Electronic Colloquium on Computational Complexity, Report No. 30 (2014)

An Approach To The Sliding Scale Conjecture Via Parallel Repetition For Low Degree Testing

Dana Moshkovitz [*]

March 5, 2014

Abstract

The Sliding Scale Conjecture in PCP states that there are PCP verifiers with a constant number of queries and soundness error that is exponentially small in the randomness of the verifier and the length of the prover's answers. The Sliding Scale Conjecture is one of the oldest open problems in PCP, and it implies hardness of approximation up to polynomial factors for problems like Max-CSP (with polynomial-sized alphabet), Directed-Sparsest-Cut and Directed-Multi-Cut.

In this work we prove:

1. The Sliding Scale Conjecture follows from a construction of a low degree test whose soundness error is exponential in the randomness of the verifier.
2. A parallel repetition theorem for low degree testing: Given a low degree test with error $|\mathbb{F}|^{-\Omega(1)}$, one can generate a repeated low degree test whose error is $|\mathbb{F}|^{-\Omega(k)}$.
3. Applying the parallel repetition theorem on a suitable low degree test, we get a low degree test with error $|\mathbb{F}|^{-\Omega(k)}$ and randomness $O(km\log|\mathbb{F}|)$. In particular, we get the first low degree test with error $\ll 1/|\mathbb{F}|$ and $O(m\log|\mathbb{F}|)$ randomness.

The missing piece for proving the Sliding Scale Conjecture is a derandomization of the parallel repetition theorem. This seems plausible given the algebraic structure of the low degree testing problem, which was utilized for derandomization in the past. **The limitation on derandomizing parallel repetition by Feige and Kilian does not rule out this approach.**

Additional contributions in this work include an analysis of the sampling properties of the incidence graph of degree-k curves and k'-tuples of points in a finite space $\mathbb{F}^m$, and a combinatorial composition lemma for PCP that abstracts the composition technique of Arora and Safra.

[*] 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 Grant Number 1218547.

ISSN 1433-8092

1 Introduction

1.1 The Sliding Scale Conjecture

The basic PCP Theorem [6, 5, 3, 2] shows that NP problems can be verified with constant soundness error $\varepsilon = \frac{1}{2}$ using a constant number q of queries to a proof of polynomial length over a constant alphabet Σ. Given a NO instance, no matter what proof is given, the verifier accepts with probability at most ε. Given a YES instance, on the other hand, there is a proof that the verifier always accepts ("perfect completeness").

For applications, one needs a stronger PCP Theorem: one with soundness error ε close to 0. While such can be obtained by sequentially repeating the basic PCP test enough times, this also increases the number of queries. A strong PCP Theorem achieves error close to 0 while keeping the number of queries constant. We remark that $\varepsilon \geq 1/|\Sigma|^q$, and, hence, the alphabet must grow to allow low error and constant number of queries.

Bellare, Goldwasser, Lund and Russell [7] conjectured in 1993 that there are PCPs with polynomially small soundness error and two queries, provided that the alphabet is of sufficiently large polynomial size. More generally, the conjecture became known as the "Sliding Scale Conjecture":

Conjecture 1.1 (Sliding Scale Conjecture)**.** *For some constant $c > 0$, for every n and $\varepsilon \geq 1/n^c$, there is an alphabet size $k = k(\varepsilon) \leq poly(1/\varepsilon)$, such that there is a PCP verifier for input size n that uses $O(\log n)$ random bits, makes constant number of queries to a proof over alphabet of size k, and has perfect completeness and soundness error ε. In notation:*

$$NP \subseteq PCP_{1,\varepsilon}[O(\log n), O(1)]_k.$$

Remark 1.1 (Polynomially small error ⇒ general error)**.** *If one constructs PCP with constant number of queries, polynomially small error and polynomially large alphabet, then one can get PCP with constant number of queries, error ε, and alphabet $poly(1/\varepsilon)$ for all $\varepsilon \geq 1/n^c$ for some constant $c > 0$. The latter follows from composition with a Hadamard based construction (More on composition in Sub-section 1.9).*

An equivalent phrasing of the conjecture is in terms of the hardness of approximating the constraint satisfaction problem (Max-CSP). The input to Max-CSP is a system of variables and constraints over them, where each constraint depends on a constant number of variables, and each variable ranges over a finite alphabet. The problem is to find an assignment to the variables that satisfies as many constraints as possible. The Sliding Scale Conjecture is that given a Max-CSP instance with polynomial-sized alphabet, it is NP-hard to distinguish between the case that all constraints can be satisfied and the case that only a fraction of $1/n^c$ of the constraints can be satisfied for some $c > 0$. The latter implies that it is NP-hard to approximate Max-CSP to within polynomial factors.

In this work we suggest an approach to proving the Sliding Scale Conjecture based on a parallel repetition theorem for low degree testing.

1.2 Applications of The Sliding Scale Conjecture to Hardness of Approximation

As discussed above, the Sliding Scale Conjecture would imply hardness of approximation to within polynomial factors for Max-CSP. It will also imply such results for constraint satisfaction problems with specific predicates, such as quadratic equations over a large finite field, and other universal predicates (i.e., predicates that can express arbitrary constraints).

Chuzhoy and Khanna [11] showed that the Sliding Scale Conjecture implies hardness to within polynomial factors for Directed-Sparsest-Cut and Directed-Multi-Cut. In Directed-Multi-Cut, we are a given an n-vertex directed graph along with source-sink pairs, and the goal is to find the minimum cardinality subset of edges whose removal separates all source-sink pairs. Directed-Sparsest-Cut has the same input, but the goal is to find a subset of edges to delete so as to minimize the ratio of the number of deleted edges to the number of source-sink pairs that are separated by this deletion. There are likely many more applications of the Sliding Scale Conjecture along these lines. We hope that this will be the subject of more research in the future.

1.3 Previous Work

Many PCP constructions aim to achieve soundness error as low as possible. They differ from one another in the tradeoff they achieve between the error and the other parameters. Ideally, one would like to achieve PCP with two queries, but it is easier to achieve constructions with more queries. Projection games are specific kind of two query PCP that is useful for hardness of approximation and hence especially desirable. In projection games the prover's answer to the first query determines at most one satisfying answer to the second query. The size, or length, of the PCP should be as low as possible as a function of the input size n. We expect the size to be polynomial in n, or even better, almost-linear $n^{1+o(1)}$. More information about PCP parameters appears in the preliminaries. The current best constructions in terms of the soundness error are as follows:

Constant number of queries. With small constant number of queries, we have PCP theorems with soundness error $2^{-(\log n)^{\beta}}$ for some constant $\beta > 0$ and 3 queries [32, 4]. The low error in these constructions was made possible thanks to improved low degree testing theorems. The low degree testing theorem of [32] was derandomized in [27], and a corresponding PCP construction with almost linear proof length $n^{1+o(1)}$ was designed in [28]. No attempt was made to optimize the exact number of queries in [28], and the number of queries became 7.

With large constant number of queries, an almost-polynomial error of $\varepsilon \leq 2^{-(\log n)^{1-\alpha}}$ is known for *any* $\alpha > 0$; the number of queries is $poly(1/\alpha)$, the alphabet is of size $2/\varepsilon$, and the proof length is polynomial in n [12].

Projection games. For projection games, there is a PCP theorem with soundness error $1/(\log n)^{\beta}$ for some constant $\beta > 0$ and almost linear proof length $n^{1+o(1)}$ [29]. By applying parallel repetition on the construction in [29], one can improve the error to $1/(\log n)^{\alpha}$ for *any* constant $\alpha > 0$, at the expense of a large polynomial proof length $n^{\Theta(1)}$ (the exponent of $n^{\Theta(1)}$ depends on α/β) [16]. The construction in [29] relies on techniques developed for PCP with a constant, larger than 2, number of queries [27, 28]. Its novelty is in a composition technique that decreases the alphabet of a PCP construction, while preserving projection and low error. The much higher error compared to that of PCPs with more queries originates in the composition technique.

Quasi-polynomial constructions. If one considers quasi-polynomial $n^{poly \log n}$ proof length, rather than polynomial, or an almost-linear, proof length, then projection games with error $1/n$

and alphabet size $n^{O(1)}$ are known – by the parallel repetition theorem of Raz [31]. Feige and Kilian [18] showed that parallel repetition could not be derandomized (i.e., the size cannot be made $n^{O(1)}$) in the setup of the latter construction.

If one considers quasi-polynomial $n^{poly \log n}$ alphabet size, rather than polynomial alphabet size, then projection games with error $1/n$ and proof length $n^{1+o(1)}$ are also known - by an algebraic construction (Manifold vs. Point; see [29]).

