, but, in fact, the verifier's view is far from one that would have been accepted. Equivalently, the PCP is a "projection game" (the equivalence between robust PCPs and projection games is spelled out in [14]). A method of composition that preserves low soundness error and projection was discovered by Moshkovitz and Raz [28], and abstracted by Dinur and Harsha [14].

Interestingly, in contrast to all those composition techniques, the Arora-Safra composition does not require that the PCP being composed is robust. This is actually an advantage, because robust PCPs (equivalently, projection games) are harder to construct than general PCPs. In the high error regime there are various techniques for "robustization" (see, e.g., [15]), but in the low error regime we do not know how to transform a general PCP verifier into a robust PCP verifier with a comparable soundness error.

The composition we define (see Section 12.5 for the definition of composition and its analysis) works in the low error regime and does not require that the PCPs being composed are robust (and, appropriately, does not guarantee that the composed PCP is robust).

1.6 Organization

We start with preliminaries regarding error correcting codes, samplers and extractors, incidence graphs, curves, surfaces and polynomials over a finite space in Section 2. We formalize low degree testing in Section 3. We analyze the sampling properties of incidence graphs in Section 4. We prove a base low degree testing theorem in Section 5. We outline our analysis in Section 6, and in the next sections we provide the proof. We show how the Sliding Scale Conjecture follows from a derandomized low degree test in Section 12. We discuss ideas for further research in Section 13.

2 Preliminaries

In this section we introduce notions and notation that we use throughout this work, including error correcting codes, samplers and extractors, incidence graphs and curves over a finite space.

Throughout this work, *k'-tuple* means an ordered set of size k'.

2.1 Error Correcting Codes

An $(n, k, d)_\Sigma$ code C is a set of $\left|\Sigma^k\right|$ strings in Σ^n, where every two different strings $x, y \in C$ *agree* on at most d of their symbols, i.e.,

$$|\{\, i \in [n] \mid x_i = y_i \}| \leq d.$$

We often associate an *encoding* function $C : \Sigma^k \to \Sigma^n$ with C. Many times it is useful that the encoding is *systematic*, i.e., the first k symbols in the encoding $C(x)$ of some $x \in \Sigma^k$ are the symbols of x.

The following code construction follows from [1] using standard techniques (concatenation):

Proposition 2.1 (Code construction)**.** *For any $0 < \delta < 1$ and natural number k, there exists an $(n, k, (1-\delta)n)_\Sigma$ code where $n = O(k/\delta^2)$ and $|\Sigma| = O(1/\delta^2)$.*

The following bound on the number of codewords that can agree with a word follows from counting (our Proposition 3.1 uses a similar argument), and is a simplified version of Johnson's bound [22]:

Proposition 2.2 (List decoding bound)**.** *Let C be an $(n, k, d)_\Sigma$ code. For every $w \in \Sigma^n$ and $\delta \geq 2\sqrt{1 - d/n}$, there exist at most $2/\delta$ codewords in C that agree on at least δ fraction with w.*

2.2 Samplers and Extractors

For a graph $G = (V, E)$ and a vertex $v \in V$, the neighborhood of v in G is $N_G(v) = \{ u \in V \mid (v, u) \in E \}$.

A *sampler* is a bi-regular bipartite graph with a large part A and a small part B, in which, for any set $B' \subseteq B$, almost every vertex in A has about $|B'|/|B|$ fraction of its neighbors landing in B':

Definition 2.1 (Sampling)**.** *For $0 < \delta, \varepsilon < 1$, we say that a bi-regular bipartite graph $G = (A, B, E)$ is a (δ, ε)-sampler if for any set $B' \subseteq B$, $\mu = |B'|/|B|$, for a uniformly distributed $a \in A$, it holds that*

$$\left| \frac{|N_G(a) \cap B'|}{|N_G(a)|} - \frac{|B'|}{|B|} \right| \leq \varepsilon,$$

with probability at least $1 - \delta$.

We call δ the *sampling error* and ε the *deviation error*.

An *extractor* is a function that maps a distribution X' with sufficient "randomness" over a large space X to a distribution that is approximately uniform over a small space Z. The mapping is probabilistic, and uses an independent uniform "seed" $y \in Y$. The function maps $x \in X$ and $y \in Y$ to $z \in Z$ and an additional $w \in W$, so (z, w) uniquely defines (x, y). The randomness of X' is measured using *min-entropy*, and is $H_\infty(X') = \log(1/\max_x \Pr[X' = x])$.

Definition 2.2 (Extractor)**.** *A 1-1 function $Ext : X \times Y \to Z \times W$ is a (δ, ε)-extractor if for any distribution X' over X, $H_\infty(|X'|) \geq \log(\delta |X|)$, the probability distribution defined*[1] *by $Ext(X', Y)$ on Z, is ε-close to uniform over Z.*

X' is called the *randomness source*. The elements in Y are called the *seeds* of the extractor. Often extractors are defined without W and without being 1-1, but incorporating W will be useful for us, and this convention – the "rotation map" of [30] – has been used in the past. We associate a bipartite graph with Ext: the graph is on vertices $X \cup Z$ and it has an edge (x, z) if there are $y \in Y$ and $w \in W$ such that $Ext(x, y) = (z, w)$. The elements in Y enumerate the neighbors of a vertex in X, while the elements of W enumerate the neighbors of a vertex in Z.

Zuckerman observed that the notions of sampler and extractor are closely related:

[1] This distribution is sampled by picking uniformly at random $x \in X'$ and $y \in Y$, computing $Ext(x, y) = (z, w)$, and outputting z.

Proposition 2.3 ([34])**.** *The following hold:*

1. *If $Ext : X \times Y \to Z \times W$ is a (δ, ε)-extractor, then the bipartite graph on $X \cup Z$ associated with it is a $(2\delta, \varepsilon)$-sampler.*
2. *If (X, Z, E) is a (δ, ε)-sampler, then a corresponding function $Ext : X \times Y \to Z \times W$ is, for any $\delta' \geq \delta$, a $(\delta', \varepsilon + \delta/\delta')$-extractor.*

2.3 Curves, Surfaces and Polynomials

Let $\mathbb{F}$ be a finite field. Let m, k and r be natural numbers. A degree-k *curve* in $\mathbb{F}^m$ is a function $c : \mathbb{F} \to \mathbb{F}^m$ such that there exist m univariate degree-k polynomials $c_1, \ldots, c_m$ where $c(t) = (c_1(t), \ldots, c_m(t))$. We often associate a curve with its image $c(\mathbb{F})$. A *line* is a degree-1 curve. A dimension-r degree-k *surface* in $\mathbb{F}^m$ is a function $s : \mathbb{F}^r \to \mathbb{F}^m$ such that there exist m r-variate degree-k polynomials $s_1, \ldots, s_m$ where $s(t_1, \ldots, t_r) = (s_1(t_1, \ldots, t_r), \ldots, s_m(t_1, \ldots, t_r))$. We often associate a surface with its image $s(\mathbb{F}^r)$. A *curve* is a dimension-1 surface.

For $T = \{t_1, \ldots, t_k\} \subseteq \mathbb{F}$ and $1 \leq i \leq k$, we use Lagrange interpolation to define $I_{T,i}$ as the degree-$(k-1)$ polynomial that is 1 on t_i and 0 on $T - \{t_i\}$:

$$I_{T,i}(t) \doteq \frac{\prod_{j \in T - \{t_i\}} (t - t_j)}{\prod_{j \in T - \{t_i\}} (t_i - t_j)}.$$

Fixing $T = \{t_1, \ldots, t_k\} \subseteq \mathbb{F}$, for every k-tuple of points $X = \{x_1, \ldots, x_k\} \subseteq \mathbb{F}^m$ we can interpolate the degree-$(k-1)$ curve c_X that passes through $x_1, \ldots, x_k$ in positions $t_1, \ldots, t_k$ as:

$$c_X(t) = \sum_{i=1}^{k} x_i \cdot I_{T,i}(t).$$

2.4 Incidence Graphs

In this work we are interested in bipartite graphs that correspond to set inclusion:

Definition 2.3 (Incidence graph)**.** *Let $\mathcal{U}$ be a set. Let A and B be families of subsets of $\mathcal{U}$. The* incidence graph $\mathcal{G}(A, B)$ *is the bipartite graph on A and B in which a vertex $a \in A$ is connected to a vertex $b \in B$ if $b \subseteq a$.*

A few examples of incidence graphs are:

1. "k-tuples vs. k'-tuples": $\mathcal{U}$ is a finite set. A is the family of all k-tuples of points in $\mathcal{U}$, and B is the family of all k'-tuples of points in $\mathcal{U}$.
2. "degree-k curves vs. k'-tuples": $\mathcal{U} = \mathbb{F}^m$ for a finite field $\mathbb{F}$ and a natural number m. A is the set of all degree-k curves in $\mathbb{F}^m$; B consists of all k'-tuples of points in $\mathbb{F}^m$.
3. "degree-k curves vs. points": A special case of "degree-k curves vs. k'-tuples" in which $k' = 1$, so B corresponds to the family of points[2] in $\mathbb{F}^m$.

[2]We will often use the shorthand $B = \mathbb{F}^m$ in this case, even though Definition 2.3 talked about B that consists of subsets.

3 Low Degree Testing

Let $\mathbb{F}$ be a finite field and let m, v, d and k be natural numbers. In this section we define low degree testing for m-variate polynomials of degree at most d over $\mathbb{F}$ by querying v-dimensional surfaces of degree at most k in $\mathbb{F}^m$, as well as querying k'-tuples of points in $\mathbb{F}^m$. We generalize to surfaces as opposed to just curves in order to capture existing tests on planes and low-dimensional subspaces. In particular, our surfaces will be spanned by lines and curves, so we can apply existing analyses of lines tests.

One is advised to think of the parameters as follows:

- $|\mathbb{F}|$ is large with respect to d and m. Typically, $|\mathbb{F}| = \text{poly}(d, m)$.
- We typically take v to be a small constant, possibly 1.
- We typically take $k \leq d$.
- We take k' to be smaller than k, but often of the same order of magnitude as k.

The set of all surfaces that may be queried is denoted $\mathcal{C}$, and the set of k'-tuples that may be queried is denoted $\mathcal{I}$. In this work $\mathcal{I}$ will always be the set of all k'-tuples of points in $\mathbb{F}^m$. For a set $S \subseteq \mathbb{F}^m$, we denote the set of all surfaces in a family $\mathcal{C}' \subseteq \mathcal{C}$ that pass through S by $\mathcal{C}'_S$.

Assignments[3] to v-dimensional surfaces of degree at most k in $\mathbb{F}^m$ are supposedly the restrictions of a single m-variate degree-d polynomial to the surface – in which case we say that they *agree* with the polynomial – and in any case are v-variate polynomials of degree at most dk over $\mathbb{F}$. Assignments to k'-tuples of points in $\mathbb{F}^m$ are supposedly the restrictions of the same m-variate degree-d polynomial to the points – in which case we say that they *agree* with the polynomial – and in any case are k'-tuples of values in $\mathbb{F}$.

A *low degree test* is specified by a verifier that makes a constant number of *queries* to surfaces and tuples, receives the assignments to the surfaces and tuples, and either accepts or rejects. The *randomness* of the low degree test is the number of random bits used by the verifier. We say that the low degree test has *perfect completeness* if the verifier always accepts if whenever it queries a surface or a tuple it gets the restriction of a single m-variate degree-d polynomial over $\mathbb{F}$. All the tests that we consider in this work have perfect completeness. We say that the tester is *uniform*, if the distribution of each of its queries is uniform over all surfaces in $\mathcal{C}$ or tuples in $\mathcal{I}$. All the tests that we consider in this work are uniform.

3.1 Initial Points

Let $q < v + k$ be a natural number. For the application to PCP we allow the embedding of q-tuples of points in $\mathbb{F}^m$ in the surfaces we consider. *Initial conditions* are given as a collection of q-tuples of points $\{(x_{i,1}, \ldots, x_{i,q})\}_{i=1}^{M}$ where $x_{i,1}, \ldots, x_{i,q} \in \mathbb{F}^m$. Typically, $M \leq |\mathbb{F}^m|$. We fix $T \subseteq \mathbb{F}^v$, $|T| = q$. We say that a family $\mathcal{C}$ of surfaces satisfies the conditions at T if each surface $c \in \mathcal{C}$ passes through $x_{i,1}, \ldots, x_{i,q}$ at positions T for some $1 \leq i \leq M$, and each q-tuple is contained this way in the same number of surfaces in $\mathcal{C}$. In PCP constructions the initial conditions are typically concentrated in a small sub-cube in $\mathbb{F}^m$, and the verifier refrains from

[3] In the context of multi-prover protocols, it is natural to consider several assignments, one for each prover, while in the context of PCP it is natural to consider a single assignment. However, even in the multi-prover context one can assume without loss of generality that there is only one assignment, provided that the test randomly picks which prover to query for each query it makes.

comparing the surfaces on them. Hence, we adapt our low degree tests as to allow "forbidden points" that the verifier does not use for comparisons:

Definition 3.1 (Forbidden points)**.** Forbidden points *are defined by a function* $Q : \mathcal{C} \to 2^{\mathbb{F}^v}$*. For a surface* $c \in \mathcal{C}$*, let* $c^{-Q} \doteq c(\mathbb{F}^v - Q(c))$*. For a family of surfaces* $\mathcal{C}$*, we will use the notation* $\mathcal{C}^{-Q}$ *to refer to* $\left\{ c^{-Q} \,\middle|\, c \in \mathcal{C} \right\}$*.*

In this work we consider forbidden points where $|Q(c)|$ is the same for all $c \in \mathcal{C}$, and we define $|Q|$ to be this number.

