Electronic Colloquium on Computational Complexity, Report No. 66 (1998)

PCP Characterizations of NP: Towards a Polynomially-Small Error-Probability

Irit Dinur* Eldar Fischer* Guy Kindler* Ran Raz† Shmuel Safra*

November 4, 1998

Abstract

This paper strengthens the low-error PCP characterization of NP, coming closer to the ultimate BGLR conjecture. Namely, we prove that witnesses for membership in any NP language can be verified with a constant number of accesses, and with an error probability exponentially small in the number of bits accessed, as long as this number is at most $(\log n)^\beta$ for *any* constant $\beta < 1$. (Compare with the BGLR conjecture, claiming the same for any $\beta \leq 1$).

Our results are in fact stronger, implying that the constant-depend Gap-Quadratic-Solvability problem with field-size $2^{\log^\beta n}$ is NP-hard. We show that given a system of quadratic polynomials each depending on a constant number of variables ranging over a field $\mathcal{F}$ ($|\mathcal{F}| \approx 2^{\log^\beta n}$ for any constant $\beta < 1$), it is NP-hard to decide between the case where there is a common root for all of the polynomials, and the case where every assignment zeros at most an $O(\frac{1}{|\mathcal{F}|})$ fraction.

At the same time, our proof presents a *direct* construction of a low-degree-test whose error-probability is exponentially small in the number of bits accessed. Such a result was previously known only by relying on recursive applications of the PCP theorem.

Introduction

Cook's characterization of NP asserts that for any NP language L, there exists a polynomial-time algorithm that, given an input string, produces a set Ψ of Boolean functions (tests) over a common set of variables, each test depending only on a constant number of the variables. The set of tests represents membership of the input in L in the sense that there exists an assignment satisfying all tests if and only if the input is in L.

A PCP characterization of NP differs from Cook's in regards to what can be guaranteed in case the input is not in the NP language L. While in Cook's characterization, no assignment would satisfy all Ψ's tests, in a PCP characterization of NP it is guaranteed that no assignment satisfies even a small fraction of Ψ in case the input is not in L.

A satisfying assignment to Ψ can be viewed as a membership-proof to the fact that the input is in L. One which, according to the PCP characterization of NP, can be probabilistically verified by randomly picking a test of Ψ and checking whether it holds – hence the term, Probabilistic Checking of Proofs. The *error probability* of a PCP system is the maximal fraction of Ψ that can be satisfied in case the input is not in L.

Introducing the PCP scheme of characterizing NP has created an avalanche of hardness results for approximation problems ([FGL+91, AS92, ALM+92, LY94, BGLR93, BGS98, Hås97], to mention only a few). For most of these applications, the characterization of NP with constant error-probability and variables of constant

* School of Mathematical Sciences, Tel Aviv University, ISRAEL
† Weizmann Inst. of Science , ISRAEL

ISSN 1433-8092

range [AS92, ALM+92] suffices. In order to prove NP-hardness of other problems however, sub-constant error-probability had turned out to be essential.

For example [LY94] and [BGLR93] were able to prove approximating SET-COVER to within logarithmic factors *almost NP-hard*, using the constant error-probability PCP characterization of NP. As to proving NP-hardness for that problem, [BGLR93] had suggested the 'sliding scale' conjecture.

The BGLR conjecture states that, even if the number of variables accessed by every test is kept constant, the error probability can be guaranteed to be polynomially small in the size of the variables' range. In other words, a membership-proof can be verified by accessing a constant number of words, such that the error probability is exponentially small in the length of a word.

One cannot expect the error-probability to be much smaller than inverse the size of the variables' range, as a random assignment would expectedly satisfy such a fraction, no matter what structure the tests take (recall that each test depends on a constant number of variables). Hence the BGLR conjecture is optimal in that sense.

It is also unlikely for the BGLR conjecture to hold for variables' range larger than polynomial, since the error-probability is unlikely to be smaller than polynomially small (if it were, then no test succeeds in case the input is not in L, making the decision between the cases easy).

The conjecture was eventually shown true for a sizeable portion of the applicable range-size in [RS97], which showed a PCP characterization of NP, where the error-probability is exponentially small in the number of bits required to represent variables, as long as that number of bits is up to $(\log n)^{\beta}$ for *some* positive constant β (see also [AS97]).