Note, however, that in those constructions the error is not polynomially small in either the proof length or the alphabet size of the PCP. Moreover, the hardness of approximation results obtained from the constructions are not NP-hardness results, but reductions from inputs of SAT of size n to super-polynomial sized inputs of the approximation problem. So, lower bounds for SAT translate to much weaker lower bounds for the approximation problem. Therefore, it is preferable to refrain from quasi-polynomial constructions.

1.4 Low Degree Testing Theorems

Our approach to the Sliding Scale Conjecture centers around low degree tests, a main component in algebraic constructions of PCP.

Fix a finite field $\mathbb{F}$, and natural numbers m and d. A low degree test is a one-round multi-prover protocol for a verifier to interact with non-communicating provers that try to convince the verifier that they agree on an m-variate polynomial p of degree at most d over $\mathbb{F}$.

An example of a low degree test is the LINE-VS.-LINE TEST which was introduced by Rubinfeld and Sudan [33]:

LINE-VS.-LINE TEST

1. Pick uniformly at random a point $x \in \mathbb{F}^m$, and two lines $\ell_1, \ell_2 \ni x$.
2. For $i \in \{1, 2\}$, query prover i about ℓ_i to get a univariate polynomial $\mathcal{A}_i(\ell_i)$ of degree at most d over $\mathbb{F}$ that is supposedly the restriction of p to ℓ_i.
3. Test whether the assignments to the two lines agree on x, i.e.[1], $\mathcal{A}_1(\ell_1)(x) = \mathcal{A}_2(\ell_2)(x)$.

Note that the verifier uses $O(m \log |\mathbb{F}|)$ random bits and that when interacting with honest provers, i.e., ones that set $\mathcal{A}_i(\ell) = p_{|\ell}$ for some degree-d polynomial p for all lines ℓ, the verifier always accepts.

Ultimately, it was shown that for sufficiently large field, all, but perhaps $1/|\mathbb{F}|^{\Omega(1)}$ fraction, of the success of the test could be traced to $\mathcal{A}_i$ agreeing with one of few polynomials of degree at most d:

Lemma 1.1 (Line vs. Line low degree testing theorem [32, 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.*

Note that if the lines are randomly partitioned into $1/\delta$ sets, and for each set there is one degree-d polynomial p_i such that all lines in the set are assigned the restriction of p_i, then the

[1] We will not bother with parameterization issues. A (supposed) restriction $\mathcal{A}_i(\ell)$ of p to ℓ should also contain the (supposed) evaluation of p on any $x \in \ell$, and we use $\mathcal{A}_i(\ell)(x)$ to denote it.

test passes with probability at least δ. This clarifies why we use $O(1/\delta)$ different polynomials to explain the success of the test. We call δ the *soundness error*: the probability that the test passes without being consistent with the list decoding of low degree polynomials. Soundness error $\delta = |\mathbb{F}|^{-\Omega(1)}$ is tight for Line-vs.-Line Test; in the sequel we will see Example 1.1 that demonstrates that. Low degree testing is formalized in Section 3.

1.5 Our Results

We prove three theorems in this work. The first reduces the task of proving the Sliding Scale Conjecture to the task of designing a low degree test with exponentially small soundness error. The second, which is the heart of our work, is a parallel repetition theorem that decreases the soundness error of "robust" low degree tests exponentially. The third is a low degree testing theorem with error $|\mathbb{F}|^{-\Omega(k)}$ and randomness $O(mk \log |\mathbb{F}|)$ obtained from applying our parallel repetition theorem on a suitable low degree test.

Theorem 1 (Minimal error low degree test ⇒ Sliding Scale Conjecture)**.** *If for every d and m, for sufficiently large field $\mathbb{F}$, there is a low degree test in which the verifier uses r random bits to generate queries for $O(1)$ provers, each responding with a string of length at most $poly(d, m, \log |\mathbb{F}|)$, and the verifier achieves soundness error $\delta = 2^{-\Omega(r)}$, then there exists $c > 0$ such that*

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

The idea of the proof – by now folklore in the PCP community – is to encode the proof of a PCP with large error by a low degree polynomial, and simulate sequential repetition of the PCP by utilizing the local testability and decodability properties of low degree polynomials. The implementation of this idea differs from previous ones in several respects, like the setting of parameters, the initial PCP, and the composition.

The parameters of the polynomials used in PCP necessarily satisfy $poly(d, m, \log |\mathbb{F}|) \geq (\log n)^{\Omega(1)}$. Hence, the low degree test required in Theorem 1 has super-polynomial alphabet size. Yet, we are able to decrease the alphabet size by means of composition, and to this end we abstract the composition technique of Arora and Safra [3]. Note that we cannot use the composition technique of the author and Raz [29] (abstracted by Dinur and Harsha [14]), as its error-alphabet tradeoff does not allow polynomially small error and polynomial alphabet size. For more details about the composition lemma in this work, see Sub-section 1.9.

The PCP theorem we deduce in Theorem 1 is not a projection PCP and does not have almost-linear proof length (and so does not settle the "Projection Games Conjecture" discussed in [26]), but we believe that it might be a stepping stone toward a construction that achieves both.

The second theorem we prove in this work, a parallel repetition theorem for low degree testing, says that a "robust" low degree test with soundness error $|\mathbb{F}|^{-\Omega(1)}$ can be transformed to a low degree test with soundness error $|\mathbb{F}|^{-\Omega(k')}$. In the next section we discuss parallel repetition and state the theorem. Applying parallel repetition on an appropriate low degree test we can prove:

Theorem 2 (Low error low degree test)**.** *There is a low degree test whose soundness error is $|\mathbb{F}|^{-\Omega(k')}$, the provers' answer size is $poly(d, k', \log |\mathbb{F}|)$, and the randomness of the verifier is $O(k'm \log |\mathbb{F}|)$.*

Theorem 2 gives the *first* low degree test with error $\ll 1/|\mathbb{F}|$ and randomness $O(m \log |\mathbb{F}|)$. It might find some applications in the future.

As usual with parallel repetition, Theorem 2 has randomness larger by a factor k' than the randomness of Line-vs.-Line Test. We suggest that an approach to proving the Sliding Scale Conjecture is to derandomize parallel repetition for low degree testing and obtain a low degree test with soundness error $|\mathbb{F}|^{-\Omega(k')}$, answer size $poly(d, k', \log|\mathbb{F}|)$, and randomness $O((k' + m)\log|\mathbb{F}|)$ for $k' = \Theta(m)$.

1.6 Parallel Repetition For Low Degree Testing

(Standard) parallel repetition is a transformation on two-prover games. Suppose that G is a game in which the verifier picks questions x and y to the provers, gets from the provers answers a and b, and decides whether to accept or reject. Let k' be a natural number. In the parallel repeated game $G^{\otimes k'}$, the verifier picks k' question pairs $x_1, y_1, \ldots, x_{k'}, y_{k'}$ independently and uniformly at random; sends $x_1, \ldots, x_{k'}$ to the first prover, and $y_1, \ldots, y_{k'}$ to the second prover; gets answers $a_1, \ldots, a_{k'}$ from the first prover and answers $b_1, \ldots, b_{k'}$ from the second prover; and accepts if it would have accepted in all k' tests.

Raz's parallel repetition theorem [31] shows that if the verifier accepts with probability at most $1 - \epsilon$ in G, then the verifier accepts with probability at most[2] $(1 - poly(\epsilon))^{-\Omega(k')}$ in $G^{\otimes k'}$.

Next we define an algebraic analog of parallel repetition for low degree testing. There are several ways to define such an analog, and we choose one in which the provers' questions retain a natural algebraic structure, and the provers' answers continue to be low dimensional low degree polynomials. Specifically, we generalize Line-vs.-Line Test to Surface-vs.-Surface Test where the surfaces have constant dimension v and degree $k \geq k'$. This allows the curves to have k' points of intersection. The verifier queries the provers on surfaces. The provers respond with v-variate polynomials of degree at most dk over $\mathbb{F}$ that are supposedly the restriction of an m-variate degree-d polynomial p to the surfaces. Then, the verifier compares their answers on the k'-intersection.

We use $\mathcal{C}$ to denote the family of surfaces. We let $\mathcal{I}$ be the family of all k'-tuples of points in $\mathbb{F}^m$. In a derandomized setting one is interested in families $\mathcal{C}$ and $\mathcal{I}$ where $|\mathcal{C}|, |\mathcal{I}| \leq |\mathbb{F}|^{O(m+k')}$. Since the family $\mathcal{I}$ specifies the number of repetitions k' in the test, we omit the $\otimes k'$ from the name of the test.

Surface-vs.-Surface Test($\mathcal{C}$,$\mathcal{I}$)

1. Pick uniformly at random k' points $\{x_1, \ldots, x_{k'}\} \in \mathcal{I}$, and two dimension-$v$ degree-k surfaces $c_1, c_2 \in \mathcal{C}$ such that $c_1, c_2 \ni x_1, \ldots, x_{k'}$.
2. For $i \in \{1, 2\}$, query prover i about c_i to get a v-variate polynomial $\mathcal{A}_i(c_i)$ of degree at most dk over $\mathbb{F}$ that is supposedly the restriction of a degree-d polynomial p to c_i.
3. Test whether $\mathcal{A}_1(c_1)(x_i) = \mathcal{A}_2(c_2)(x_i)$ for all $1 \leq i \leq k'$.