3.2 A Variety of Low Degree Testers

The low degree testers that we consider in this work are:

1. Surface-vs.-Surface Test: compares two surfaces that intersect in a k'-tuple. This is a generalization of Curve-vs.-Curve Test.
2. Surface-vs.-Surface-on-Point Test: compares two surfaces that intersect on a k'-tuple, but only on a random point in the k'-tuple.
3. Surfaces Test: compares three surfaces and four k'-tuples on the three surfaces.

Surface-vs.-Surface Test is parameterized by two families of surfaces, $\mathcal{C}_1$ and $\mathcal{C}_2$, forbidden points $Q_i : \mathcal{C}_i \to 2^{\mathbb{F}}$, and a family $\mathcal{I}$ of tuples. In this tester and in similar testers: If $\mathcal{C}_2$ and Q_2 are omitted, it should be understood that $\mathcal{C}_2 = \mathcal{C}_1$ and $Q_2 = Q_1$.

Surface-vs.-Surface Test$(\mathcal{C}_1, \mathcal{C}_2, Q_1, Q_2, \mathcal{I})$

1. Pick uniformly $c_1 \in \mathcal{C}_1$; pick $S \in \mathcal{I}$ uniformly such that $S \subseteq c_1^{-Q_1}$; pick $c_2 \in \mathcal{C}_2$ uniformly such that $S \subseteq c_2^{-Q_2}$.
2. Check that $\mathcal{A}(c_1)(x) = \mathcal{A}(c_2)(x)$ for every $x \in S$.

When the surfaces are one dimensional, we refer to the test as Curve-vs.-Curve Test. When the curves are lines, we refer to the test as Line-vs.-Line Test. Surface-vs.-Surface-on-Point Test is similar to Surface-vs.-Surface Test, except that it only compares the two surfaces on a random point in their intersection. It is mainly useful an auxiliary test for the analysis:

Surface-vs.-Surface-on-Point Test$(\mathcal{C}_1, \mathcal{C}_2, Q_1, Q_2, \mathcal{I})$

1. Pick uniformly $c_1 \in \mathcal{C}_1$; pick $S \in \mathcal{I}$ uniformly such that $S \subseteq c_1^{-Q_1}$; pick $c_2 \in \mathcal{C}_2$ uniformly such that $S \subseteq c_2^{-Q_2}$.
2. Pick uniformly at random a point $x \in S$.
3. Check that $\mathcal{A}(c_1)(x) = \mathcal{A}(c_2)(x)$.

When the intersections between curves are points (i.e., $\mathcal{I} = \mathbb{F}^m$, $k' = 1$), Surface-vs.-Surface-on-Point Test and Surface-vs.-Surface Test are equivalent. Surfaces Test queries three surfaces from a family $\mathcal{C}$ with forbidden points $Q : \mathcal{C} \to 2^{\mathbb{F}}$, and four k'-tuples from a family $\mathcal{I}$.

SURFACES TEST$(\mathcal{C}, Q, \mathcal{I})$

1. Pick uniformly and independently at random a surface $c_1 \in \mathcal{C}$, tuples $S_0, S_1 \in \mathcal{I}$, $S_0, S_1 \subseteq c_1^{-Q}$, $S_0 \cap S_1 = \phi$, a surface $c_2 \in \mathcal{C}$, $S_1 \subseteq c_2^{-Q}$, a tuple $S_2 \in \mathcal{I}$, $S_2 \subseteq c_2^{-Q}$, $S_1 \cap S_2 = \phi$, a surface $c_3 \in \mathcal{C}$, $S_2 \subseteq c_3^{-Q}$, and a tuple $S_3 \subseteq c_3^{-Q}$, $S_3 \in \mathcal{I}$, $S_2 \cap S_3 = \phi$.

2. Check that $\mathcal{A}(c_1)$ agrees on S_0 with $\mathcal{A}(S_0)$; $\mathcal{A}(c_1)$ and $\mathcal{A}(c_2)$ agree on S_1 with $\mathcal{A}(S_1)$; $\mathcal{A}(c_2)$ and $\mathcal{A}(c_3)$ agree on S_2 with $\mathcal{A}(S_2)$; $\mathcal{A}(c_3)$ agrees on S_3 with $\mathcal{A}(S_3)$.

One could consider a variant of SURFACES TEST that queries only curves and not k'-tuples, but the test we defined is easier to analyze, and hence we prefer it.

3.3 Low Degree Testing Theorems: Proximity and List Decoding

Let $\varepsilon > 0$ be a function of $|\mathbb{F}|$, d and m (typically $\varepsilon \approx d/|\mathbb{F}|$). Let d' be a natural number (typically $d' \approx d$). There are several soundness guarantees we consider for low degree tests:

- *Surface (Tuple) Proximity:* Let $\gamma' : [0,1] \to [0,1]$ (typically, $\gamma'(\gamma) = \gamma - \varepsilon$). For every $\gamma \geq \varepsilon$, if the verifier accepts with probability γ, then there exists an m-variate polynomial of degree at most d' over $\mathbb{F}$ that agrees with $\gamma' = \gamma'(\gamma)$ fraction of the surfaces in $\mathcal{C}$ (resp., tuples in $\mathcal{I}$). To denote that this statement holds we write $\mathsf{AgrErr}^{\mathcal{C}}_{\gamma\to\gamma',d\to d'}(Test) \leq \varepsilon$ (resp., $\mathsf{AgrErr}^{\mathcal{I}}_{\gamma\to\gamma',d\to d'}(Test) \leq \varepsilon$).

- *Surface (Tuple) List decoding:* Let $l : [0,1] \to \mathbb{N}$ (typically, $l(\gamma) = O(1/\gamma)$). For every $\gamma \geq \varepsilon$, there exist m-variate polynomials $p_1, \ldots, p_l$, $l = l(\gamma)$, of degree at most d' over $\mathbb{F}$ such that the probability that the verifier accepts yet the assignments to the surfaces (resp., tuples) it picked do not agree with one of $p_1, \ldots, p_l$, is at most γ. To denote that this statement holds we write $\mathsf{ListErr}^{\mathcal{C}}_{l,d'}(Test) \leq \varepsilon$ (resp., $\mathsf{ListErr}^{\mathcal{I}}_{l,d'}(Test) \leq \varepsilon$).

It is straightforward to show that a low degree testing theorem in list decoding form implies a theorem in proximity form, since one of the polynomials in the list has to agree with at least $\gamma'(\gamma) \doteq (\gamma - \varepsilon)/l(\gamma)$ fraction of the tuples. Next we show that the other direction holds as well, i.e., from a low degree testing in proximity form, one can deduce the list decoding form. Below we outline the argument for tuples, since this is what we will use later.

First, we need the following proposition which uses the error correction properties of polynomials, and the sampling properties of the family of all k'-tuples of points in $\mathbb{F}^m$. The Proposition extends Proposition 2.2.

Proposition 3.1 (Short list decoding)**.** *For $\delta_0 = (d'/|\mathbb{F}|)^{k'}$, for every assignment $\mathcal{A}$ of elements in $\mathbb{F}^{k'}$ to tuples in $\mathcal{I}$, and any $\delta \geq 2\sqrt{\delta_0}$, there are at most $2/\delta$ m-variate polynomials $p_1, \ldots, p_l$ of degree at most d' over $\mathbb{F}$, such that*

$$\Pr_{S\in\mathcal{I}}\left[\mathcal{A}(S) \equiv p_{|S}\right] > \delta.$$

Proof. Assume on way of contradiction that there are different m-variate polynomials $p_1, \ldots, p_l$ of degree at most d' over $\mathbb{F}$ with $\Pr_{c\in\mathcal{C}}\left[\mathcal{A}(c) \equiv p_{|c}\right] > \delta$ for $l = 1 + \lfloor 2/\delta \rfloor$.

For $1 \leq i < j \leq l$, the polynomials p_i and p_j can agree on at most $d'/|\mathbb{F}|$ fraction of the points in $\mathbb{F}^m$. For at most $\delta_0 = (d'/|\mathbb{F}|)^{k'}$ fraction the tuples, the polynomials p_i and p_j agree on the tuple.

By inclusion-exclusion, the number of tuples that agree with one of $p_1, \ldots, p_l$ can be lower bounded by:

$$l\delta \left|\mathcal{I}\right| - \binom{l}{2}\delta_0 \left|\mathcal{I}\right|.$$

We have $l\delta > 2$ and $\binom{l}{2} \leq 1/\delta_0$, which implies that $|\mathcal{I}| > |\mathcal{I}|$ – contradiction! □

Proposition 3.2 (Proximity ⇒ List decoding)**.** *Let $\delta' = \gamma'(\delta) - |\mathbb{F}|^{-k'}$. Assume that $\delta' \geq 2(d'/|\mathbb{F}|)^{k'/2}$. Then, for any low degree tester* Test*,*

$$AgrErr^{\mathcal{I}}_{\gamma\to\gamma',d\to d'}(\textsc{Test}) \leq \delta \;\Rightarrow\; ListErr^{\mathcal{I}}_{2/\delta',d'}(\textsc{Test}) \leq \delta$$

Proof. Let $\delta^* = \gamma'(\delta)$. Let $\delta' = \delta^* - |\mathbb{F}|^{-k'}$, so $\delta' \geq 2(d'/|\mathbb{F}|)^{k'/2}$. Let $p_1, \ldots, p_l$ be all the m-variate polynomials of degree at most d over $\mathbb{F}$ that agree with $\mathcal{A}$ on at least δ' fraction of the tuples $S \in \mathcal{I}$. By Proposition 3.1, we have $l \leq 2/\delta'$. We will upper bound by δ the probability that the test passes, yet the verifier picks $S \in \mathcal{I}$ such that $\mathcal{A}(S) \notin \{p_{1|S}, \ldots, p_{l|S}\}$ (this will imply the lemma). Assume, toward a contradiction, that this is not the case.

For every tuple $S \in \mathcal{I}$ such that $\mathcal{A}(S) \in \{p_{1|S}, \ldots, p_{l|S}\}$, define $\mathcal{A}^*(S)$ to be a random element in $\mathbb{F}^{k'}$. By our assumption, the probability that Test passes for $\mathcal{A}^*$ is at least δ. Since $AgrErr^{\mathcal{I}}_{\gamma\to\gamma',d\to d'}(\textsc{Test}) \leq \varepsilon$, there is an m-variate polynomial p^* of degree at most d' over $\mathbb{F}$ that agrees with $\mathcal{A}^*$ on at least $\gamma'(\delta) = \delta^*$ fraction of the tuples $S \in \mathcal{I}$. The probability that $\mathcal{A}^*(S) = p^*_{|S}$ on those tuples S for which $\mathcal{A}^*(S)$ was chosen randomly is $|\mathbb{F}|^{-k'}$, and thus p^* must agree with $\mathcal{A}$ on at least $\delta^* - |\mathbb{F}|^{-k'} = \delta'$ fraction of the tuples. Thus $p^* \equiv p_j$ for some $1 \leq j \leq l$, and p_j agrees with $\mathcal{A}^*$ with probability at least δ' over the tuples. This is a contradiction! □

3.4 A Theorem And A Conjecture

In this work we give the first low degree test whose soundness error can be made $|\mathbb{F}|^{-k}$ for arbitrarily large $k \geq 1$. The low degree testing theorem follows from applying our direct product theorem (Theorem 6.1) on the low degree testing theorem in Section 5. The family of surfaces used is specified in Section 5 as well.