Our Main Results

In this paper, we improve the above result in regards to the range of β for which it holds, substituting "some constant" with "any constant smaller than 1".
Our results are, in fact, stronger: we show that given a system of quadratic polynomials each depending on a constant number of variables ranging over a field $\mathcal{F}$ ($|\mathcal{F}| \approx 2^{\log^{\beta} n}$ for any constant $\beta < 1$), it is NP-hard to decide whether there is a common root to all of the polynomials, or if every assignment zeros no more than an $O(\frac{1}{|\mathcal{F}|})$ fraction.

Our proof also gives a direct construction of a low-degree-test whose error-probability is exponentially small in the number of bits accessed, as discussed below.

Related Results

As mentioned above, the BGLR conjecture claims a PCP characterization of NP whose error-probability is exponentially small in the number of bits accessed, as long as that number is at most $O(\log n)$.

Focusing on the smallest possible error-probabilities (i.e. polynomial), we note that even when allowing super-polynomial reductions, no PCP result obtains an error-probability polynomially small in the size of the generated instance when restricted to a constant number of accesses. For example, the repetition lemma of [Raz98] shows that with two accesses to $O(\log n)$ bits, the error-probability is polynomially small, but only in the size of the original input, while the size of the construction is $n^{\log n}$. Similarly, the multi-linear extension of [BFL91], yields a system with a $\frac{1}{n}$ error-probability but whose size is $n^{\log n}$. In fact, in any known reduction, there is always a factor of at least $\log^{\epsilon} n$ in the exponent that separates the error-probability from the size of the generated instance. Reaching an error-probability polynomial in the size of the generated instance is an important open problem. Such a characterization of NP would improve hardness results for several problems. For example, NP-hardness for approximating the 'Monotone-Minimum-Satisfying-Assignment' problem (that is closely related to approximating the length of propositional proofs [ABMP98]) has been shown [DS98] via a reduction from PCP (such that the

non-approximation ratio is preserved). Hence a polynomially small error-probability PCP characterization of NP would immediately imply that it is NP-hard to polynomially approximate the length of propositional proofs.

[RS97] managed to keep the exponential relation between the number of bits accessed and the error-probability, thus showing the BGLR conjecture true for variables' range-size of up to $2^{\log^{\beta} n}$ for *some* constant $\beta < 1$. For larger β (any constant $\beta < 1$) [RS97] showed a system whose error probability is $2^{-\log^{\beta} n}$, yet without the exponential relation between the number of bits accessed and the error-probability (the number of bits accessed was $O(\log^{\beta} n \cdot \text{poly} \log\log n)$). This factor of $\text{poly} \log\log n$ is signifi cant when viewing, for example, the result in terms of Gap-Quadratic-Solvability. The result of [RS97], if it were to be translated to Gap-Quadratic-Solvability terms, would at best give an equation system with each equation depending on $O(\text{poly} \log\log n)$ variables. In comparison, our result translates to a quadratic equation system with the same error-probability, but where every equation depends on a *constant* $O(\frac{1}{1-\beta})$ number of variables.

An essential component of the known proofs for PCP characterizations of NP, involves a consistency test for low-degree-polynomials (low-degree-test). A low-degree-test of small error-probability seems necessary in order to show a PCP characterization of NP of small error-probability. Such low-degree-tests with small error-probability appeared in [RS97, AS97].

Technique

We use the general framework of [AS92, ALM+92, RS97] for our proof, however separate the proof into two independent parts: the sum-check part and the LDF-reader part (LDF = low-degree-function). We replace the generalized form of the composition paradigm utilized in previous PCP proofs, with a concrete representation of an LDF and a direct low-degree-test whose error-probability is exponentially small in the number of bits accessed. This gives our proof the form of a series of manipulations of quadratic equations, and produces the claimed NP-hardness for Gap-Quadratic-Solvability with a fi eld of size $2^{\log^{\beta} n}$ for any constant $\beta < 1$.

The general composition-recursion paradigm calls for the recursive application of a PCP theorem on tests of a previously obtained PCP system. Our technique is more subtle, replacing that general paradigm with a concrete representation.