Remark 1.2 (Why degree-k and not dimension-k?)**.** *The reason that we use degree-k rather than dimension-k is that in the latter the number of possible responses of the provers is approximately $|\mathbb{F}|^{d^k}$. This large alphabet is the reason that k-dimensional subspaces are usually considered in PCP for $k = \Theta(1)$ [32, 27]. In contrast, in this work we consider $k = \Theta(m)$, which would*

[2] Here $poly(\cdot)$ and $\Omega(\cdot)$ hide large constants and a dependence on the answer size of the provers in G. Better constants were achieved by Holenstein [21]. For projection games there is an even better dependence on the acceptance probability in G, and no dependence on the answer size [30].

eventually allow us to get error that is polynomially small in $|\mathbb{F}^m|$. Alphabet $|\mathbb{F}|^{d^{\Theta(m)}}$ is not sufficiently smaller than the trivial alphabet $|\mathbb{F}|^{d^m}$ that corresponds to the provers simply sending the entire m-variate degree-d polynomial to the verifier.

In analogy to standard parallel repetition, one might expect Surface-vs.-Surface Test to have error $|\mathbb{F}|^{-\Omega(k')}$. Yet, it turns out that the error – not only does not decrease with k' – but actually increases with k'. Next we show a prover strategy that makes the verifier accept with probability $\approx k'/|\mathbb{F}|$, even though the strategy does not agree with any low degree polynomial on a substantial fraction of the surfaces.

Example 1.1. *Per point $x \in \mathbb{F}^m$, pick uniformly and independently at random an m-variate polynomial p_x of degree at most d. In addition pick uniformly at random a permutation on the points in $\mathbb{F}^m$. Prover $i \in \{1,2\}$, given a surface c_i in $\mathbb{F}^m$, picks the first point $x \in c_i$ according to the permutation, and outputs the restriction of p_x to c. The following hold for this strategy:*

- *There is no single low degree polynomial that agrees with the answers of the provers on a fraction $\gg 1/|\mathbb{F}^m|$ of the surfaces.*
- *The probability that one of the k' joint elements of c_1 and c_2 turns out to be the smallest element according to the random permutation among the $2|\mathbb{F}|^v - k'$ elements in $c_1 \cup c_2$, and hence that the verifier accepts in the repeated test, is $k'/(2|\mathbb{F}|^v - k')$.*

For $k' = 1$ we can obtain a strategy with success probability $1/|\mathbb{F}|$ that is not close to a low degree polynomial by picking for each surface $c \in \mathcal{C}$ an assignment independently at random - so the value assigned to each $x \in c$ is uniformly random over $\mathbb{F}$.

Example 1.1 reflects the difference between parallel repetition for PCP and parallel repetition for testing, which is as follows. In parallel repetition for PCP one shows that a prover strategy for the repeated game that makes the verifier accept with probability at least $\delta^{\Theta(k')}$ implies the *existence* of a (possibly very different) prover strategy for the original game that makes the verifier accept with probability at least δ. On the other hand, in parallel repetition for testing one shows that the successful prover strategy for the repeated test is itself close to satisfying the tested property. The paper [13] shows that indeed parallel repetition for PCP fails to satisfy this stronger property: a successful strategy for the repeated game is not close to a repeated successful strategy for the original game.

Next we modify the repeated test to avoid the problem raised in Example 1.1. We follow a similar fix in a work by Impagliazzo, Kabanets and Wigderson [23] and include a constant number of additional provers. We set the degree of the surfaces to $k \geq 2k'$:

Surfaces Test$(\mathcal{C}, \mathcal{I})$

1. Pick uniformly and independently at random $S_0, S_1, S_2, S_3 \in \mathcal{I}$. Pick three degree-$k$ surfaces $c_1, c_2, c_3 \in \mathcal{C}$ such that $c_1 \supseteq S_0, S_1$, $c_2 \supseteq S_1, S_2$, $c_3 \supseteq S_2, S_3$.
2. Send each of $S_0, S_1, S_2, S_3, c_1, c_2, c_3$ to a different prover.
   - For each surface c_i receive a v-variate polynomial p_i of degree at most dk that is supposed to be the restriction of a polynomial p to c_i.
   - For each tuple S_j receive 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 $x \in S_j$.

Note that the definition of the repeated test is designed to simplify our analysis, rather than save in the number of provers.

Note how Example 1.1 fails: S_1 and S_2 are picked independently from the large $\mathcal{I}$. To be consistent with the prover who received c_1 as in Example 1.1, the prover who received c_2 should decide about the assignment to c_2 based on S_1. Similarly, to be consistent with the prover who received c_3 as in Example 1.1, the prover who received c_2 should decide about the assignment to c_2 based on S_2. Unfortunately, the prover who received c_2 cannot do both, and will hence either be inconsistent with at least one of the other provers.

Indeed, we can show that the repeated SURFACES TEST has error $|\mathbb{F}|^{-\Omega(k')}$, provided that the family of surfaces $\mathcal{C}$ is such that SURFACE-VS.-SURFACE TEST has "robust" error $|\mathbb{F}|^{-\Omega(1)}$. For a family of surfaces $\mathcal{C}'$ and a set of points $S \subseteq \mathbb{F}^m$, let $\mathcal{C}'_S$ denote all the surfaces in $\mathcal{C}'$ that contain all the points in S.

Definition 3 (Robust low degree test)**.** *We say that a low degree test* TEST *that compares surfaces $\mathcal{C}$ on k' tuples of points has (δ, β)-robust error ε if for any $S \subseteq \mathbb{F}^m$, $|S| \leq \beta k'$, for any sub-family $\mathcal{C}' \subseteq \mathcal{C}$, where $|\mathcal{C}'_S| \geq \delta\,|\mathcal{C}_S|$,* TEST *on $\mathcal{C}'_S$ has error at most ε.*

Theorem 4 (Parallel repetition for low degree testing)**.** *Assume that $\mathbb{F}$ is a sufficiently large field (polynomial size) in d, m and k, and that k is large enough (linear size) in k'. If* SURFACE-VS.-SURFACE TEST$(\mathcal{C}, \mathbb{F}^m)$ *has $(|\mathbb{F}|^{-\Omega(k')}, |\mathbb{F}|^{-\Omega(1)})$-robust error at most $|\mathbb{F}|^{-\Omega(1)}$, then* SURFACES TEST$(\mathcal{C}, \mathcal{I})$ *has error at most $|\mathbb{F}|^{-\Omega(k')}$.*

We need robustness of the base test because of our definition of the repeated test. In standard parallel repetition, the questions to the provers are products of the questions in the basic game. Hence, the following assertion, which is crucial for the validity of parallel repetition, holds: conditioned on likely questions and answers that lead to success of the test in up to $\beta' k'$ of the tests, the remaining questions are such that one can apply the analysis of the basic game. Here there is no obvious distinction between questions in the repeated game and questions in the basic game (both sets of questions consist of surfaces), and the premise is meant to ensure the assertion. In Section 6 we show how one can deduce a robust low degree test as in the premise of our parallel repetition from the existing analysis of LINE-VS.-LINE TEST.

Our work was inspired by the previous work of Impagliazzo, Kabanets and Wigderson [23] who focused on direct product testing by querying sized-k sets that intersect in k'-tuples of points. Although not presented this way, the analysis in [23] can also be thought of as showing a parallel repetition theorem for testing. Moreover, with some extra work, the analysis in [23] could be adapted to show a parallel repetition theorem for low degree testing. However, the analysis in [23] could only achieve soundness error $2^{-\Omega(k')}$ for $k' \leq \sqrt{k}$. This soundness error falls short of what we would expect from a parallel repetition because:

- The base of the exponent is $1/2$ instead of the error of the original test, which is roughly $1/|\mathbb{F}|$ for low degree testing.
- The exponent is the number of repetitions k' only when $k' \leq \sqrt{k}$, instead of for any k' up to a linear function in k.

Obtaining the right base for the exponent in parallel repetition – the one that corresponds to the error of the base test – is often referred to as "parallel repetition for low error" as it allows one to benefit from a low error of the original test. Parallel repetition for low error has been a challenge even for standard parallel repetition, as Raz's proof [31] does not extend in this way. The recent analysis of Dinur and Steurer [16] was the first to prove a parallel repetition theorem for low error. In this work we obtain a similar theorem for low degree testing through combinatorial techniques different from those used in [16].