Theorem 3.1 (Low error low degree testing theorem)**.** *Let $\mathbb{F}$ be a finite field that is large enough (polynomial size) with respect to m, k', q and d, and fix initial conditions $\{(x_{i,1}, \ldots, x_{i,q})\}_{i=1}^{M} \subseteq (\mathbb{F}^m)^q$. Then, there is a family $\mathcal{C}$ of surfaces, $|\mathcal{C}| \leq M\,|\mathbb{F}|^{O(mk')}$, that satisfies the initial conditions with forbidden points $Q : \mathcal{C} \to 2^{\mathbb{F}}$; in which the surfaces are of degree $k = \Theta(k' + q)$ and dimension $v = O(1)$; and it holds:*

$$\mathsf{ListErr}^{\mathcal{C}}_{|\mathbb{F}|^{O(k')},dk}(\textsc{Surfaces Test}(\mathcal{C}, Q, (\mathbb{F}^m)^{k'})) \leq |\mathbb{F}|^{-\Omega(k')}.$$

We conjecture that there is a low degree test whose soundness error can be made $\approx 1/|\mathbb{F}^m|$ when the randomness is only $O(m \log |\mathbb{F}|)$ (note that the verifier can only access a number of curves and tuples that is exponential in its randomness). As we show in Section 12, the conjecture would imply the Sliding Scale Conjecture:

Conjecture 3.1 (Derandomized low degree test conjecture)**.** *Let $\mathbb{F}$ be a finite field that is large enough (polynomial size) with respect to m, k', q and d, and fix initial conditions $\{(x_{i,1}, \ldots, x_{i,q})\}_{i=1}^{M} \subseteq (\mathbb{F}^m)^q$. Then, there exist:*

1. *A family $\mathcal{C}$ of surfaces that satisfies the initial conditions, and in which the surfaces are of degree $k = poly(k', q, d)$ and dimension $v = O(1)$;*

2. *A family $\mathcal{I}$ of k'-tuples of points in $\mathbb{F}^m$;*

3. *A low degree tester* TEST *that uses $O((m+k')\log|\mathbb{F}| + \log M)$ random bits to make $O(1)$ queries to $\mathcal{C}$ and $\mathcal{I}$, so*

$$\mathsf{ListErr}^{\mathcal{C}}_{|\mathbb{F}|^{O(k')},\mathrm{poly}(d,k)}(\textsc{Test}) \leq |\mathbb{F}|^{-\Omega(k')}.$$

4 Curve-Tuple Sampling

In this section we explore the sampling properties of the "degree-k curves vs. k'-tuples" ($k > k'$) incidence graph[4]. This argument yields low sampling error for $k' = \Theta(1)$. Then, we show that the "degree-k curves vs. k'-tuples" graph can be viewed as a k'-product of the "degree-k curves vs. 1-tuples". We use this connection to argue that the graph for k'-tuples has essentially the same sampling error as the graph for 1-tuples, albeit with larger deviation.

Interestingly, while the larger deviation is too large for the "k-tuples vs. k'-tuples" graph (the graph relevant to IKW [21]), it is sufficiently small when k-tuples are replaced with degree-k curves.

4.1 The k/k'-wise Independence Argument

Let $B \subseteq (\mathbb{F}^m)^{k'}$, $|B| = \mu\left|(\mathbb{F}^m)^{k'}\right|$. Pick $c \in \mathcal{C}$ uniformly at random. For a k' tuple $T = \{t_1, \ldots, t_{k'}\} \subseteq \mathbb{F}$, let X_T indicate whether $(c(t_1), \ldots, c(t_{k'})) \in B$, and let $\hat{X}_T = X_T - \mu$. Since each k'-tuple appears in the same number of curves, we have $\mathbf{E}\left[\hat{X}_T\right] = 0$. Define $\hat{X} \doteq \binom{|\mathbb{F}|}{k'}^{-1} \sum_{T \subseteq \mathbb{F}} \hat{X}_T$.

Proposition 4.1 (l'th Moment)**.** *Let $2 \leq l \leq k/k'$.*

$$\mathbf{E}\left[\hat{X}^l\right] \quad \leq \quad |\mathbb{F}|^{-l/2}\, k^l \mu(l+1).$$

Proof.

$$\begin{aligned}
\mathbf{E}\left[\hat{X}^l\right] &= \binom{|\mathbb{F}|}{k'}^{-l} \cdot \mathbf{E}\left[\left(\sum_{T \subseteq \mathbb{F}} \hat{X}_T\right)^l\right] \\
&= \binom{|\mathbb{F}|}{k'}^{-l} \cdot \mathbf{E}\left[\sum_{T_1,\ldots,T_l \subseteq \mathbb{F}} \hat{X}_{T_1} \cdots \hat{X}_{T_l}\right] \\
&= \binom{|\mathbb{F}|}{k'}^{-l} \cdot \sum_{T_1,\ldots,T_l \subseteq \mathbb{F}} \mathbf{E}\left[\hat{X}_{T_1} \cdots \hat{X}_{T_l}\right] \qquad (1)
\end{aligned}$$

[4]Our proof readily extends from curves to surfaces.

For every l pairwise disjoint $T_1, \ldots, T_l \subseteq \mathbb{F}$, we have that $\hat{X}_{T_1}, \ldots, \hat{X}_{T_l}$ are independent. Hence, if among $T_1, \ldots, T_l \subseteq \mathbb{F}$ there is at least one $1 \leq i \leq l$ such that T_i is disjoint from the other T_j for $j \neq i$, we have

$$\mathbf{E}\left[\hat{X}_{T_1} \cdots \hat{X}_{T_l}\right] = \mathbf{E}\left[\hat{X}_{T_i}\right] \cdot \mathbf{E}\left[\hat{X}_{T_1} \cdots \hat{X}_{T_{i-1}} \hat{X}_{T_{i+1}} \cdots \hat{X}_{T_l}\right] = 0.$$

Therefore, the only terms that survive in (1) are those where every T_i has non-empty intersection with $\bigcup_{j \neq i} T_j$ (for this we need $l \geq 2$). Their number is bounded by $\binom{|\mathbb{F}|}{k'}^l |\mathbb{F}|^{-l/2} k^l$, since the T_i's pick $lk' \leq k$ elements from $\mathbb{F}$ with multiplicities, and at least $l/2$ times the element that is picked is one of at most k elements. Each term can be bounded by:

$$\begin{aligned}
\mathbf{E}\left[\hat{X}_{T_1} \cdots \hat{X}_{T_l}\right] &\leq \Pr\left[\exists i \hat{X}_{T_i} = 1 - \mu\right] \cdot (1 - \mu) + \Pr\left[\forall i \hat{X}_{T_i} = -\mu\right](-\mu)^l \\
&\leq l \cdot \mu \cdot 1 + 1 \cdot \mu \\
&= \mu(l + 1).
\end{aligned}$$

The proposition follows. □

As a corollary we get a proof of the sampling property of $\mathcal{G}(\mathcal{C}, (\mathbb{F}^m)^{k'})$:

Proposition 4.2. *For $l \leq k/k'$, the incidence graph $\mathcal{G}(\mathcal{C}, (\mathbb{F}^m)^{k'})$ is $\mu |\mathbb{F}|^{-l/2} k^l (l+1)\varepsilon^{-l}$-sampling.*

Proof. By Markov's inequality,

$$\Pr\left[\hat{X} \geq \varepsilon\right] \leq \Pr\left[\hat{X}^l \geq \varepsilon^l\right] \leq \frac{\mathbf{E}\left[\hat{X}^l\right]}{\varepsilon^l}.$$

The proposition follows from Proposition 4.1. □

In the sequel we will also need an analysis of the sampling properties of $\mathcal{G}(\mathcal{C}_S, \mathcal{I}_S)$ where $\mathcal{C}_S$ is the family of all the degree-k curves through a small set of points $S \subseteq \mathbb{F}^m$, $|S| \ll k'$, and $\mathcal{I}_S$ is the family of all k'-tuples of points in $\mathbb{F}^m$ that contain the points in S. Such an analysis follows along the same lines as above.

For $k' = 1$ we recover the standard upper bound of $\approx |\mathbb{F}|^{-k}$ on the sampling error of "degree-k curves vs. points".

Corollary 4.3. *For sufficiently large field $|\mathbb{F}| = \Theta(k)$, the "degree-k curves vs. points" incidence graph is an $(|\mathbb{F}|^{-\Theta(k)}, |\mathbb{F}|^{-\Theta(1)})$-extractor.*

However, for larger k' we get a much weaker upper bound of $\approx |\mathbb{F}|^{-l}$.

It is instructive to have an example in mind for when a sampling error of $\approx |\mathbb{F}|^{-l} k^l$ occurs:

Example 4.1. *Let $X \subseteq \mathbb{F}^m$ be a set of points of fraction $\mu = |X| / |\mathbb{F}^m|$ to be determined later; let $B \subseteq (\mathbb{F}^m)^{k'}$ be the family of k'-tuples in which the lexicographically first element lands in X; let $A \subseteq \mathcal{C}$ be the family of degree-k curves in which the first l points according to the lexicographic order land in X. Then the probability mass of B is μ; the probability mass of A is μ^l; given that $c \in A$ and $S \subseteq c$, $S \in (\mathbb{F}^m)^{k'}$, the probability that $S \in B$ is[5] $\approx lk'/|\mathbb{F}|$. Pick μ so $\mu < lk'/|\mathbb{F}| - \varepsilon$. The sampling error is roughly $(k/|\mathbb{F}| - \varepsilon)^l$.*

Note that Example 4.1 works only when $\varepsilon < lk'/|\mathbb{F}|$. For larger $\varepsilon = |\mathbb{F}|^{-\Theta(1)}$ we will be able to get error $\approx |\mathbb{F}|^{-k}$ rather than $\approx |\mathbb{F}|^{-l}$ in Section 4.2.

[5] In contrast, for "k-tuples vs. k'-tuples" the probability would have been $\approx lk'/k$; the difference is crucial for understanding why our approach in Section 4.2 works for the algebraic case, but not for direct products.

4.2 Extractor Product

In this section we define a replacement product operation on extractors, and use it to prove a much lower sampling error for "degree-k curves vs. k'-tuples" than the one proved in Proposition 4.2. Replacement product turns out to have been defined before in [11].

Replacement product is a generalization of a widely-used transformation by Wigderson and Zuckerman [33]. Both transformations take two extractors Ext_1 and Ext_2 and generate a new extractor whose output is the multiplication of the output of Ext_1 and the output of Ext_2. The new extractor requires independent seeds for Ext_1 and Ext_2. The difference between our operation and the Wigderson-Zuckerman one is that WZ require Ext_2 to work for the same domain as Ext_1, and handle a lower min-entropy than Ext_1. In our operation the domain of Ext_2 is potentially much smaller than the domain of Ext_1, and there is no similar demand on the min-entropy of Ext_2. This allows Ext_2 to have a smaller seed, and in certain settings may allow for exhaustive search of a construction of Ext_2 with optimal parameters.

Definition 4.1 (Replacement product for extractors)**.** *Suppose $Ext : X_1 \times Y_1 \to Z_1 \times X_2$ is an extractor, and $\{Ext_z : X_2 \times Y_2 \to Z_2 \times W_2\}_{z \in Z_1}$ is a family of extractors. $Ext \otimes \{Ext_z\} : X_1 \times (Y_1 \times Y_2) \to (Z_1 \times Z_2) \times W_2$ is defined as follows: assume $Ext(x_1, y_1) = (z_1, x_2)$ and $Ext_{z_1}(x_2, y_2) = (z_2, w_2)$, then $(Ext \otimes \{Ext_z\})(x_1, y_1, y_2) = (z_1, z_2), w_2$.*

The bipartite graph associated with the product extractor can be constructed as follows: Take the bipartite graph associated with Ext_1, and replace every vertex $z \in Z_1$ with a copy of the extractor Ext_2, by identifying the Ext_1 neighbors of z with elements of X_2, and connecting them to elements in $\{z\} \times Z_2$ according to Ext_2.

The next lemma states that the product of two extractors is also an extractor

Lemma 4.4 (Replacement product lemma)**.** *If $Ext : X_1 \times Y_1 \to Z_1 \times X_2$ is a $(\delta_1, \varepsilon_1)$-extractor and $\{Ext_z\}_{z \in Z_1} : X_2 \times Y_1 \to Z_2 \times W_2$ is a family of $(\delta_2, \varepsilon_2)$-extractors, then $Ext \otimes \{Ext_z\}$ is a (δ, ε)-extractor for $\delta \geq \max\{\delta_1, \delta_2\}$ and $\varepsilon \geq \varepsilon_1 + \varepsilon_2 + \delta_2/\delta$.*

Proof. Let X be a distribution over X_1 with $H_\infty(X) \geq \log(\delta |X_1|)$, and let us show that the distribution defined by $(Ext \otimes \{Ext_z\})(X, Y_1, Y_2)$ over $Z_1 \times Z_2$ is ε-close to uniform.

Consider $z_1 \in Z_1$ whose probability according to $Ext(X, Y_1)$ is at least $(\delta_2/\delta) \cdot (1/|Z_1|)$. Let X_{z_1} be the distribution over X_2 that assigns each $x \in X_2$ its probability according to X conditioned on z_1 being chosen. Since $H_\infty(X) \geq \log(\delta |X_1|)$, the probability of any element according to X_{z_1} is at most $(1/(\delta |X_1|)) \cdot (1/|Y_1|) \cdot (\delta |Z_1| / \delta_2) = |Z_1| / (\delta_2 |X_1| |Y_1|) = 1/(\delta_2 |X_2|)$. Hence, $H_\infty(X_{z_1}) \geq \log(\delta_2 |X_2|)$. By the property of Ext_{z_1}, the distribution defined by $Ext_{z_1}(X_{z_1}, Y_2)$ over Z_2 is ε_2-close to uniform.

The total probability according to $Ext(X, Y_1)$ on $z_1 \in Z_1$ whose probability according to $Ext(X, Y_1)$ is less than $(\delta_2/\delta) \cdot (1/|Z_1|)$ is less than δ_2/δ.

By the property of Ext, the distribution defined by $Ext(X, Y_1)$ on Z_1 is ε_1-close to uniform.