Obtaining the right exponent in parallel repetition – one that is linear in k – was achieved in PCP by Raz's proof. For testing, this remained an intriguing open problem following [23]. We show that in the algebraic setting one can achieve the right exponent. To this end, we analyze the sampling properties of the incidence graph of degree-k curves and k'-tuples in $\mathbb{F}^m$. See Sub-section 1.8 for more details.

Two more remarks about the relation of our work to [23] are in order:

1. Our analysis (in particular, Sections 8, 9, 10 and 11) can be adapted to analyzing the direct product test, and can be thought of as a simplification of the analysis in [23].

2. The work [23] contains a weak derandomization through linear subspaces, but its parameters are too weak for the Sliding Scale Conjecture. See Remark 1.2.

1.7 Derandomized parallel repetition?

Uri Feige [17] observed that the limitation on derandomized parallel repetition for PCP in his work with Kilian [18] implies a certain limitation for parallel repetition of low degree testing. Our repeated test avoids this limitation.

Feige and Kilian's argument is designed for two prover games. It holds for protocols that have a small "degree" (the degree is the maximal number of possible questions to one prover given a question to the other prover). The argument is that the provers can guess a few of the questions to one of the provers with significant probability (thanks to the small degree). Assuming they succeeded, they narrowed down the space of possible remaining questions to the prover substantially. Typically, in this situation they can agree on a strategy for the rest of the questions.

Feige observed that this argument continues to hold for more than two provers, provided that the provers can all guess the questions of one prover with reasonably large probability (See Section 13 for more details). In Surfaces Test we defined, some of the provers are completely independent, and the test is designed so that the provers are unable to guess the questions of any one prover with significant probability.

In the paper we prove the soundness of the non-derandomized Surfaces Test. In the proof we use the strong testing guarantee to argue that the strategies of the independent provers must be consistent globally. Moreover, we use the algebraic setting to guarantee error low enough for the Sliding Scale Conjecture. While the proof that we show in this paper is for a non-derandomized parallel repetition, we see no reason why an argument based on a testing guarantee and low error could not be carried out for a suitable derandomized repeated test. So long as the $O(1)$ tuples picked by the verifier are uniform and independent over a sufficiently large family $\mathcal{I}$ of tuples, $|\mathcal{I}| \geq |\mathbb{F}^m|$, the Feige-Kilian limitation is avoided.

Motivated by the goal of derandomizing Surfaces Test, we suggest the following open problem:

> **The Intersecting Surfaces Problem:** Is there a family $\mathcal{C}$ of degree-k dimension-$O(1)$ surfaces in $\mathbb{F}^m$, and a family $\mathcal{I}$ of k'-tuples of points in $\mathbb{F}^m$, where $|\mathcal{C}|, |\mathcal{I}| \leq$

$|\mathbb{F}|^{O(m+k')}$, the incidence graph $\mathcal{G}(\mathcal{C}, \mathcal{I})$ is $|\mathbb{F}|^{-\Omega(k')}$-sampling, and the incidence graph $\mathcal{G}(\mathcal{I}, \mathbb{F}^m)$ is δ-dispersing for $\delta(\mu) = (\mu + |\mathbb{F}|^{-\Omega(1)})^{\Omega(k')}$?

A positive answer to the Intersecting Surfaces Problem would be a step in the direction we suggest in this paper, while a negative answer would rule out certain proof strategies for analyzing low degree tests.

As usual with derandomization of low degree testing – a random set of $poly(|\mathbb{F}^m|)$ degree-k surfaces is not expected to form intersecting surfaces. In fact, it is unlikely for such surfaces to have intersections in $\Omega(k')$ points at all! The small number of surfaces only guarantees intersections in $O(1)$ points, even though the degree of the surfaces allows intersections in k' points. The Intersecting Surfaces Problem calls for surfaces that have some "random" properties while also being "structured" in their intersections.

We note that there are explicit families $\mathcal{C}$ of size $|\mathbb{F}|^{O(m+k')}$ such that the incidence graph $\mathcal{G}(\mathcal{C}, \mathbb{F}^m)$ is $|\mathbb{F}|^{-\Theta(k')}$-sampling. Such were constructed by Guo [19] based on techniques by Ta-Shma and Umans [35] and Guruswami, Umans and Vadhan [20].

1.8 The Sampling Properties of Curves and Tuples

The incidence graph of "degree-k curves vs. k'-tuples" 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. Our analysis of parallel repetition relies on the sampling (or extractor) properties of the incidence graph "degree-k curves vs. k'-tuples". We say that the graph is a (δ, ε)-sampler if for any subset S of the k'-tuples, all curves, but at most δ fraction, have $\mu \pm \varepsilon$ fraction of their points in S, where μ is the overall fraction of tuples in S. We call δ the *sampling error*, and we call ε the *deviation error* (For background on sampling, extractors and incidence graphs, see the preliminaries). Roughly speaking, our analysis of the parallel repetition theorem shows that the error of the repeated low degree test is $\delta + \varepsilon^{k'}$.

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 [23] that the graph "k-tuples vs. k'-tuples" has sampling error $\delta = exp(-k/k')$ with a small constant deviation error ε. Indeed, the reason for the error $exp(-\sqrt{k})$ in [23] is taking $k' = \sqrt{k}$, as to balance $\delta = exp(-k/k')$ and $\varepsilon^{k'} = exp(-k')$.

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". The product we use is a replacement product for extractors, which turns out to have also appeared in a previous work [10]. The product graph essentially inherits the low sampling error of the initial graph, and hence the sampling error of "degree-k curves vs. k'-tuples" is similar to that of "degree-k curves vs. points".

Interestingly, this approach does not apply to the "k-tuples vs. k'-tuples" incidence graph relevant for [23]. The reason is that the deviation error accumulated in the k' applications of the product builds up, and the initial deviation error is not sufficiently low to withstand that.

1.9 Abstraction of Arora-Safra Composition

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. Their work led 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 [34], Dinur-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 the author and Raz [29], and was 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.

As a side-benefit of our construction, we provide an abstract version of the Arora-Safra composition. This lemma (Lemma 4.2) 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).

1.10 Organization

We start with preliminaries regarding PCP verifiers and multi-prover games, 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, and show how the Sliding Scale Conjecture follows from a derandomized low degree test in Section 4. For our parallel repetition theorem we will need to analyze the sampling properties of curves and tuples. We do so via a general paradigm for extractor products in Section 5. We address the premise of our parallel repetition theorem (a robust low degree test), and how it follows from existing analyses of low degree tests in Section 6. We outline the proof of our parallel repetition theorem in Section 7, and in the next sections we provide the proof. We discuss Feige's limitation on derandomized low degree testing in Section 13, and ideas for further research in Section 14.

2 Preliminaries

In this section we introduce notions and notation that we use throughout this work, including PCP verifiers and multi-prover games, 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 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 5 (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 2.1. *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 2.2 (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 6 (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 7 (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 8 (Low error PCP Theorem [32, 4, 12]).

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

2.2 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 2.3 (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 [24]:

Proposition 2.4 (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.*

2.3 Samplers, Dispersers 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 9 (Sampling). *For $\delta : [0,1] \times (0,1) \to (0,1)$, we say that a bi-regular bipartite graph $G = (A, B, E)$ is δ-sampling 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 $1 - \delta$ where $\delta = \delta(\mu, \varepsilon)$.

We call δ the *sampling error* and ε the *deviation error*. We will also use the terminology (δ, ε)-*sampling* when the condition in Definition 9 holds for specific δ and ε and for all $0 \leq \mu \leq 1$.

A one-sided version of sampling is given in the following definition:

Definition 10 (Dispersing). *For $\delta : [0,1] \to (0,1)$, we say that a bi-regular bipartite graph $G = (A, B, E)$ is δ-dispersing if for any set $B' \subseteq B$, $\mu = |B'| / |B|$, when one picks uniformly at random $a \in A$, the probability that all of A's neighbors land in B' is at most $\delta = \delta(\mu)$.*

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 randomness of X' is measured using *min-entropy*, and is $H_\infty(X') = \log(1/\max_x \Pr[X' = x])$.

Definition 11 (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*[3] *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 similar conventions have 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 $Ext(x, y) = (z, w)$.

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

Proposition 2.5 ([37]). *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 $(2\delta, \varepsilon)$-sampling.*

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.*

2.4 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:

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

[3] 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.

2.5 Incidence Graphs

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

Definition 12 (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": U is a finite set. A is the family of all k-tuples of points in U, and B is the family of all k'-tuples of points in U.

2. "degree-k curves vs. k'-tuples": $U = \mathbb{F}^m$ 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[4] in $\mathbb{F}^m$.

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 $\mathbb{F}^m$. So, for example, in Line-vs.-Line Test, $v = 1$ and $k = k' = 1$.