Overall, we can upper bound the distance of the distribution defined by $(Ext \otimes \{Ext_z\})(X, Y_1, Y_2)$ over $Z_1 \times Z_2$ from uniform by $\varepsilon_1 + \delta_2/\delta + \varepsilon_2$. □

Corollary 4.5. *Let $0 < \varepsilon < 1$. Set $\delta(|\mathbb{F}|, k, \varepsilon) = |\mathbb{F}|^{-k/2} k^k (k+1) \varepsilon^{-k}$. Then, for all k', the incidence graph "degree-k curves vs. k'-tuples" is a $(\delta_{k'}/\varepsilon, 2k'\varepsilon)$-extractor for*

$$\delta_{k'} = (|\mathbb{F}| - k' + 1)^{-(k-k'+1)/2} k^k (k+1) \varepsilon^{-k}.$$

Proof. The proof is by induction on k'. For $k' = 1$ the claim follows from Proposition 4.2 that analyzes the sampling properties of the incidence graph "degree-k curves vs. points" and Proposition 2.3 that converts samplers to extractors. Assume that the claim is true for $k'-1$, and let us prove it for k'.

For every $(k'-1)$-tuple of points $S \subseteq \mathbb{F}^m$, consider the "(degree-k curves through S) vs. points" incidence graph, where every curve through S is connected to all the points on it except for those in S. Similarly to Proposition 4.2, this incidence graph is a $(\delta_{k'}, \varepsilon)$-extractor. Moreover, for different S's we get isomorphic incidence graphs.

We can view the incidence graph "degree-k curves vs. k'-tuples" as the product of the incidence graph "degree-k curves vs. $(k'-1)$-tuples" and the family of incidence graphs "(degree-k curves through S) vs. points" for different $(k'-1)$-tuples $S \subseteq \mathbb{F}^m$. By the induction hypothesis, the first is a $(\delta_{k'-1}/\varepsilon, 2(k'-1)\varepsilon)$-extractor. Each graph in the second family is a $(\delta_{k'}, \varepsilon)$-extractor. By Lemma 4.4 and since $\delta_{k'} \geq \delta_{k'-1}$, the product graph is a $(\delta_{k'}/\varepsilon, 2k'\varepsilon)$-extractor. □

Note that the statement of Corollary 4.5 is meaningful for sufficiently large $\mathbb{F}$ with respect to k'. For such we can take $\varepsilon = |\mathbb{F}|^{-\Theta(1)}$ and have a deviation $2k'\varepsilon = |\mathbb{F}|^{-\Theta(1)}$.

Corollary 4.6. *There exists $a \geq 1$, such that for sufficiently large $k \geq a \cdot k'$, for sufficiently large $|\mathbb{F}| \geq a \cdot k$, the incidence graph "degree-k curves vs. k'-tuples" is a $(|\mathbb{F}|^{-\Theta(k)}, |\mathbb{F}|^{-\Theta(1)})$-extractor.*

Another application of the replacement product that we will need later is to analyzing the incidence graph of "pairs of degree-k curves vs. pairs of points". One side of this graph contains all pairs of degree-k curves in $\mathbb{F}^m$, while the other side contains all the pairs of points in $\mathbb{F}^m$. A pair of curves is connected to a pair of points if the first curve contains the first point and the second curve contains the second point.

Corollary 4.7. *The incidence graph "pairs of degree-k curves vs. pairs of points" is a $(|\mathbb{F}|^{-\Theta(k)}, |\mathbb{F}|^{-\Theta(1)})$-extractor.*

Proof. This "pairs of degree-k curves vs. pairs of points" graph is the product of the following incidence graphs:

- "pairs of degree-k curves vs. points", where a pair of curves is connected to a point on the first curve.
- The family of graphs "pairs of degree-k curves where the first curve passes through a point z vs. points" where $z \in \mathbb{F}^m$ and a pair of curves, where the first curve must pass through z, is connected to a point on the second curve.

By Corollary 4.3, both graphs are $(|\mathbb{F}|^{-\Theta(k)}, |\mathbb{F}|^{-\Theta(1)})$-extractors (note that the second curve of the pair in the first graph and the first curve of the pair in the second graph do not damage the extractor property). By Lemma 4.4, the product is an $(|\mathbb{F}|^{-\Theta(k)}, |\mathbb{F}|^{-\Theta(1)})$-extractor as well. □

5 The Base Low Degree Test

In this section we show that a "robust" low degree testing theorem for Surface-vs.-Surface follows from the low degree testing theorem for Line-vs.-Line Test [4].

Lemma 5.1 (Line vs. Line low degree testing theorem [4])**.** *Assume that $\mathbb{F}$ is a large enough field (polynomial size) with respect to d and m. For some $\delta = |\mathbb{F}|^{-\Omega(1)}$, for any prover strategies*

$\mathcal{A}_1$, $\mathcal{A}_2$, there are m-variate polynomials $p_1, \ldots, p_l$, $l \leq O(1/\delta)$, of degree at most d over $\mathbb{F}$, such that the probability that the Line-vs.-Line Test *passes but $\mathcal{A}_1(\ell_1)$ is not one of $p_{1|\ell_1}, \ldots, p_{l|\ell_1}$ (similarly for $\mathcal{A}_2(\ell_2)$), is at most δ.*

Our "robust" low degree testing theorem holds when restricting to surfaces that pass through given k'' points $S \subseteq \mathbb{F}^m$, and even further when restricting to a sub-family of fraction $\delta = |\mathbb{F}|^{-\Theta(k)}$ of the surfaces that pass through S. We use $\mathcal{C}_S$ to denote those surfaces in $\mathcal{C}$ that pass through S. The construction relies on the curve-tuple sampling proved in Section 4.

We focus on the following family $\mathcal{C}$ of 3-dimensional surfaces in $\mathbb{F}^m$: Fix initial conditions

$$\{(x_{i,1}, \ldots, x_{i,q})\}_{i=1}^{M} \subseteq (\mathbb{F}^m)^q.$$

For every $1 \leq i \leq M$, add all the 3-dimensional surfaces of the form

$$s(t_1, t_2, t_3) = c_1(t_1) + t_3 c_2(t_2),$$

where c_1 is a curve of degree at most $q + k$ and passes through $x_{i,1}, \ldots, x_{i,q}$, and c_2 is a curve of degree at most k. The forbidden points $Q : \mathcal{C} \to 2^{\mathbb{F}}$ rule out the initial points embedded in the curves. The number of curves is $M \cdot |\mathbb{F}|^{O((q+k)m)}$. Each surface is the union of $|\mathbb{F}|^2$ lines $x + ty$ where $x \in \mathbb{F}^m$ is on the curve c_1, and $y \in \mathbb{F}^m$ is on the curve c_2.

Proposition 5.2 (Base test)**.** *Assume that $\mathbb{F}$ is a large enough field (polynomial size) with respect to m, d, k and $|Q|$. Assume that k is large enough (linear size) in k''. There are $\delta = |\mathbb{F}|^{-\Omega(k)}$ and $\varepsilon = |\mathbb{F}|^{-\Omega(1)}$ that satisfy:*

For any $\mathcal{C}' \subseteq \mathcal{C}$ and $S' \subseteq \mathbb{F}^m$, $|S'| \leq k''$, such that $\left|\mathcal{C}'_{S'}\right| \geq \delta \left|\mathcal{C}_{S'}\right|$,

$$ListErr^{\mathcal{C}'_{S'}}_{|\mathbb{F}|^{O(1)}, d(q+k)}(\textsc{Surface-vs.-Surface Test}(\mathcal{C}'_{S'}, Q, \mathbb{F}^m)) \leq |\mathbb{F}|^{-\Omega(1)}.$$

Proof. The proposition is proved by using a successful strategy for Surface-vs.-Surface to succeed in Line-vs.-Line Test, applying Lemma 5.1 to deduce agreement of the assignments to lines with a low degree polynomial, and arguing that this agreement extends to surfaces.

Fix $\mathcal{C}'$ and S' as in the premise. Fix a strategy for Surface-vs.-Surface Test$(\mathcal{C}'_{S'}, Q, \mathbb{F}^m)$. We use it to devise a strategy for Line-vs.-Line Test as follows. Given a line $\ell = \{x + ty \mid t \in \mathbb{F}\}$ where x is uniformly distributed in $\mathbb{F}^m$ and y is uniformly and independently distributed in $\mathbb{F}^m - \{0\}$, we consider a uniform surface in $\mathcal{C}'_{S'}$ of the form $s = \{c_1(t_1) + t \cdot c_2(t_2) \mid t_1, t_2, t \in \mathbb{F}\}$ such that $x \in c_1(\mathbb{F})$ and $y \in c_2(\mathbb{F})$, $\ell \subseteq s^{-Q}$. The assignment $\mathcal{A}(\ell)$ is the restriction to ℓ of the assignment to the surface.

By Corollary 4.7, and using that $|\mathbb{F}|$ is sufficiently larger than $|Q|$, the distribution of the surfaces picked by the above process during a run of Line-vs.-Line Test is $|\mathbb{F}|^{-\Theta(1)}$-close to the distribution of surfaces in Surface-vs.-Surface Test$(\mathcal{C}'_{S'}, Q, \mathbb{F}^m)$. Hence, except for probability $|\mathbb{F}|^{-\Theta(1)}$, Line-vs.-Line Test succeeds whenever Surface-vs.-Surface Test does.

Applying Lemma 5.1, for $\delta = |\mathbb{F}|^{-\Theta(1)}$ there are m-variate polynomials $p_1, \ldots, p_l$, $l \leq O(1/\delta)$, of degree at most $d(q + k)$ over $\mathbb{F}$, such that the probability that Line-vs.-Line Test passes, yet it picks a line ℓ_1 such that $A(\ell_1)$ is not one of $p_{1|\ell_1}, \ldots, p_{l|\ell_1}$, is at most δ.

Assume that Surface-vs.-Surface Test passes with probability sufficiently large $|\mathbb{F}|^{-\Theta(1)}$ (otherwise we are done). Since $|\mathbb{F}|$ is sufficiently large with respect to d, q and k, for a uniform surface in $\mathcal{C}'_{S'}$

$$s = \{c_1(t_1) + t \cdot c_2(t_2) \mid t_1, t_2, t \in \mathbb{F}\},$$

except for probability $|\mathbb{F}|^{-\Theta(1)}$, for more than $d(q+k)l/|\mathbb{F}|$ fraction of the choices of $t_1, t_2 \in \mathbb{F}$, $c_2(t_2) \neq \vec{0}$, on the restriction of s to the line $\{c_1(t_1) + t \cdot c_2(t_2) \mid t \in \mathbb{F}\}$ the assignment $\mathcal{A}(s)$ agrees with one of $p_{1|\ell_1}, \ldots, p_{l|\ell_1}$. For such s's it holds that $\mathcal{A}(s) \in \{p_{1|s}, \ldots, p_{l|s}\}$. The lemma follows. □

6 Setup For Direct Product Theorem

In this section we formally define our direct product theorem, and outline its analysis.

The theorem assumes that a test that compares two surfaces on a point has error $|\mathbb{F}|^{-\Omega(1)}$ even when restricting to a sub-family of surfaces, and shows that a test that compares two surfaces on k' points has error $|\mathbb{F}|^{-\Omega(k')}$.

Theorem 6.1 (Direct product for low degree testing)**.** *Assume that $\mathbb{F}$ is large enough (polynomial size) in d', m, k', $|Q|$, and that k is large enough (linear size) in k'. Then, the assumption about the base test implies the conclusion about the repeated test:*

Base Test: *Assume that there exists a constant $0 < \beta' < 1$ such that for every set $S' \subseteq \mathbb{F}^m$, $|S'| \leq \beta' k'$, for every $\mathcal{C}' \subseteq \mathcal{C}$, $|\mathcal{C}'| \geq |\mathbb{F}|^{-\beta' k'} |\mathcal{C}|$, and $Q' : \mathcal{C}' \to 2^{\mathbb{F}}$, $|Q'| \leq |Q| + \beta' k'$:*

$$ListErr^{\mathcal{C}'_{S'}}_{|\mathbb{F}|^{-\Omega(1)},d'}(\textsc{Surface-vs.-Surface Test}(\mathcal{C}'_{S'}, Q', \mathbb{F}^m)) \leq |\mathbb{F}|^{-\Theta(1)}.$$

Product Test: *Then, there exists $\delta = |\mathbb{F}|^{-\Omega(k')}$, such that*

$$ListErr^{\mathcal{C}}_{|\mathbb{F}|^{O(k')},d'}(\textsc{Surfaces Test}(\mathcal{C}, Q, \mathcal{I})) \leq \delta.$$

For simplicity, we prove Theorem 6.1 for curves rather than surfaces. Our arguments readily extend to surfaces.