One is advised to think of the parameters as follows:

- $|\mathbb{F}|$ is large with respect to d and m. Typically, $|\mathbb{F}| = poly(d, m)$.
- 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 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 $\mathbb{F}^m$. In a derandomized test, $\mathcal{I}$ would contain fewer k'-tuples. Ideally, $|\mathcal{I}|, |\mathcal{C}| \leq |\mathbb{F}|^{O(m+k')}$. Assignments[5] 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

[4] We will often use the shorthand $B = \mathbb{F}^m$ in this case, even though Definition 12 talked about B that consists of subsets.

[5] 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.

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 surface 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 $\{(x_{i,1}, \ldots, x_{i,q})\}_{i=1}^{M} \subseteq (\mathbb{F}^m)^q$. Typically, $M \leq |\mathbb{F}^m|$. We fix $T \subseteq \mathbb{F}$, $|T| = q$. 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 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 13 (Forbidden points)**.** Forbidden points *are defined by a function* $Q : \mathcal{C} \to 2^{\mathbb{F}}$*. For a curve* $c \in \mathcal{C}$*, let* $c^{-Q} \doteq c(\mathbb{F} - Q(c))$*. For a family of curves* $\mathcal{C}$*, 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 and Line-vs.-Line 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 $\mathcal{C}_2$, a distribution $\mathcal{P}$ over pairs in $\mathcal{C}_1 \times \mathcal{C}_2$, forbidden points $Q_i : \mathcal{C}_i \to 2^{\mathbb{F}}$, and a family $\mathcal{I}$ of tuples. Throughout this work, we only consider distributions $\mathcal{P}$ where the two surfaces are independent given their intersection, and where the distribution on each of the surfaces has sufficient min-entropy. 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 $Q_2 = Q_1$. If $\mathcal{P}$ is omitted, it should be understood that it is the uniform distribution over pairs of surfaces in $\mathcal{C}_1 \times \mathcal{C}_2$ that intersect in a tuple from $\mathcal{I}$.

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

1. Pick $(c_1, c_2) \in \mathcal{C}_1 \times \mathcal{C}_2$ from the distribution $\mathcal{P}$ and a tuple $S \in \mathcal{I}$ such that $S \subseteq c_1^{-Q_1}, c_2^{-Q_2}$.

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 surfaces are lines, we refer to the test as Line-vs.-Line Test.

Curve-vs.-Curve-on-Point Test is similar to Curve-vs.-Curve Test, except that it only compares the two curves on a random point in their intersection. It is mainly useful an auxiliary test for the analysis:

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

1. Pick $(c_1, c_2) \in \mathcal{C}_1 \times \mathcal{C}_2$ from the distribution $\mathcal{P}$ and a tuple $S \in \mathcal{I}$ such that $S \subseteq c_1^{-Q_1}, c_2^{-Q_2}$.
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 surfaces 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 $\mathcal{I}$. It satisfies that the k'-tuples are independent, thus ruling out cheating strategies as in Example 1.1.

Surfaces Test$(\mathcal{C}, Q, \mathcal{I})$

1. Pick uniformly at random four tuples $S_0, S_1, S_2, S_3 \in \mathcal{I}$. Pick curves $c_1, c_2, c_3 \in \mathcal{C}$ such that c_1^{-Q} contains S_0; c_1^{-Q}, c_2^{-Q} contain S_1; c_2^{-Q}, c_3^{-Q} contain S_2; c_3^{-Q} contains S_3.
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 surfaces 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 2.4.

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:

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

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 $\mathcal{A}(S) \notin \left\{p_{1|S}, \ldots, p_{l|S}\right\}$ (this will imply the lemma). Assume, toward a contradiction, that this is not the case.

For every tuple $S \in \mathcal{I}$ such that $\mathcal{A}(S) \in \left\{p_{1|S}, \ldots, p_{l|S}\right\}$, 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 $\approx 1/|\mathbb{F}^m|$. The low degree testing theorem follows from applying our parallel repetition theorem (Theorem 17) on the low degree testing theorem in Section 6. The family of surfaces used is specified in Section 6 as well.

Theorem 14 (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:*

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

where $\mathcal{P}$ is a distribution over pairs of surfaces in $\mathcal{C}$.

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 4, 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* $\textsc{Test}$ *that uses $O((m+k')\log|\mathbb{F}|+\log M)$ random bits to make $O(1)$ queries to $\mathcal{C}$ and $\mathcal{I}$, so*

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

4 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. 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 8:

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

By sequential repetition of this PCP $O(\sqrt{\log n / \log\log n})$ times (see Proposition 2.2) 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 2.1). Ultimately, we want to prove a PCP theorem with polynomially small error and polynomial alphabet size:

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

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 4.1), we get a construction with sub-exponential alphabet:

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

In the second application (see Section 4.2), we get a construction[6] with poly-logarithmic randomness, poly-logarithmically small soundness error, and quasi-polynomial alphabet:

$$NP \subseteq PCP_{1,2^{-\Omega(\log^{2v} n)}}\left[O((\log n)^{4v-1}), O(1)\right]_{2^{\Theta(\log^{4v-1} n)}}. \tag{3}$$

Our final construction (1) is obtained from composing (2) as an outer construction and (3) as inner construction. The idea is that construction (3) is invoked on n' which is about logarithmic in the alphabet size of (2), 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.

4.1 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 surfaces and tuples of points. Each surface 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 surface.

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

[6] In fact, as we explain in Section 4.2, we need a stronger guarantee, namely a "decoding verifier".

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 $\{(x_{w,1}, \ldots, x_{w,q})\}_w$. That is, every surface in $\mathcal{C}$ contains a q-tuple of queries, and every q-tuple of queries appears on the same number of surfaces in $\mathcal{C}$.

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 surfaces 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 surface). Reject if the low degree testing verifier rejects.

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 surface).

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 4.1 (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 (2)):

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

4.2 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 15 (Decoding verifier)**.** *We say that a verifier V is a decoding verifier with error probability ε and list size l for* $\textsc{Sat}_N$*, 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 [25]. 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 (3)):

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

4.3 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 [34], Dinur-Reingold [15], Ben-Sasson et al [8] and Dinur-Harsha [14] (which is an abstraction of the composition of the author and Raz [29]), 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 $\textsc{Sat}_n$ 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 2.3. 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.

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 4.2 (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 2.4), 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 $1 \le p \le l_2$ such that $\pi_{w_1,p}$ agrees with π_1 on at least $(1-\varepsilon_2)/l_2 \ge \varepsilon$ 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 $\varepsilon_2 \cdot \varepsilon^{q_1} \ge \varepsilon_1$. □

By composing the verifier from Section 4.1 and the decoding verifier from Section 4.2, we get Theorem 1.

5 Curve-Tuple Sampling

In this section we explore the sampling properties of the "degree-k curves vs. k'-tuples" incidence graph[7], which will be used in our parallel repetition proof. We start with an argument based on k/k'-wise independence. 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 [23]), it is sufficiently small when k-tuples are replaced with degree-k curves.

5.1 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 5.1 (l'th Moment)**.** *Let $l \le k/k'$.*

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

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

Proof.

$$\begin{aligned}\mathbf{E}\left[\hat{X}^l\right] &= \binom{|\mathbb{F}|}{k'}^{-l} \cdot \mathbf{E}\left[\left(\sum_{T\subseteq\mathbb{F}} \hat{X}_T\right)^l\right] \\ &= \binom{|\mathbb{F}|}{k'}^{-l} \cdot \mathbf{E}\left[\sum_{T_1,\ldots,T_l\subseteq\mathbb{F}} \hat{X}_{T_1}\cdots\hat{X}_{T_l}\right] \\ &= \binom{|\mathbb{F}|}{k'}^{-l} \cdot \sum_{T_1,\ldots,T_l\subseteq\mathbb{F}} \mathbf{E}\left[\hat{X}_{T_1}\cdots\hat{X}_{T_l}\right] \qquad (4)\end{aligned}$$

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 $1 \le i \le l$ such that T_i is disjoint from the other T_j for $j \ne i$, 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 (4) are those where every T_i has non-empty intersection with $\bigcup_{j\ne i} T_j$ (for this we need $l \ge 2$). Their number is bounded by $\binom{|\mathbb{F}|}{k'}^l |\mathbb{F}|^{-l/2} k^l$. Each can be bounded by:

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

The proposition follows. □

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

Proposition 5.2. *For $l \le k/k'$, 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,

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

The proposition follows from Proposition 5.1. □

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". 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 5.1. *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 B. 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*[8] $\approx lk'/|\mathbb{F}|$*. Pick μ so $\mu < lk'/|\mathbb{F}| - \varepsilon$. The sampling error is roughly $(k/|\mathbb{F}| - \varepsilon)^l$.*

Note that Example 5.1 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 5.2.

5.2 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 5.2. Replacement product turns out to have been defined before in [10].