The heart of the proof of Theorem 6.1 is an analysis of the two query Curve-vs.-Curve Test. This analysis only gives a guarantee about a tiny portion of all curves in $\mathcal{C}$. We then use the extra queries in Curves Test and the error correction properties of polynomials to give a guarantee about a sizable portion of $\mathcal{C}$. The analysis consists of the following three parts:

1. Analysis of Curve-vs.-Curve Test (Sections 7, 8 and 9): This is the heart of the analysis. Use the base test to deduce that success $|\mathbb{F}|^{-\beta k'}$ in Curve-vs.-Curve Test gives rise to a set S' of $\approx \beta k'$ points in $\mathbb{F}^m$ and a low degree polynomial that agrees with $\approx |\mathbb{F}|^{-\beta k'}$ fraction of the curves in $\mathcal{C}_{S'}$ (recall that $\mathcal{C}_{S'}$ are the curves in $\mathcal{C}$ that contain S'). To reduce to the base test we use the sampling properties of the "degree-k curves vs. k'-tuples of points" incidence graph (this is the Wide Agreement Lemma, which is in Section 8).

2. Weak analysis of Curves Test (Section 10): Use the analysis of Curve-vs.-Curve Test from the previous item to deduce that success $|\mathbb{F}|^{-\beta k'}$ in Curves Test gives rise to a low degree polynomial that agrees with $\approx |\mathbb{F}|^{-\beta k'}$ fraction of the k'-tuples of points in $\mathbb{F}^m$. Here we use the extra queries of the verifier in Curves Test to go from a structural conclusion about the assignment to $\mathcal{C}_{S'}$, a tiny portion of all $\mathcal{C}$, to a structural conclusion about the assignment to a sizable portion of all k'-tuples of points in $\mathbb{F}^m$.

3. Strong analysis of Curves Test (Section 11): Use the weak analysis of Curves Test to deduce that success $|\mathbb{F}|^{-\beta k'}$ in Curves Test gives rise to a low degree polynomial that agrees with $\approx |\mathbb{F}|^{-\beta k'}$ fraction of the curves in $\mathcal{C}$. Here we rely on the list decoding argument in Section 3.3.

7 Identifying a Successful Sub-Test

In this section we show that success probability $\gamma \gg |\mathbb{F}|^{-k'}$ of Curve-vs.-Curve Test implies a much higher success probability $\gg 1/|\mathbb{F}|$ of Curve-vs.-Curve-on-Point Test over a small subset of the curves that have some points in common and agree on them. We will later use this lemma where the initial family of curves does not necessarily contain all degree-k curves. We therefore use $\mathcal{C}^*$ to denote the initial set of curves, and $\mathcal{C}$ to denote the family of all degree-k curves in $\mathbb{F}^m$.

Lemma 7.1 (Sub-Test Lemma)**.** *Suppose that the probability that* Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ *passes is at least* $4|\mathbb{F}|^{-\beta k'}$*. For any* $k'' \geq 1$ *if* $0 < \beta' < 1$ *satisfies* $\beta' > 8\beta k'/k''$*, then there exist*

- $S' \subseteq \mathbb{F}^m$, $|S'| = k''$; $\left|\mathcal{C}^*_{S'}\right| \geq |\mathbb{F}|^{-\beta k'} \cdot (|\mathcal{C}^*| / |\mathcal{C}|) \cdot |\mathcal{C}_{S'}|$;
- $\mathcal{C}' \subseteq \mathcal{C}_{S'}$, $|\mathcal{C}'| \geq |\mathbb{F}|^{-\beta k'} \cdot \left|\mathcal{C}^*_{S'}\right|$;
- $Q' : \mathcal{C}' \to 2^{\mathbb{F}}$, $|Q'| \leq |Q| + k''$;

such that Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ *accepts with probability at least* $(1/2)|\mathbb{F}|^{-\beta'}$.

Proof. Pick uniformly $c \in \mathcal{C}^*$ and $S' \subseteq c^{-Q}$, $|S'| = k''$, such that conditioned on Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ picking c and $S \subseteq c^{-Q}$, $S \supseteq S'$, the test passes with probability at least $|\mathbb{F}|^{-\beta k'}$. Let $a_1, \ldots, a_{k''} \in \mathbb{F}$ be the assignments of $\mathcal{A}(c)$ to the points in S'. Let $\mathcal{C}' \subseteq \mathcal{C}^*_{S'}$ contain all those curves $c' \in \mathcal{C}^*_{S'}$ that assign $a_1, \ldots, a_{k''}$ to the points in S'. We know that $|\mathcal{C}'| \geq |\mathbb{F}|^{-\beta k'} \cdot \left|\mathcal{C}^*_{S'}\right|$. The probability that $\left|\mathcal{C}^*_{S'}\right| < |\mathbb{F}|^{-\beta k'} \cdot (|\mathcal{C}^*| / |\mathcal{C}|) \cdot |\mathcal{C}_{S'}|$ is at most $|\mathbb{F}|^{-\beta k'}$.

We'll show that conditioned on Curve-vs.-Curve Test picking $c_1, c_2 \in \mathcal{C}'$ and $S' \subseteq S \subseteq c_1^{-Q}, c_2^{-Q}$, it holds that $\mathcal{A}(c_1)$, $\mathcal{A}(c_2)$ agree on more than $|\mathbb{F}|^{-\beta'}$ fraction of the $x \in S$ except with probability at most $|\mathbb{F}|^{-\beta' k''/2}$ over the choice of c, S', as well as over the randomness in Curve-vs.-Curve Test. Let us call this property "high probability agreement". If $\mathcal{A}(c_1)$, $\mathcal{A}(c_2)$ agree on at most $|\mathbb{F}|^{-\beta'}$ fraction of the $x \in S$, then – since S' is uniform inside S – the probability that S' falls inside the agreement points is at most $|\mathbb{F}|^{-\beta' k''}$. By Bayes' law, and since the probability of agreement on S' is at least $|\mathbb{F}|^{-\beta k'}$, conditioning on agreement on S' the probability that $\mathcal{A}(c_1)$, $\mathcal{A}(c_2)$ agree on at most $|\mathbb{F}|^{-\beta'}$ fraction of the $x \in S$ is at most $|\mathbb{F}|^{-(\beta' k'' - \beta k')} \leq |\mathbb{F}|^{-\beta' k''/2}$.

Fix c and S' such that $\left|\mathcal{C}^*_{S'}\right| \geq |\mathbb{F}|^{-\beta k'} \cdot (|\mathcal{C}^*| / |\mathcal{C}|) \cdot |\mathcal{C}_{S'}|$ and high probability agreement holds where the probability is only taken over the randomness in the test. Let Q' inherit the forbidden points of Q in addition to the points of S'.

The distribution $\mathcal{D}_1$ of curves and tuples in Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ is different from their distribution $\mathcal{D}_2$ in Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ conditioned on the curves being in $\mathcal{C}'$ and tuples falling to $\mathcal{I}_{S'}$. In particular, $\mathcal{D}_1$ may place much of the probability on very few tuples that most of their curves fall in $\mathcal{C}'$ (since the probability of a tuple is proportional to *square* the fraction of its curves that fall into $\mathcal{C}'$), while in $\mathcal{D}_2$, assuming that $\mathcal{C}'$ is large enough, the distribution of tuples is close to uniform over all tuples in $\mathcal{I}_{S'}$. However, for any event E, the probability of E according to $\mathcal{D}_1$ is at most $|\mathbb{F}|^{\beta k'}$ times its probability according to $\mathcal{D}_2$. In particular, this is true for the event that $\mathcal{A}(c_1)$, $\mathcal{A}(c_2)$ agree on more than $|\mathbb{F}|^{-\beta'}$ fraction of the $x \in S$. Therefore, the probability that Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ accepts is at least $(1/2)|\mathbb{F}|^{-\beta'}$. □

8 From Large Intersection To One Point Intersection

In this section we prove the Wide Agreement Lemma discussed in the introduction. We show that if the Curve-vs.-Curve-on-Point Test passes with good probability when the intersection between curves contains $(k' - k'')$ points, then the Curve-vs.-Curve Test passes with comparably good probability when the intersection between curves contains just one point. The proof relies on the sampling property of the "degree-k curves vs. k'-tuples" incidence graph.

Lemma 8.1 (Wide Agreement Lemma)**.** *Assume the setup of Lemma 7.1, and in particular that for S', $\mathcal{C}'$ and Q' as there, the probability that* Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ *accepts is at least $(1/2)\,|\mathbb{F}|^{-\beta'}$. Further, assume that β and β' are such that for every $(k''+1)$-tuple of points $S'' \subseteq \mathbb{F}^m$, the incidence graph $\mathcal{G}(\mathcal{C}_{S''}, \mathcal{I}_{S''})$ is (δ, ε)-sampling for*

$$\varepsilon \leq (1/12) \cdot |\mathbb{F}|^{-\beta'}.$$

$$\delta \leq \zeta^2\, |\mathbb{F}|^{-2\beta k'}\, (|\mathcal{C}^*| \,/\, |\mathcal{C}|)\varepsilon^2,$$

Then, Curve-vs.-Curve Test$(\mathcal{C}', Q', \mathbb{F}^m)$ *passes with probability at least $(1/2)\,|\mathbb{F}|^{-\beta'} - 3\varepsilon$.*

Proof. For a point $x \in \mathbb{F}^m$ let $\mathcal{C}_{S'\cup\{x\}}$ be the curves $c \in \mathcal{C}_{S'}$ such that $x \in c^{-Q'}$, and let $\mathcal{I}_{S'\cup\{x\}}$ be the tuples in $\mathcal{I}_{S'}$ that pass through x. Since S' is fixed throughout the proof, we use $\mathcal{C}_x$ to denote $\mathcal{C}_{S'\cup\{x\}}$ and we use $\mathcal{I}_x$ to denote $\mathcal{I}_{S'\cup\{x\}}$.

Let $p(x)$ be the probability that a uniform $c \in \mathcal{C}'$ contains x as $x \in c^{-Q'}$. For a possible assignment $a \in \mathbb{F}$ to x, let $\mathcal{C}_{x,a}$ be the family of curves in $\mathcal{C}_x$ with $\mathcal{A}(c)(x) = a$. Per $S \in \mathcal{I}_{S'}$, let $\mu_{x,a}(S)$ be the fraction of curves in $\mathcal{C}_{x,a}$ among the curves in $\mathcal{C}' \cap \mathcal{C}_x$ that contain S. Since every curve $c \in \mathcal{C}' \cap \mathcal{C}_x$ contains the same number of k'-tuples in $\mathcal{I}_x$,

$$\mathop{\mathbf{E}}_{S\in\mathcal{I}_x}[\mu_{x,a}(S)] = \frac{|\mathcal{C}' \cap \mathcal{C}_{x,a}|}{|\mathcal{C}' \cap \mathcal{C}_x|}.$$

The probability that Curve-vs.-Curve Test$(\mathcal{C}', Q', \mathbb{F}^m)$ passes is given by

$$\sum_{x\in\mathbb{F}^m} p(x) \cdot \sum_{a\in\mathbb{F}} \left(\frac{|\mathcal{C}' \cap \mathcal{C}_{x,a}|}{|\mathcal{C}' \cap \mathcal{C}_x|}\right)^2. \tag{2}$$

The probability that Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I})$ passes is given by

$$\sum_{x\in\mathbb{F}^m} p(x) \cdot \sum_{a\in\mathbb{F}} \sum_{c\in\mathcal{C}'\cap\mathcal{C}_{x,a}} \frac{1}{|\mathcal{C}' \cap \mathcal{C}_x|} \cdot \mathop{\mathbf{E}}_{S\in\mathcal{I}_x : S\subseteq c}[\mu_{x,a}(S)]. \tag{3}$$

By the sampling property of "degree-k curves vs. $(k'+1)$-tuples", all curves $c \in \mathcal{C}_x$, except for at most $\delta\,|\mathcal{C}_x|$ curves which we denote B_x, have

$$\mathop{\mathbf{E}}_{S\in\mathcal{I}_x : S\subseteq c}[\mu_{x,a}(S)] = \frac{|\mathcal{C}' \cap \mathcal{C}_{x,a}|}{|\mathcal{C}' \cap \mathcal{C}_x|} \pm \varepsilon.$$

Let $G \subseteq \mathbb{F}^m$ be the points $x \in \mathbb{F}^m$ for which $|\mathcal{C}' \cap \mathcal{C}_x| > (\delta/\varepsilon)\,|\mathcal{C}_x|$. Since $|\mathcal{C}'| \geq (\delta/\varepsilon^2)\,|\mathcal{C}_{S'}|$, for $x \notin G$ it holds:

$$p(x) = \frac{|\mathcal{C}' \cap \mathcal{C}_x|}{|\mathcal{C}'|} \leq \frac{\delta}{\varepsilon}\frac{|\mathcal{C}_x|}{|\mathcal{C}'|} \leq \frac{\delta}{\varepsilon}\frac{|\mathcal{C}_x|}{(\delta/\varepsilon^2)\,|\mathcal{C}_{S'}|} = \varepsilon\frac{|\mathcal{C}_x|}{|\mathcal{C}_{S'}|}.$$

Hence, the contribution to (3) from points $x \notin G$ is at most

$$\sum_{x \in \mathbb{F}^m} \varepsilon \frac{|\mathcal{C}_x|}{|\mathcal{C}_{S'}|} \leq \varepsilon.$$