Replacement product is a generalization of a widely-used transformation by Wigderson and Zuckerman [36]. 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 16 (Replacement product for extractors)**.** *Suppose $Ext_1 : X_1 \times Y_1 \to Z_1 \times X_2$ and $Ext_2 : X_2 \times Y_2 \to Z_2 \times W_2$ are extractors. $Ext_1 \otimes Ext_2 : X_1 \times (Y_1 \times Y_2) \to (Z_1 \times Z_2) \times (X_2 \times W_2)$ is defined as follows: assume $Ext_1(x_1, y_1) = (z_1, x_2)$ and $Ext_2(x_2, y_2) = (z_2, w_2)$, then $(Ext_1 \otimes Ext_2)(x_1, y_1, y_2) = (z_1, z_2), (x_2, w_2)$.*

In contrast, the Wigderson-Zuckerman operation takes $Ext_2 : X_1 \times Y_2 \to Z_2 \times W_2$, and sets $Ext(x_1, y_1, y_2) = (z_1, z_2), (w_1, w_2)$ if $Ext_1(x_1, y_1) = (z_1, w_1)$ and $Ext_2(x_1, y_2) = (z_2, w_2)$.

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 5.3 (Replacement product lemma)**.** *If Ext_1 is a $(\delta_1, \varepsilon_1)$-extractor and Ext_2 is a $(\delta_2, \varepsilon_2)$-extractor, then $Ext_1 \otimes Ext_2$ 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 $(Ext_1 \otimes Ext_2)(X, Y_1, Y_2)$ over $Z_1 \times Z_2$ is ε-close to uniform.

Consider $z_1 \in Z_1$ whose probability according to $Ext_1(X, Y_1)$ 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 $(1/(\delta |X_1|)) \cdot (\delta |Z_1| / \delta_2) = |Z_1| / \delta_2 |X_1| = 1/(\delta_2 |X_2|)$. Hence, $H_\infty(X_{z_1}) \geq \log(\delta_2 |X_2|)$. By the property of Ext_2, the distribution defined by $Ext_2(X_{z_1}, Y_2)$ over Z_2 is ε_2-close to uniform.

[8] 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 5.2 works for the algebraic case, but not for direct products.

The total probability according to $Ext_1(X, Y_1)$ on $z_1 \in Z_1$ whose probability according to $Ext_1(X, Y_1)$ is less than $(\delta_2/\delta) \cdot (1/|Z_1|)$ is less than δ_2/δ.

By the property of Ext_1, the distribution defined by $Ext_1(X, Y_1)$ on Z_1 is ε_1-close to uniform.

Overall, we can upper bound the distance of the distribution defined by $(Ext_1 \otimes Ext_2)(X, Y_1, Y_2)$ over $Z_1 \times Z_2$ from uniform by $\varepsilon_1 + \delta_2/\delta + \varepsilon_2$. □

Corollary 5.4. *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 5.2 that analyzes the sampling properties of the incidence graph "degree-k curves vs. points" and Proposition 2.5 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 5.2, 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 incidence graph "(degree-k curves through a $(k' - 1)$-tuple) vs. points". By the induction hypothesis, the first is a $(\delta_{k'-1}/\varepsilon, 2(k' - 1)\varepsilon)$-extractor. The second is a $(\delta_{k'}, \varepsilon)$-extractor. By Lemma 5.3 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 5.4 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 $2k'\varepsilon = |\mathbb{F}|^{-\Theta(1)}$.

6 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 (Lemma 1.1). By "robust" we refer to the fact that the low degree testing theorem holds when restricting the family of surfaces to *any* sub-family of fraction $\delta = |\mathbb{F}|^{-\Theta(k)}$ of the surfaces, and even further to any such family and to surfaces that pass through given k'' points in $\mathbb{F}^m$. The construction relies on the curve-tuple sampling proved in Section 5. As explained in the introduction, the robust version is required for our parallel repetition theorem.

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 surfaces, and possibly other points. The number of surfaces 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 6.1 (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. Fix $\mathcal{C}'$ as in the premise. The proposition is proved by reduction to Line-vs.-Line Test. Consider the following variant of Line-vs.-Line Test:

$$\textsc{Line-vs.-Line Test}^{\uparrow}(\mathcal{C}'_{S'}, Q, \mathbb{F}^m)$$

1. Pick at random surfaces $s, s' \in \mathcal{C}'_{S'}$ that intersect in a point $x \in s^{-Q}, s'^{-Q}$.

2. Pick at random $\ell \subseteq s$, $\ell' \subseteq s'$ so $x \in \ell, \ell'$.

3. Define $\mathcal{A}(\ell) \equiv \mathcal{A}(s)_{|\ell}$ and $\mathcal{A}(\ell') \equiv \mathcal{A}(s')_{|\ell'}$.

4. Check that $\mathcal{A}(\ell)(x) = \mathcal{A}(\ell')(x)$.

Take $\varepsilon = |\mathbb{F}|^{-0.4}$. By the curve-tuple sampling property, the distribution of the lines picked in $\textsc{Line-vs.-Line Test}^{\uparrow}(\mathcal{C}'_{S'}, Q, \mathbb{F}^m)$ is ε-close to the distribution of the lines picked in Line-vs.-Line Test.

Applying Lemma 1.1, there exists $\delta_0 = |\mathbb{F}|^{-\Omega(1)}$ and there are m-variate polynomials $p_1, \ldots, p_l$, $l \leq O(1/\delta_0)$, of degree at most $d(q+k)$ over $\mathbb{F}$, such that the probability that $\textsc{Line-vs.-Line Test}^{\uparrow}(\mathcal{C}'_{S'}, Q, \mathbb{F}^m)$ passes, yet $A_1(\ell_1)$ is not one of $p_{1|\ell_1}, \ldots, p_{l|\ell_1}$, is at most δ_0.

If $\textsc{Surface-vs.-Surface Test}(\mathcal{C}'_{S'}, Q, \mathbb{F}^m)$ passes on surfaces s, s', then $\textsc{Line-vs.-Line Test}^{\uparrow}(\mathcal{C}'_{S'}, Q, \mathbb{F}^m)$ passes on lines ℓ, ℓ'. Moreover, the probability that $\mathcal{A}(s)$ is not one of $p_{1|s}, \ldots, p_{l|s}$, yet for a random $\ell \subseteq s$ it holds that $\mathcal{A}(s)_{|\ell}$ is one of $p_{1|\ell}, \ldots, p_{l|\ell}$, is at most $|\mathbb{F}|^{-\Theta(1)}$. □

7 Setup For Parallel Repetition Theorem

In this section we formally define our parallel repetition theorem, and outline its analysis. Our parallel repetition theorem assumes that a test that compares two surfaces on a point has error $|\mathbb{F}|^{-\Omega(1)}$, and shows that an appropriately defined k'-repeated test has error $|\mathbb{F}|^{-\Omega(k')}$. The guarantee about the base test has to hold for any sufficiently large sub-test, and in Section 6 we showed that existing low degree tests can be adapted so they satisfy this requirement. The formal statement of the parallel repetition theorem is as follows:

Theorem 17 (Parallel repetition 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 $Q' : \mathcal{C}' \to 2^{\mathbb{F}}$, $|Q'| \leq |Q| + \beta' k'$:*

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

Repeated 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 17 for curves rather than surfaces. Our arguments readily extend to surfaces.

The heart of the proof of Theorem 17 is an analysis of the two query Curve-vs.-Curve Test. As is expected from Example 1.1, 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 8, 9 and 10): 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 – the main challenge in any parallel repetition analysis – we use the sampling properties of the "degree-k curves vs. k'-tuples of points" incidence graph.

2. Weak analysis of Curves Test (Section 11): 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 12): 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.

8 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. 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.

The higher success probability is obtained by conditioning, as in the following lemma:

Pick $c_1, c_2 \in \mathcal{C}^*$, $S \in \mathcal{I}$, $S \in c_1^{-Q}, c_2^{-Q}$, as in Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$. Pick a uniformly random permutation of the points in the tuple S. For $i = 1, \ldots, k'$, let agr_i denote the event that $\mathcal{A}(c_1)$ and $\mathcal{A}(c_2)$ agree on the i'th point in S.

Lemma 8.1 (Conditional success)**.** *Suppose that the probability that* Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ *passes is at least* $|\mathbb{F}|^{-\beta k'}$. *Let* $0 < \beta' < \beta$. *Then, there exists* $0 \le k'' < \beta' k'$ *such that* $\Pr\left[agr_{k''+1} | \wedge_{i=1}^{k''} agr_i\right] \ge |\mathbb{F}|^{-\beta/\beta'}$.