The contribution to (3) from points $x \in G$ and curves $c \in B_x$ is at most

$$\sum_{x \in G} p(x) \cdot \frac{|\mathcal{C}' \cap \mathcal{C}_x \cap B_x|}{|\mathcal{C}' \cap \mathcal{C}_x|} \leq \sum_{x \in G} p(x) \cdot \frac{\delta \, |\mathcal{C}_x|}{(\delta/\varepsilon) \, |\mathcal{C}_x|} \leq \varepsilon.$$

Hence, we can upper bound the probability in (3) by:

$$\begin{aligned}
(3) \quad &\leq \quad 2\varepsilon + \sum_{x \in G} p(x) \cdot \sum_{a \in \mathbb{F}} \sum_{c \in \mathcal{C}' \cap \mathcal{C}_{x,a} - B_x} \frac{1}{|\mathcal{C}' \cap \mathcal{C}_x|} \cdot \mathop{\mathbf{E}}_{S \in \mathcal{I}_x : S \subseteq c} [\mu_{x,a}(S)] \\
&\leq \quad 2\varepsilon + \sum_{x \in G} p(x) \cdot \sum_{a \in \mathbb{F}} \sum_{c \in \mathcal{C}' \cap \mathcal{C}_{x,a}} \frac{1}{|\mathcal{C}' \cap \mathcal{C}_x|} \cdot \left(\frac{|\mathcal{C}' \cap \mathcal{C}_{x,a}|}{|\mathcal{C}' \cap \mathcal{C}_x|} + \varepsilon \right) \\
&\leq \quad 3\varepsilon + \sum_{x \in \mathbb{F}^m} p(x) \cdot \sum_{a \in \mathbb{F}} \left(\frac{|\mathcal{C}' \cap \mathcal{C}_{x,a}|}{|\mathcal{C}' \cap \mathcal{C}_x|} \right)^2
\end{aligned}$$

The lemma follows from (2). □

9 Curve vs. Curve Analysis

In this section we apply the machinery we developed to this point, as well as the guarantee about the base test, to start from a noticeable success probability of Curve-vs.-Curve Test and get a small set of points and a low degree polynomial that agrees with a noticeable fraction of the curves through the points.

Lemma 9.1 (Analysis of Curve-vs.-Curve Test)**.** *Assume that $|\mathbb{F}|$ is a sufficiently large polynomial of d, m, k. Let $0 < \beta' < \beta < 1$.*

Base Test: *Let $\mathcal{C}' \subseteq \mathcal{C}$ be such that for every $\beta' k'$-tuple of points $S' \subseteq \mathbb{F}^m$ where $\left|\mathcal{C}'_{S'}\right| \geq \delta \, |\mathcal{C}_{S'}|$, for every $Q' : \mathcal{C}' \to 2^{\mathbb{F}}$, $|Q'| \leq q$:*

$$AgrErr^{\mathcal{C}'}_{\gamma \to \gamma', d \to d'}(\text{Curve-vs.-Curve Test}(\mathcal{C}', Q', \mathbb{F}^m)) \leq \gamma_0.$$

Assumptions:

- $$\gamma_0 \leq (1/4) \cdot |\mathbb{F}|^{-\beta'},$$
- *For every $S'' \subseteq \mathbb{F}^m$, $|S''| \leq \beta' k' + 1$, the incidence graph $\mathcal{G}(\mathcal{C}_{S''}, \mathcal{I}_{S''})$ is (δ, ε)-sampling for δ and ε as in Lemma 8.1.*

Repeated Test: *If Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ passes with probability at least $4 \, |\mathbb{F}|^{-\beta k'}$, then there exists a set $S' \subseteq \mathbb{F}^m$, $|S'| \leq \beta' k'$ and an m-variate polynomial of degree at most d' over $\mathbb{F}$ that agrees with at least $\gamma'((1/4) \cdot |\mathbb{F}|^{-\beta'}) \cdot |\mathbb{F}|^{-2\beta k'} \left(|\mathcal{C}^*| \,/\, |\mathcal{C}|\right) |\mathcal{C}_{S'}|$ fraction of the curves in $\mathcal{C}^*_{S'}$.*

Proof. Assume that Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ passes with probability at least $4\left|\mathbb{F}\right|^{-\beta k'}$. Let $0 < \beta' < \beta$ and $k'' < \beta' k'$ be as in Lemma 7.1. By Lemma 7.1, there exist $S' \in (\mathbb{F}^m)^{k''}$; $\mathcal{C}' \subseteq \mathcal{C}_{S'}$; $Q' : \mathcal{C}' \to 2^{\mathbb{F}}$; such that $|\mathcal{C}'| = \left|\mathcal{C}'_{S'}\right| \geq \left|\mathbb{F}\right|^{-2\beta k'} \left(|\mathcal{C}^*| / |\mathcal{C}|\right) |\mathcal{C}_{S'}|$, and the probability that Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ passes is at least $(1/2)\left|\mathbb{F}\right|^{-\beta'}$. By Lemma 8.1, Curve-vs.-Curve Test$(\mathcal{C}', Q', \mathbb{F}^m)$ passes with probability at least $(1/4) \cdot \left|\mathbb{F}\right|^{-\beta'}$. By our assumption on the base test, there is an m-variate polynomial p of degree at most d' over $\mathbb{F}$ that agrees with a set of curves $\mathcal{C}_p$ of fraction $\gamma' = \gamma'((1/4) \cdot \left|\mathbb{F}\right|^{-\beta'})$ in $\mathcal{C}'$. We have

$$|\mathcal{C}_p| \geq \gamma' \left|\mathcal{C}'\right| \geq \gamma' \left|\mathbb{F}\right|^{-2\beta k'} \left(|\mathcal{C}^*| / |\mathcal{C}|\right) |\mathcal{C}_{S'}| .$$

□

10 Curves Test Analysis

In this section we use the analysis of Curve-vs.-Curve Test in Section 9 to derive an analogous conclusion for Curves Test. Thanks to the third query in Curves Test, this time we get a conclusion about the agreement of a low degree polynomial with a large portion of all k'-tuples, not just those that contain a small set S' of points. In the next sections we will extend this to argue about agreement with a large portion of all curves.

Lemma 10.1 (Analysis of Curves Test)**.** *Using the notation of Lemma 9.1, and under its assumptions about the parameters and the base test:*

Repeated Test:

$$AgrErr^{\mathcal{I}}_{\gamma \to \left|\mathbb{F}\right|^{-\Omega(k')}, d \to d'}(\text{Curves Test}(\mathcal{C}, Q, \mathcal{I})) \leq 2\left|\mathbb{F}\right|^{-\beta k'} .$$

Proof. Assume that Curves Test$(\mathcal{C}, Q, \mathcal{I})$ passes with probability at least $2\left|\mathbb{F}\right|^{-\beta k'}$. Let $\mathcal{C}^*$ be the family of curves $c \in \mathcal{C}$, such that with probability at least $\left|\mathbb{F}\right|^{-\beta k'}$ over the choice of $S \subseteq c^{-Q}$, it holds that $\mathcal{A}(c)$ and $\mathcal{A}(S)$ agree. We have $|\mathcal{C}^*| \geq \left|\mathbb{F}\right|^{-\beta k'} |\mathcal{C}|$, and Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ passes with probability at least $\left|\mathbb{F}\right|^{-\beta k'}$.

By Lemma 9.1, there exists an m-variate polynomial p of degree at most d' over $\mathbb{F}$, such that with probability at least $\gamma'((1/4) \cdot \left|\mathbb{F}\right|^{-\beta'}) \cdot \left|\mathbb{F}\right|^{-3\beta k'} |\mathcal{C}_{S'}|$ over $c \in \mathcal{C}^*_{S'}$, it holds that $\mathcal{A}(c) \equiv p_{|c}$. The lemma follows since every k'-tuple that does not intersect S' is contained in the same number of curves in $\mathcal{C}_{S'}$, which means that $\mathcal{A}(c_2)$ agrees with at least $\left|\mathbb{F}\right|^{-\beta k'}$ fraction of the S_2's. □

11 Concluding The Analysis

In this section we get a list of polynomials that explains almost all of the success of Curves Test.

Lemma 11.1 (decoding → list decoding)**.** *Under the assumptions of Lemma 10.1, there exists* $\delta_0 = \left|\mathbb{F}\right|^{-\Omega(k')}$*, such that*

$$ListErr^{\mathcal{C}}_{2/\delta_0, dk}(\text{Curves Test}(\mathcal{C}, Q, \mathcal{I})) \leq \delta_0.$$

Proof. By Lemma 10.1 and the list decoding transformation of Lemma 3.2. □

From the list decoding that explains the success of assignments to tuples we can get a list decoding that explains that success of assignments to curves:

Lemma 11.2. *Using the assumptions and notation of Lemma 11.1, and assuming that*[6] $(\delta_0^{3/2}/2)\cdot\binom{|\mathbb{F}|-|Q|}{k'} > \binom{dk}{k'}$,

$$ListErr^{\mathcal{C}}_{2/\delta_0,dk}(\textsc{Curves Test}(\mathcal{C},Q,\mathcal{I})) \leq \delta_0.$$

Proof. For c_1, S_1, c_2, S_2, c_3 picked in Curves Test, we set:

- AGR: $\mathcal{A}(c_3)_{|S_2} \equiv \mathcal{A}(S_2)$;
- $TEXP_i$: $\mathcal{A}(S_2) \equiv p_{i|S_2}$;
- $TEXP$: $\bigvee_{i=1}^{l} TEXP_i$;
- $CEXP_i$: $\mathcal{A}(c_3) \equiv p_{i|c_3}$;
- $CEXP$: $\bigvee_{i=1}^{l} CEXP_i$;

By Lemma 11.1,

$$\Pr\left[AGR \wedge \neg TEXP\right] \leq \delta.$$

Hence,

$$\Pr_{c_3}\left[\Pr_{S_2}\left[AGR \wedge \neg TEXP\right] \geq \sqrt{\delta}\right] \leq \sqrt{\delta}.$$

Consider a curve c_3 such that $\Pr_{S_2}\left[AGR \wedge TEXP\right] \geq \sqrt{\delta}$. Then, there exists $1 \leq i \leq l$, such that $\Pr_{S_2}\left[TEXP_i\right] \geq \sqrt{\delta}/l$. Since the premise of the lemma guarantees that $\sqrt{\delta}/l$ fraction of the tuples on a curve must span more than dk points, we have $CEXP_i$. Hence,

$$\begin{aligned}
\Pr\left[AGR \wedge \neg CEXP\right] &\leq \Pr\left[AGR \wedge \Pr_{S_2}\left[AGR \wedge TEXP\right] < \sqrt{\delta}\right] \\
&\leq \Pr\left[AGR \wedge (\Pr_{S_2}\left[AGR\right] < 2\sqrt{\delta})\right] + \Pr\left[\Pr_{S_2}\left[AGR \wedge \neg TEXP\right] \geq \sqrt{\delta}\right] \\
&\leq 2\sqrt{\delta} + \sqrt{\delta}.
\end{aligned}$$

□

Putting together Lemma 11.2, appropriately adjusted to surfaces, and Proposition 5.2 establishing the base test, Theorems 3.1 and 1.1 follow. The proof of Theorem 1.2, establishing a property testing algorithm based on Surfaces Test, follows as well. Given a function $f : \mathbb{F}^m \to \mathbb{F}$ to be tested for low degree, the algorithm invokes Surfaces Test, queries the evaluations of f on the points contained in surfaces picked by Surfaces Test, and checks whether they correspond to assignments that would have passed the test. If f is a polynomial of degree at most d, then the algorithm always accepts. Suppose that the test accepts with sufficiently large probability $|\mathbb{F}|^{-\Theta(k)}$. Then, as we saw, the restriction of f to $|\mathbb{F}|^{-\Theta(k)}$ fraction of the surfaces is consistent with some low degree polynomial. By Corollary 4.3 adapted to surfaces (and provided that the fraction $|\mathbb{F}|^{-\Theta(k)}$ is sufficiently large), f must agree with the polynomial on $1 - |\mathbb{F}|^{-\Omega(1)}$ fraction of the points.

[6] When considering v-dimensional surfaces rather than curves, the condition becomes $(\delta_0^{3/2}/2)\cdot\binom{|\mathbb{F}|-|Q|}{k'} > \binom{dk|\mathbb{F}|^{v-1}}{k'}$.

12 From Derandomized Low Degree Test to Sliding Scale Conjecture

12.1 Preliminaries on Probabilistically Checkable Proofs

A PCP verifier is an NP verifier that has polynomially many tests, each depending on a bounded number of queries to the proof. A random test (even though it involves only a bounded number of queries!) predicts correctly the outcome of the verification with good probability.

Definition 12.1 (PCP verifier)**.** *For c, s, r, q, Σ that are functions of n, the class $PCP_{c,s}[r, q]_\Sigma$ contains all languages L that have verifiers that on input x of size n use r random bits to make q queries to a proof over alphabet Σ, and satisfy:*

- Completeness: *For every $x \in L$, there exists a proof π such that the verifier accepts with probability at least c.*
- Soundness: *For every $x \notin L$, for any purported proof π, the verifier accepts with probability at most s.*

Σ is called the *alphabet* of the proof. It Σ is omitted, the understanding is that $\Sigma = \{0, 1\}$. Often we only specify the size of Σ, in which case it is understood that $\Sigma = \{1, \ldots, |\Sigma|\}$. The *size* of the PCP (equivalently, the *proof length*) can be bounded by $2^r q$. If on inputs x of size n we have $2^r q = n^{1+o(1)} poly(1/\varepsilon)$, then we say that the PCP is of *almost linear size*. If $c = 1$ we say that the verifier has *perfect completeness*. In this work we will only consider verifiers with perfect completeness. The fraction s is called the *soundness error* of the verifier, or simply the *error*. We have the following lower bounds on the error:

Proposition 12.1. *If $s < 2^{-r}$ or $s < |\Sigma|^{-q}$, then $PCP_{c,s}[r, q]_{|\Sigma|} \subseteq P$.*

In other words, for $r = O(\log n)$ the error can be at best polynomially small in n, and to achieve error s with a constant number of queries, one has to take the alphabet to be at least $(1/s)^{\Omega(1)}$.

Given a PCP verifier, one can generate a new PCP verifier with lower error and more queries by sequentially repeating the test of the original verifier. The new verifier can be implemented in a randomness-efficient manner, yielding the following:

Proposition 12.2 (Sequential repetition)**.** *For every $\varepsilon = \varepsilon(n) > 0$, there is $k = \Theta(\log_{1/s}(1/\varepsilon))$, such that*

$$PCP_{1,s}[r, q]_{|\Sigma|} \subseteq PCP_{1,\varepsilon}[O(r + k), qk]_{|\Sigma|}.$$

We say that a PCP verifier is a *projection PCP verifier* (or that the PCP is a *projection game*) if the verifier makes $q = 2$ queries, and given the answer to the first query, there is at most one accepting answer to the second query.

A different perspective on PCP verifiers is given by the notion of *multi-prover interactive proofs* or *multi-prover games*:

Definition 12.2 (MIP)**.** *We say that a language L has an MIP protocol with parameters c, s, r, q, Σ, if there is a protocol in which a verifier interacts with q non-interacting provers, uses r random bits to decide on queries to the q provers; the provers respond with replies taken from an alphabet Σ.*

- Completeness: *For every $x \in L$, there exists a strategy to the provers such that the verifier accepts with probability at least c.*
- Soundness: *For every $x \notin L$, for any strategy to the provers, the verifier accepts with probability at most s.*

One can view any MIP protocol as a PCP verifier, and vice versa. The proof for the PCP verifier consists of writing down, for each of the MIP's protocol q provers, its replies on all the possible questions of the verifier.

The PCP Theorem states that probabilistic checking of proofs can always be done with constant number of queries:

Theorem 12.1 (PCP Theorem [6, 5, 3, 2])**.** $NP \subseteq PCP_{1,\frac{1}{2}}[O(\log n), O(1)]$.

Various works amplify the soundness error of the basic PCP theorem. We will use a PCP theorem with low error:

Theorem 12.2 (Low error PCP Theorem [29, 4, 12])**.**

$$NP \subseteq PCP_{1,2/|\Sigma|}\left[O(\log n), O(1)\right]_{|\Sigma|}, \textit{where } \log|\Sigma| = \sqrt{\log n \log\log n}.$$

12.2 From Derandomized Low Degree Test to Sliding Scale Conjecture

In this section we show how a derandomized low degree test as in Conjecture 3.1 implies the Sliding Scale Conjecture, hence proving Theorem 1.3. The idea of the proof is to use the low degree test for simulating sequential repetition. This idea has been used in many works before, however, there are a few differences between the current proof and previous works: (1) We start with a low error PCP by Dinur et al [12], and our choice of parameters is unusual; (2) We formulate and use a new abstraction of the composition theorem of Arora-Safra [3].

Our construction is as follow. We start with an instantiation of the low error PCP from Theorem 12.2:

$$NP \subseteq PCP_{1,2/|\Sigma|}\left[O(\log n), O(1)\right]_{|\Sigma|}, \text{where } \log|\Sigma| = \sqrt{\log n \log\log n}.$$

By sequential repetition of this PCP $O(\sqrt{\log n/\log\log n})$ times (see Proposition 12.2) we get:

$$NP \subseteq PCP_{1,1/n}\left[O(\log n), O(\sqrt{\log n/\log\log n})\right]_{|\Sigma|}.$$

We wish to decrease the number of queries to a constant without hurting the soundness error or the randomness too much. Recall that to allow that we have to increase the alphabet appropriately (see Proposition 12.1). Ultimately, we want to prove a PCP theorem with polynomially small error and polynomial alphabet size:

$$NP \subseteq PCP_{1,1/n^{\Omega(1)}}\left[O(\log n), O(1)\right]_{n^{O(1)}}. \tag{4}$$

From this, one can get "sliding-scale", i.e., error ε with alphabet size $poly(1/\varepsilon)$ by composition with a Hadamard/quadratic functions-based construction.

In the next section we describe the algebraic framework for converting a PCP verifier with many queries to a PCP verifier with a constant number of queries based on the local testing and decoding properties of low degree polynomials. This framework is invoked twice, with two

different settings of parameters. In the first application (see Section 12.3), we get a construction with sub-exponential alphabet (v is a constant):

$$NP \subseteq PCP_{1,1/n^{\Omega(1)}}\left[O(\log n), O(1)\right]_{2^{2^{\Theta((\log n)^{1/2v})}}}. \tag{5}$$

In the second application (see Section 12.4), we get a construction[7] with poly-logarithmic randomness, poly-logarithmically small soundness error, and quasi-polynomial alphabet:

$$NP \subseteq PCP_{1,2^{-\Omega(\log^{2v} n)}}\left[O((\log n)^{4v-1}), O(1)\right]_{2^{\Theta(\log^{4v-1} n)}}. \tag{6}$$

Our final construction (4) is obtained from composing (5) as an outer construction and (6) as inner construction. The idea is that construction (6) is invoked on n' which is about logarithmic in the alphabet size of (5), i.e., $n' = 2^{\Theta((\log n)^{1/2v})}$, so $(\log n')^{2v} = O(\log n)$. The final construction inherits its soundness error from both the outer and inner constructions, but inherits its alphabet only from the inner construction.

12.3 Query Reduction Using Polynomials

We assume a PCP verifier V_1 that uses r random bits to make q queries to a proof over alphabet Σ. The verifier has perfect completeness and soundness error ε. We show how to simulate V_1 using a new verifier V_2 that makes only $O(1)$ queries to a proof over a larger alphabet.

The general idea is this: The proof for V_2 contains a (supposed) encoding of V_1's proof as a low degree polynomial. The encoding is given by the restrictions of the polynomial to curves and tuples of points. Each curve goes through q points that represent q queries of V_1 on some randomness string. The verifier V_2 locally tests the encoding by making only $O(1)$ queries using the low degree test, and achieves low soundness error. The verifier V_2 locally decodes the q queries required for V_1 by making a single query to a curve.

The details are as follows: Let $N = 2^r q$ be the maximal length of a proof accessible by a verifier with 2^r possible tests, each accessing q locations in the proof. Let m, h be natural numbers for which $h^m = N$. Denote $d \doteq m(h-1)$. Let $\mathbb{F}$ be a finite field of characteristic two and size $|\mathbb{F}| \geq poly(d, |\Sigma|)$ for a sufficiently large polynomial as in Conjecture 3.1. Let $H \subseteq \mathbb{F}$, $|H| = h$, and associate $\{1, \ldots, N\}$ with H^m. Let $S \subseteq \mathbb{F}$, $|S| = |\Sigma|$, and associate Σ with S.

For a string $\pi \in \Sigma^N$, let $p_\pi : \mathbb{F}^m \to \mathbb{F}$ be the m-variate polynomial of degree at most $h-1$ in each of its variables for which $p_\pi(x) = \pi(x)$ for every $x \in H^m$.

For randomness $w \in \{0,1\}^r$, let $(x_{w,1}, \ldots, x_{w,q}) \in (H^m)^q$ be the q-tuple of points corresponding to the queries of V_1 on randomness w. Let $\mathcal{C}$ be a family of v-dimensional surfaces that pass through $\{(x_{w,1}, \ldots, x_{w,q})\}_w$ as discussed in Section 5.

The verifier V_2 is as follows:

VERIFIER V_2

Prescribed proof: As specified by the low degree test; supposedly the restrictions of p_π to curves in $\mathcal{C}$ and k'-tuples in $\mathcal{I}$.
Test:

1. Simulate the verifier of the low degree test; let $x_{w,1}, \ldots, x_{w,q} \in \mathbb{F}^m$ be the initial points picked by the verifier (embedded in a curve). Reject if the low degree testing verifier rejects.

[7] In fact, as we explain in Section 12.4, we need a stronger guarantee, namely a "decoding verifier".

2. Let $v_1, \ldots, v_q \in \mathbb{F}$ be the evaluations received on $x_{w,1}, \ldots, x_{w,q}$ (embedded in the assignment for the curve).

3. Reject if it is not the case that $v_1, \ldots, v_q \in S$.

4. Apply V_1 on randomness w and answers $v_1, \ldots, v_q$. Reject if V_1 rejects; accept otherwise.

The verifier V_2 uses $O(\log |\mathcal{C}|)$ random bits to make $O(1)$ queries to a proof over alphabet $\mathbb{F}^{\binom{d'+v}{v}}$. It has perfect completeness. It remains to prove soundness.

Lemma 12.3 (PCP Soundness)**.** *There are $\gamma, \gamma' = |\mathbb{F}|^{-\Omega(k')}$ for which: if there is a proof that makes V_2 accept with probability more than γ', then there is a proof that makes V_1 accept with probability more than γ.*

Proof. Assume on way of contradiction that there is no proof that makes V_1 accept with probability more than γ (to be fixed later). Apply the soundness of the low degree test for an appropriate parameter $\varepsilon = |\mathbb{F}|^{-\Omega(k')}$, and let $p_1, \ldots, p_l$ be the polynomials list decoding, $l = |\mathbb{F}|^{O(k')}$. Let $\pi_1, \ldots, \pi_l$ be the proofs that correspond to $p_1, \ldots, p_l$: For every $i \in \{1, \ldots, N\}$, the i'th position of π_j is $p_j(i)$ if $p_j(i) \in S$, and an arbitrary symbol otherwise (Recall that we associate $\{1, \ldots, N\}$ with H^m).

There are two cases in which V_2 accepts:

1. The low degree test passes although it is not the case that $v_1 = p_i(x_{w,1}), \ldots, v_q = p_i(x_{w,q})$ for some $1 \leq i \leq l$. By the low degree test soundness guarantee, this happens with probability at most ε.

2. $v_1 = p_i(x_{w,1}), \ldots, v_q = p_i(x_{w,q})$ for some $1 \leq i \leq l$, and $v_1, \ldots, v_q \in S$, and V_1 accepts π_i on randomness w. By the soundness of V_1, for every $1 \leq i \leq l$, this happens with probability at most γ. Thus, the probability it happens for some $1 \leq i \leq l$ is at most $l\gamma$.

This means that V_2 accepts with probability at most $\varepsilon + l\gamma$. Pick $\gamma = |\mathbb{F}|^{-\Omega(k')}$ so $\gamma' \doteq \varepsilon + l\gamma = |\mathbb{F}|^{-\Omega(k')}$. □

Settings of Parameters (toward (5)):

- $m, k', q, k = \Theta((\log n)^{1-1/2v})$.
- $h, |\mathbb{F}| = 2^{\Theta((\log n)^{1/2v})}$.
- $|\mathbb{F}^m| = n^{\Theta(1)}$.
- $|\mathbb{F}|^{-\Omega(k')} = 1/n^{\Omega(1)}$.
- $|\Sigma| = 2^{2^{\Theta((\log n)^{1/2v})}}$.
- $|\mathcal{C}| \leq n^{O(1)}$.

12.4 Decoding verifier

We can adapt the algebraic construction from the previous section into a "decoding" verifier, i.e., a verifier that, if it does not reject, outputs symbols from a list decoding of proofs. This variant is required for the composition scheme:

Definition 12.3 (Decoding verifier)**.** *We say that a verifier V is a decoding verifier with error probability ε and list size l for SAT_N, if the following holds: On input a formula φ on N variables, and a collection of u-tuples of variables,*

- Completeness: *For every assignment π that satisfies φ, there is a proof that V never rejects. Moreover, given access to this proof, V outputs $(x_{i_1}, v_1), \ldots, (x_{i_u}, v_u)$, where $(x_{i_1}, \ldots, x_{i_u})$ is uniformly distributed u-tuple from the given collection, and $v_1 = \pi(x_{i_1}), \ldots, v_u = \pi(x_{i_u})$.*
- Soundness: *For every proof for V, there are assignments $\pi_1, \ldots, \pi_l$ that satisfy φ, such that the probability that V does not reject and outputs $(x_{i_1}, v_1), \ldots, (x_{i_u}, v_u)$, so none of $\pi_1, \ldots, \pi_l$ satisfies $v_1 = \pi(x_{i_1}), \ldots, v_u = \pi(x_{i_u})$, is at most ε.*

For a large enough (polynomial size) field $\mathbb{F}$ with respect to q, one can obtain a decoding verifier with error $|\mathbb{F}|^{-\Omega(1)}$ and list size $|\mathbb{F}|^{O(1)}$ from the standard Sum-Check construction and the Line-vs.-Line Test [23]. Applying our query reduction technique on this decoding verifier, one obtains a decoding verifier with error probability $|\mathbb{F}|^{-\Omega(k')}$ and list size $|\mathbb{F}|^{O(k')}$.