Proof. By the chain rule,

$$|\mathbb{F}|^{-\beta k'} \le \Pr\left[\wedge_{i=1}^{\beta' k'} agr_i\right] = \Pr\left[agr_1\right] \cdot \Pr\left[agr_2 | agr_1\right] \cdots \Pr\left[agr_{\beta' k'} | \wedge_{i=1}^{\beta' k' - 1} agr_i\right]. \qquad (5)$$

Hence, there must exist $0 \le k'' < \beta' k'$ such that $\Pr\left[agr_{k''+1} | \wedge_{i=1}^{k''} agr_i\right] \ge |\mathbb{F}|^{-\beta k'/(\beta' k')} = |\mathbb{F}|^{-\beta/\beta'}$. □

Lemma 8.2 (Sub-Test Lemma)**.** *Suppose that the probability that* Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ *passes is at least* $|\mathbb{F}|^{-\beta k'}$. *Let* $0 < \beta' < \beta$ *and* k'' *be as in Lemma 8.1. Set* $\zeta = (1/4)\,|\mathbb{F}|^{-\beta k'/2-\beta/2\beta'}$. *Then, there exist*

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

such that the probability that Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ *passes is at least* $(3/4)\,|\mathbb{F}|^{-\beta/\beta'}$.

Proof. Pick a k''-tuple S' of points in $\mathbb{F}^m$ and assignments over $\mathbb{F}$ to the points in S' by generating a pair of curves as in Curve-vs.-Curve Test conditioned on $\wedge_{i=1}^{k''} agr_i$, and observing the k'' points of agreement and their assignments. Let $\mathcal{C}' \subseteq \mathcal{C}_{S'}$ be the set of curves that contain the k''-tuple we picked and agree with the assignment we generated for it. For $c \in \mathcal{C}'$, get $Q'(c)$ by adding to the forbidden points $Q(c)$ the locations of the points in the k''-tuple.

The distribution of curves in Curve-vs.-Curve Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ for random S' and $\mathcal{C}'$ is the same as the distribution of curves in Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ conditioned on $\wedge_{i=1}^{k''} agr_i$. Hence, from Lemma 8.1, the expected success probability of Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ is at least $|\mathbb{F}|^{-\beta/\beta'}$.

The probability that $S' \in \mathcal{I}$ is picked in Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ is $\left|\mathcal{C}^*_{S'}\right|^2 / \sum_{S \in \mathcal{I}} |\mathcal{C}^*_S|^2$. Conditioning on $\wedge_{i=1}^{k''} agr_i$, this probability is multiplied by at most $|\mathbb{F}|^{\beta k'}$.

By convexity, we can bound:

$$\frac{1}{|\mathcal{I}|} \sum_{S \in \mathcal{I}} |\mathcal{C}^*_S|^2 \geq \left(\frac{1}{|\mathcal{I}|} \sum_{S \in \mathcal{I}} |\mathcal{C}^*_S| \right)^2 = \left(\frac{|\mathcal{C}^*|}{|\mathcal{C}|} \cdot \frac{1}{|\mathcal{I}|} \sum_{S \in \mathcal{I}} |\mathcal{C}_S| \right)^2 = \left(\frac{|\mathcal{C}^*|}{|\mathcal{C}|} \, |\mathcal{C}_{S'}| \right)^2,$$

where the last inequality follows since $|\mathcal{C}_S|$ is the same for all $S \in \mathcal{I}$. Let

$$\mathcal{I}' \doteq \left\{ S \in \mathcal{I} \mid |\mathcal{C}^*_S| < \zeta \frac{|\mathcal{C}^*|}{|\mathcal{C}|} \, |\mathcal{C}_S| \right\}.$$

The probability that we picked $S' \in \mathcal{I}'$ is at most

$$\sum_{S' \in \mathcal{I}'} |\mathbb{F}|^{\beta k'} \cdot \frac{\left|\mathcal{C}^*_{S'}\right|^2}{\sum_{S \in \mathcal{I}} \left|\mathcal{C}^*_S\right|^2} \leq \sum_{S' \in \mathcal{I}'} |\mathbb{F}|^{\beta k'} \cdot \frac{1}{|\mathcal{I}|} \cdot \left(\frac{\zeta(|\mathcal{C}^*| \,/\, |\mathcal{C}|)\, |\mathcal{C}_{S'}|}{|\mathcal{C}^*|\, |\mathcal{C}_{S'}| \,/ |\mathcal{C}|} \right)^2 \leq |\mathbb{F}|^{\beta k'} \cdot \zeta^2.$$

We have $\mathbf{E}\left[|\mathcal{C}'|\right] \geq |\mathbb{F}|^{-k''} \left|\mathcal{C}^*_{S'}\right|$. Thus, similarly to the previous argument, the probability that $|\mathcal{C}'| < \zeta\,|\mathbb{F}|^{-k''} \left|\mathcal{C}^*_{S'}\right|$ is at most $|\mathbb{F}|^{\beta k'} \cdot \zeta^2$.

Therefore, there exist S' and $\mathcal{C}'$ for which the probability that Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ passes is at least $(3/4)\,|\mathbb{F}|^{-\beta/\beta'}$, while $\left|\mathcal{C}^*_{S'}\right| \geq \zeta \cdot (|\mathcal{C}^*| \,/\, |\mathcal{C}|) \cdot |\mathcal{C}_{S'}|$ and $|\mathcal{C}'| \geq \zeta\,|\mathbb{F}|^{-k''} \cdot \left|\mathcal{C}^*_{S'}\right|$. □

9 From Large Intersection To One Point Intersection

In this section 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. In other words, while Example 1.1 shows that the provers have an advantage due to the larger intersection, we are able to bound this advantage. The proof relies on the sampling property of the "degree-k curves vs. k'-tuples" incidence graph.

Lemma 9.1 (Sampler Argument). *Assume the setup of Lemma 8.2, and in particular that for S', $\mathcal{C}'$ and Q' as there, the probability that* Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ *passes is at least $|\mathbb{F}|^{-\beta/\beta'}$. 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*

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

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

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

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

$$\sum_{x\in\mathbb{F}^m} p(x) \cdot \sum_{a\in\mathbb{F}} \left(\frac{|\mathcal{C}' \cap \mathcal{C}_{x,a}|}{|\mathcal{C}' \cap \mathcal{C}_x|}\right)^2. \tag{6}$$

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

$$\sum_{x\in\mathbb{F}^m} p(x) \cdot \sum_{a\in\mathbb{F}} \sum_{c\in\mathcal{C}'\cap\mathcal{C}_{x,a}} \frac{1}{|\mathcal{C}' \cap \mathcal{C}_x|} \cdot \mathop{\mathbf{E}}_{S\in\mathcal{I}_x : S\subseteq c}[\mu_{x,a}(S)]. \tag{7}$$

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 (7) 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 (7) 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 (7) by:

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

The lemma follows from (6). □

10 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 10.1 (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:

- $$\gamma_0 \leq (1/2) \cdot |\mathbb{F}|^{-\beta/\beta'},$$
- *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 9.1.*

Repeated Test: *If Curve-vs.-Curve Test$(\mathcal{C}^*, Q, \mathcal{I})$ passes with probability at least $|\mathbb{F}|^{-\beta k'}$, 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 $\gamma'((1/2) \cdot |\mathbb{F}|^{-\beta/\beta'}) \cdot \delta$ 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 $|\mathbb{F}|^{-\beta k'}$. Let $0 < \beta' < \beta$ and $k'' < \beta' k'$ be as in Lemma 8.1. By Lemma 8.2, there exist $S' \in (\mathbb{F}^m)^{k''}$; $\mathcal{C}' \subseteq \mathcal{C}_{S'}$; $Q' : \mathcal{C}' \to 2^{\mathbb{F}}$; such that $|\mathcal{C}'| = \left|\mathcal{C}'_{S'}\right| \geq \delta \left|\mathcal{C}_{S'}\right|$, and the probability that Curve-vs.-Curve-on-Point Test$(\mathcal{C}', Q', \mathcal{I}_{S'})$ passes is at least $(3/4)\,|\mathbb{F}|^{-\beta/\beta'}$. By Lemma 9.1, Curve-vs.-Curve Test$(\mathcal{C}', Q', \mathbb{F}^m)$ passes with probability at least $(1/2) \cdot |\mathbb{F}|^{-\beta/\beta'}$. 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 $\gamma' = \gamma'((1/2) \cdot |\mathbb{F}|^{-\beta/\beta'})$ in $\mathcal{C}'$. We have

$$|\mathcal{C}_p| \geq \gamma' \left|\mathcal{C}'\right| \geq \gamma' \delta \left|\mathcal{C}_{S'}\right|.$$

□

11 Curves Test Analysis

In this section we use the analysis of Curve-vs.-Curve Test in Section 10 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 11.1 (Analysis of Curves Test)**.** *Using the notation of Lemma 10.1, and under its assumptions about the parameters and the base test:*

Repeated Test:

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

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

By Lemma 10.1, there exists an m-variate polynomial p of degree at most d' over $\mathbb{F}$, such that with probability at least $\gamma'((1/2) \cdot |\mathbb{F}|^{-\beta/\beta'}) \cdot \delta$ 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 $|\mathbb{F}|^{-\beta k'}$ fraction of the S_2's. □

12 Concluding The Analysis

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

Lemma 12.1 (decoding → list decoding)**.** *Under the assumptions of Lemma 11.1, there exists $\delta_0 = |\mathbb{F}|^{-\Omega(k')}$, such that*

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

Proof. By Lemma 11.1 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 12.2. *Using the assumptions and notation of Lemma 12.1, and assuming that*[9] $(\delta_0^{3/2}/2) \cdot \binom{|\mathbb{F}|-|Q|}{k'} > \binom{dk}{k'}$,

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

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 12.1,

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

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,

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

□

13 Limitations On Derandomized Parallel Repetition

Uri Feige [17] observed that his work with Kilian about the limitations of derandomized parallel repetition in PCP [18] implies a certain limitation also for derandomizing Surfaces Test. In this section we discuss Feige's argument. We start with a combinatorial lemma that follows from [18]. At a first glance, it seems like this lemma yields a provers strategy for Surfaces Test that is far from any low degree polynomial, yet passes the test with too high of a probability. On a closer inspection, it turns out that **the lemma says nothing when the $O(1)$ intersections between the surfaces are picked independently from $\mathcal{I}$.** The lemma does rule out several other formulations of a derandomized Surfaces Test that look natural a-priori.