Setting of parameters (toward (6)):

- $m = \sqrt{\log n}$.
- $h, |\mathbb{F}| = 2^{\Theta(\sqrt{\log n})}$.
- $u = \Theta(1)$.
- $k', k, q = \Theta((\log n)^{2v-1/2})$.
- $|\mathbb{F}^m| = poly(n)$.
- $|\mathbb{F}|^{\Theta(k')} = 2^{\Theta((\log n)^{2v})}$.
- $|\Sigma| \leq 2^{O(\sqrt{\log n})}$.
- $|\mathcal{C}| \leq 2^{O((\log n)^{2v})} poly(M)$.

12.5 Composition

Using a PCP verifier with low error ε but large alphabet Σ, and a decoding verifier for input size $n' \approx \log |\Sigma|$ with low error ε and small alphabet Σ', one can obtain a PCP verifier with error $O(\varepsilon)$ and alphabet Σ'. The technique, called composition, was first introduced by Arora and Safra in their breakthrough PCP paper [3]. The next lemma describes an abstract interpretation of the Arora-Safra composition.

Interestingly, while this composition lemma is in the same spirit as the combinatorial composition lemmas of Szegedy [32], Dinur-Reingold [15], Ben-Sasson et al [8] and Dinur-Harsha [14] (which is an abstraction of the composition of the author and Raz [28]), it differs from them

in its parameters and in its requirements from the initial verifiers. It preserves low error like the composition lemma of [14], but it does not require the initial verifiers to be robust. Its disadvantage is that the number of queries increases and (naturally) the output verifier is not robust.

We compose a verifier and a decoding verifier as follows. Let V_{out} be a PCP verifier for SAT_n that uses r_1 random bits to make q_1 queries to a proof over alphabet Σ_1 and achieves perfect completeness and soundness error ε_1. Let C be an error correcting code for encoding symbols from Σ_1, whose parameters are $(n', \log|\Sigma_1|, (1-\varepsilon^2/4)n')_S$ as in Proposition 2.1. For every randomness $w \in \{0,1\}^{r_1}$, consider the formula φ_w over variables $x_{1,1}, \ldots, x_{1,n'}, \cdots, x_{q_1,1}, \ldots, x_{q_1,n'}$, each ranging over S, such that φ_w is satisfied iff the variables correspond to $C(v_1), \ldots, C(v_{q_1})$ where $v_1, \ldots, v_{q_1} \in S$ are values that would make V_{out} accept on randomness w. Consider the collection of q_1-tuples $\{(x_{1,i}, \ldots, x_{q_1,i})\}_{i=1}^{n'}$. Suppose that for all $w \in \{0,1\}^{r_1}$, on input φ_w and the collection we defined, a decoding verifier V_{in} uses r_2 random bits to make q_2 queries to a proof over alphabet Σ_2 and achieves error probability ε_2 with list size l_2. Note that with the specified initial points, the verifier V_{in} decodes a uniformly random symbol $i \in [n']$ from each of the q_1 encodings of the queries of V_{out} on randomness w.

The composed verifier is as follows:

VERIFIER V

Prescribed proof:

- A proof π_1 for V_{out} written over the alphabet S, where each symbol in Σ_1 is encoded using C. We denote the length of π_1 by N_1.
- Per random string $w \in \{0,1\}^{r_1}$, the prescribed proof π_w of V_{in} for the formula φ_w, the collection of q_1-tuples we defined above, and the satisfying assignment corresponding to π_1.

1. Pick uniformly at random $w_1 \in \{0,1\}^{r_1}$.
2. Simulate V_{in} on π_{w_1}. If V_{in} rejects, reject. Otherwise, V_{in} decodes q_1 symbols $v_1, \ldots, v_{q_1} \in S$ that are supposed to equal certain symbols in π_1. If they are not equal, reject.
3. If none of the tests above rejects, accept.

In the lemma below we analyze the composed verifier. Note that we think of q_1, q_2 that are constants.

Lemma 12.4 (Composition)**.** *Suppose that $\varepsilon_2^{1/q_1}(1-\varepsilon_2)/l_2 \geq 2\varepsilon_1^{1/q_1}$ and that $\varepsilon \leq 2(\varepsilon_1/\varepsilon_2)^{1/q_1}$. The composed verifier uses $r_1 + r_2$ random bits to make $q_1 + q_2$ queries to a proof over alphabet Σ_2 and achieves perfect completeness and soundness error $O(\varepsilon_2)$.*

Proof. Without loss of generality, we assume that every symbol of π_1 is accessed by V_{out} with the same probability.

The randomness, number of queries, alphabet and perfect completeness of V are evident. Let us argue soundness. Assume that V accepts with probability at least $2\varepsilon_2$. We will argue that there exists a proof for V_{out} that makes it accept with probability at least ε_1.

For every $i \in [N_1]$, pick uniformly at random a symbol $\sigma \in \Sigma_1$ among the ones whose encoding agrees with $\pi_1(i)$ on at least ε fraction. By Johnson's bound (see Proposition 2.2), there are

at most $2/\varepsilon$ such symbols. We will argue that the expected probability that V_{out} accepts is at least ε_1. It will follow that there exists a proof with success probability at least ε_1.

For at least ε_2 fraction of the choices of w_1, with probability at least ε_2, the verifier V_{in} accepts the proof π_{w_1} and decodes q_1 symbols from π_1, one per query of V_{out} on randomness w_1. By the soundness of V_{in}, for all those w_1's, there must exist $\pi_{w_1,1}, \ldots, \pi_{w_1,l_2} \in S^{n'}$ that satisfy φ_{w_1}, and, with probability at least $1 - \varepsilon_2$, the proof π_1 agrees with one of $\pi_{w_1,1}, \ldots, \pi_{w_1,l_2}$ on the q_1 S-symbols V_{in} decodes. Hence, there must be $1 \leq p \leq l_2$ such that $\pi_{w_1,p}$ agrees with π_1 on at least $(1 - \varepsilon_2)/l_2 \geq \varepsilon$ fraction of S-symbols from each one of the q_1 symbols V_{out} queries on randomness w_1. The probability that all q_1 symbols in V_{out}'s probabilistic proof agree with $\pi_{w_1,p}$ is at least ε^{q_1}. Thus, the expected fraction of w_1's for which V_{out} accepts is at least $\varepsilon_2 \cdot \varepsilon^{q_1} \geq \varepsilon_1$. □

By composing the verifier from Section 12.3 and the decoding verifier from Section 12.4, we get Theorem 1.3.

13 Further Research

We hope that the approach for proving the Sliding Scale Conjecture suggested in this paper will eventually result in a proof of the conjecture. This would follow from a more randomness-efficient tester, either for low degree polynomials or for a modified code (e.g., "folded" low degree extension or some enhanced polynomial encoding such as multiplicity code). Any approach that works by amplification of error is subject to limitations as in [18, 26].

Acknowledgements

Dana Moshkovitz is thankful to Uri Feige, Ariel Gabizon, Shafi Goldwasser, Praladh Harsha, Sanjeev Khanna, Madhu Sudan, Ran Raz, Ronen Shaltiel, Chris Umans, Salil Vadhan, Avi Wigderson and Henry Yuen for extremely helpful discussions at various stages of this work.

References

[1] N. Alon, J. Bruck, J. Naor, M. Naor, and R. Roth. Construction of asymptotically good, low-rate error-correcting codes through pseudo-random graphs. *IEEE Transactions on Information Theory*, 38:509–516, 1992.

[2] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.

[3] S. Arora and S. Safra. Probabilistic checking of proofs: a new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.

[4] S. Arora and M. Sudan. Improved low-degree testing and its applications. *Combinatorica*, 23(3):365–426, 2003.

[5] L. Babai, L. Fortnow, L. A. Levin, and M. Szegedy. Checking computations in polylogarithmic time. In *Proc. 23rd ACM Symp. on Theory of Computing*, pages 21–32, 1991.

[6] L. Babai, L. Fortnow, and C. Lund. Nondeterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.

[7] M. Bellare, S. Goldwasser, C. Lund, and A. Russell. Efficient probabilistically checkable proofs and applications to approximations. In *Proc. 25th ACM Symp. on Theory of Computing*, pages 294–304, 1993.

[8] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. Vadhan. Robust PCPs of proximity, shorter PCPs, and applications to coding. *SIAM Journal on Computing*, 36(4):889–974, 2006.

[9] E. Ben-Sasson, M. Sudan, S. P. Vadhan, and A. Wigderson. Randomness-efficient low degree tests and short PCPs via epsilon-biased sets. In *Proc. 34th ACM Symp. on Theory of Computing*, pages 612–621, 2003.

[10] M. Braverman and A. Garg. Small value parallel repetition for general games. In *Proc. 48th ACM Symp. on Theory of Computing*, pages 335–340, 2015.

[11] M. R. Capalbo, O. Reingold, S. P. Vadhan, and A. Wigderson. Randomness conductors and constant-degree lossless expanders. In *IEEE Conference on Computational Complexity*, page 15, 2002.

[12] I. Dinur, E. Fischer, G. Kindler, R. Raz, and S. Safra. PCP characterizations of NP: Toward a polynomially-small error-probability. *Computational Complexity*, 20(3):413–504, 2011.

[13] I. Dinur and E. Goldenberg. Locally testing direct product in the low error range. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, pages 613–622, 2008.

[14] I. Dinur and P. Harsha. Composition of low-error 2-query PCPs using decodable PCPs. In *Proc. 50th IEEE Symp. on Foundations of Computer Science*, pages 472–481, 2009.

[15] I. Dinur and O. Reingold. Assignment testers: Towards a combinatorial proof of the PCP theorem. *SIAM Journal on Computing*, 36(4):975–1024, 2006.

[16] I. Dinur and D. Steurer. Analytical approach to parallel repetition. In *Proc. 47th ACM Symp. on Theory of Computing*, 2014.

[17] I. Dinur and D. Steurer. Direct product testing. In *Computational Complexity Conference*, 2014.

[18] U. Feige and J. Kilian. Impossibility results for recycling random bits in two-prover proof systems. In *Proc. 27th ACM Symp. on Theory of Computing*, pages 457–468, 1995.

[19] P. Gemmell, R. J. Lipton, R. Rubinfeld, M. Sudan, and A. Wigderson. Self-testing/correcting for polynomials and for approximate functions. In *Proc. 23rd ACM Symp. on Theory of Computing*, pages 32–42, 1991.

[20] O. Goldreich and S. Safra. A combinatorial consistency lemma with application to proving the PCP theorem. *SIAM J. Comput.*, 29(4):1132–1154, 2000.

[21] R. Impagliazzo, V. Kabanets, and A. Wigderson. New direct-product testers and 2-query PCPs. *SIAM Journal on Computing*, 41(6):1722–1768, 2012.

[22] S. M. Johnson. A new upper bound for error-correcting codes. *IRE Transactions on Information Theory*, pages 203–207, 1962.

[23] D. Moshkovitz. Lecture notes in probabilistically checkable proofs. Available on the author's webpage.

[24] D. Moshkovitz. An approach to the sliding scale conjecture via parallel repetition for low degree testing. Technical Report 30, ECCC, 2014.

[25] D. Moshkovitz. Parallel repetition from fortification. In *Proc. 55th IEEE Symp. on Foundations of Computer Science*, 2014.

[26] D. Moshkovitz, G. Ramnarayan, and H. Yuen. A no-go theorem for derandomized parallel repetition: Beyond feige-kilian, 2015.

[27] D. Moshkovitz and R. Raz. Sub-constant error low degree test of almost-linear size. *SIAM Journal on Computing*, 38(1):140–180, 2008.

[28] D. Moshkovitz and R. Raz. Two query PCP with sub-constant error. *Journal of the ACM*, 57(5), 2010.

[29] R. Raz and S. Safra. A sub-constant error-probability low-degree test and a sub-constant error-probability PCP characterization of NP. In *Proc. 29th ACM Symp. on Theory of Computing*, pages 475–484, 1997.

[30] O. Reingold, S. P. Vadhan, and A. Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. *Annals of Mathematics*, 155(1):157–187, 2002.

[31] R. Rubinfeld and M. Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.

[32] M. Szegedy. Many-valued logics and holographic proofs. In J. Weidermann, P. van Emde Boas, and M. Nielsen, editors, *Automata, Languages and Programming, 26th International Colloquium, ICALP 2007. Lecture notes in Computer Science*, pages 676–686. Springer-Verlag, 1999.

[33] A. Wigderson and D. Zuckerman. Expanders that beat the eigenvalue bound: Explicit construction and applications. *Combinatorica*, 19:245–251, 1993.

[34] D. Zuckerman. Randomness-optimal oblivious sampling. *Random Structures and Algorithms*, 11(4):345–367, 1997.