Consider the graph $G_{\mathcal{C}} = (\mathcal{C}, E_{\mathcal{C}})$ that has an edge $e = (c_1, c_2)$ if c_1 and c_2 intersect by a k'-tuple from $\mathcal{I}$. We denote the intersection tuple by $\mathcal{I}[e]$. The heart of Feige's argument is

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

that with significant probability over the random choice of an edge, the intersection between the surfaces falls in a small set of points associated with the surfaces. Below think of $s = (1/|\mathbb{F}|^{\omega(1)}) \cdot |\mathbb{F}^m|$ and $\gamma_s = |E_{\mathcal{C}}|^{-\omega(1/m)}$. Recall that we are after polynomially small error $|E_{\mathcal{C}}|^{-\Theta(1)}$.

Lemma 13.1 (Follows from Theorem 6 in [18]). *Let $s > 2|\mathbb{F}|^3$. There exists an assignment $F_{\mathcal{C}}$ that assigns each surface $c \in \mathcal{C}$ a set $F_{\mathcal{C}}(c) \subseteq \mathbb{F}^m$ of at most s points, such that for at least $\gamma_s \doteq |E|^{-2\log|\mathbb{F}|/\log(s/2|\mathbb{F}|)}$ fraction the edges $e = (c_1, c_2) \in E_{\mathcal{C}}$, the surfaces c_1 and c_2 are assigned the same set, and the set contains $\mathcal{I}[e]$.*

For every tuple $S \in \mathcal{I}$, assign a set $F_{\mathcal{I}}(S)$ as follows: pick at random $e = (c_1, c_2) \in E_{\mathcal{C}}$ with $\mathcal{I}[e] = S$. If c_1 and c_2 are assigned the same set $F_{\mathcal{C}}(c_1) = F_{\mathcal{C}}(c_2)$ and this set contains S, assign it to S. Otherwise, assign S to S.

Consider the following provers strategy: for every set $F \subseteq \mathbb{F}^m$ of at most s points, let $p[F]$ be a random m-variate polynomial of degree at most d over $\mathbb{F}$. For every $c \in \mathcal{C}$, let $\mathcal{A}(c)$ be the restriction of $p[F_{\mathcal{C}}(c)]$ to c. For every tuple $S \in \mathcal{I}$, let $\mathcal{A}(S)$ be the restriction of $p[F_{\mathcal{I}}(S)]$ to S.

Let us make the following assumption, which holds for the family of all k'-tuples of points in $\mathbb{F}^m$, and it is reasonable to expect that it holds for $\mathcal{I}$ as well:

Definition 18 (Tuple Dispersing). *We say that $\mathcal{I}$ is tuple dispersing if for $s = (1/|\mathbb{F}|^{\omega(1)}) \cdot |\mathbb{F}^m|$, for any set of s points, at most $|\mathbb{F}|^{-\omega(k')}$ fraction of the k'-tuples are contained in the set.*

Under the tuple dispersing assumption, the provers strategy we defined does not agree with any low degree polynomial on a fraction $|\mathbb{F}|^{-\Omega(k')}$ of the tuples. A-priori, it seems like Lemma 13.1 ought to imply that the strategy succeeds in Surfaces Test with probability $poly(\gamma_s)$ (hence demonstrating that polynomially small error cannot be attained for Surfaces Test). However, this does *not* hold as long as the $O(1)$ intersections between the surfaces in Surfaces Test are chosen independently from $\mathcal{I}$.

The strategy we defined does succeed with probability $poly(\gamma_s)$ for some formulations of a derandomized Surfaces Test:

- Given c_2, the choice of $S_1, S_2 \subseteq c_2$ in Surfaces Test is uniformly random. For such a test, the strategy we defined succeeds with probability at least γ_s^2.

While this may seem like a reasonable formulation a-priori, it is a misleading one, as it picks intersecting S_1 and S_2 with probability $\approx 1/|\mathbb{F}|$. For such a test, a provers strategy like the one in Example 1.1 succeeds with a high probability $\approx 1/|\mathbb{F}|$.

- Given c_2, the choice of $S_1, S_2 \subseteq c_2$ in Surfaces Test is uniformly random conditioned on $S_1 \cap S_2 \neq \phi$. For such a test, the strategy we defined succeeds with probability at least $poly(\gamma_s)$.

This formulation was considered by Feige [17]. Note that the framework we outlined above does not fall into this criterion.

14 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 derandomizing our parallel repetition for low degree testing, either for SURFACES TEST or for a different test, for either low degree polynomials or for a modified code (e.g., "folded" low degree extension or some enhanced polynomial encoding such as multiplicity code). The reader who is interested in pursuing this direction is advised to read Section 13, where we discuss a certain limitation on a derandomized SURFACES TEST, as well as a way around it. Additional limitations on derandomizing the test would be very interesting as well, as they might lead to better hypotheses.

There are many potential applications to handling the projection case, and settling the Projection Games Conjecture [26], with, or without, almost linear proof length. To achieve a projection PCP, one would have to adapt our low degree testing theorem to projection, and also devise a composition technique that works for polynomially small error and projection. Other generalizations of the Sliding Scale Conjecture concern verifiers that make linear tests (with imperfect completeness), and "smooth" verifiers, i.e., verifiers in which the legit views of the verifier form an error correcting code [22].

Acknowledgements

Dana Moshkovitz is thankful to Uri Feige, Ariel Gabizon, 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. Julia Chuzhoy pointed out the applications of the Sliding Scale Conjecture to DIRECTED-SPARSEST-CUT and DIRECTED-MULTI-CUT [11].

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. 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.

[11] J. Chuzhoy and S. Khanna. Polynomial flow-cut gaps and hardness of directed cut problems. *Journal of the ACM*, 56(2), 2009.

[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. The structure of winning strategies in parallel repetition games. In Maria J. Serna, Ronen Shaltiel, Klaus Jansen, and Jos D. P. Rolim, editors, *APPROX-RANDOM*, volume 6302 of *Lecture Notes in Computer Science*, pages 518–530. Springer, 2010.

[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. *CoRR*, abs/1305.1979, 2013.

[17] U. Feige. On intersection graphs and the Z-test. Private communication, 2011.

[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] Z. Guo. Randomness-efficient curve samplers. In *RANDOM*, pages 575–590, 2013.

[20] V. Guruswami, C. Umans, and S. Vadhan. Unbalanced expanders and randomness extractors from parvaresh–vardy codes. *Journal of the ACM*, 56(4), 2009.

[21] T. Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(1):141–172, 2009.

[22] J. Holmerin and S. Khot. A new PCP outer verifier with applications to homogeneous linear equations and max-bisection. In *Proc. 36th ACM Symp. on Theory of Computing*, pages 11–20, 2004.

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

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

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

[26] D. Moshkovitz. The projection games conjecture and the NP-hardness of $\ln n$-approximating set-cover. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 15th International Workshop, APPROX 2012*, volume 7408, pages 276–287, 2012.

[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. Sub-constant error probabilistically checkable proof of almost-linear size. *Computational Complexity*, 19(3):367–422, 2010.

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

[30] A. Rao. Parallel repetition in projection games and a concentration bound. *SIAM Journal on Computing*, 40(6):1871–1891, 2011.

[31] R. Raz. A parallel repetition theorem. In *SIAM Journal on Computing*, volume 27, pages 763–803, 1998.

[32] 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.

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

[34] 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.

[35] A. Ta-Shma and C. Umans. Better lossless condensers through derandomized curve samplers. In *Proc. 47th IEEE Symp. on Foundations of Computer Science*, pages 177–186. IEEE Computer Society, 2006.

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

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