An essential component of the known proofs for PCP characterizations of NP, involves a consistency test for low-degree-polynomials (low-degree-test). A low-degree-test of small error-probability seems necessary in order to show a PCP characterization of NP of small error-probability. Such low-degree-tests with small error-probability appeared in [RS97, AS97]. [RS97] introduced a new low-degree-test (plane-vs.-plane) and showed it to be of small error-probability. [AS97] proved the previously used (line-vs.-point) test to be of small error-probability.

The error probability of all these tests is still not exponentially small in the number of bits accessed. In fact, in any such direct low-degree-test comparing subspaces (lines, planes, etc.) for consistency, the error-probability must be at least polynomial in the number of bits accessed.

One can attain a more complex test with an exponentially small error-probability by utilizing the composition technique, applying the entire PCP theorem to the known low-degree-test. Our proof, in contrast, avoids the general composition-recursion paradigm, utilizing instead a tree-representation for LDFs for which a direct low-degree-test is presented, whose error-probability is exponentially small in the number of bits accessed, while maintaining a constant number of accesses.

In order to achieve the result pertaining to Quadratic-Solvability (which is stronger than the general PCP characterization of NP) we take special measures to keep the equations quadratic throughout the proof. The proof starts from the Gap-Quadratic-Solvability problem (where no bound is set for the number of variables each equation depends on), shown NP-hard in [HPS93], and takes the form of a series of manipulations of quadratic equations. By regarding a quadratic equation as a linear equation whose variables are quadratic products, the proof manages to appear as a series of *linear* manipulations of equations, with an independent multiplication-test part, and thus the

quadratic degree of the equations is maintained throughout the proof. Observe, for example, that the sum-check technique is linear in the sense that its tests involve no multiplication whatsoever, yet since it is invoked over a quadratic equation, the resulting equations are quadratic. The equations maintain their form due to the fact that, in contrast to previous PCP proofs, our proof uses *the same field throughout* (note that a linear equation over one field is non-linear over another).

This 'linearity with an independent quadratic part', enables termination of the recursion (resembling [ALM+92]) but with a smaller blow-up in size, thus yielding the lower error probability $2^{-\log^{\beta} n}$ for any constant $\beta < 1$. Thus we also attain the NP-hardness of constant-depend Gap-Quadratic-Solvability with a field of size $2^{\log^{\beta} n}$ for any constant $\beta < 1$.

This paper presents a complete proof of the PCP theorem, relying only on the low-degree-test of [RS97].

Gap Quadratic Solvability

Our theorem is in fact stronger than the PCP characterization claimed, imposing some additional structure on the tests. We define the gap version of the quadratic solvability maximization problem as follows,

Definition 1 ($QS_n(D, F, \epsilon)$) *An instance of the Gap Quadratic Solvability problem (with parameters n, D, F, ϵ) is a set of quadratic equations over a finite field $\mathcal{F} = Z_p$ with $F \leq |\mathcal{F}| \leq 2F$, such that the size of the system is n, and each equation depends on D variables at most. The problem is to distinguish between the following two cases:*

Yes. There is an assignment to the variables that satisfies all of the equations.

No. Every assignment to the variables satisfies no more than an ϵ fraction of the equations.

For instances that fall in neither of the cases, any outcome is acceptable.

We often use values for $\epsilon = \epsilon(n)$, $D = D(n)$, and $F = F(n)$ that are functions of n. The NP-hardness of $QS_n(n, F(n), O((F(n))^{-1}))$ is proven in [HPS93] for any function $F(n) < n$, by a simple pseudo-random amplification technique. In this paper we prove the following theorem.

Theorem 1 (Main Theorem) $\forall \delta, c > 0$, $\exists\, c_1, c_2 > 0$ *(constants) such that $QS_n(c_1, F(n), \frac{c_2}{F(n)})$ is NP-hard, for $F(n) = 2^{c(\log n)^{1-\delta}}$.*

This theorem is in fact true (and obtained via the same techniques) for any function $F(n) < 2^{c(\log n)^{1-\delta}}$. In this paper, however, we focus only on the largest field size.

We would like to suggest the following conjecture, which can be viewed as the QS parallel for the [BGLR93] conjecture:

Conjecture 2 $\exists c_1, c_2, c_3 > 0$ *constants, such that $QS_n(c_2, n^{c_1}, c_3 \cdot n^{-c_1})$ is NP-hard.*

Structure of the Paper

The proof of the main theorem (theorem 1) proceeds by a series of manipulations to reduce a $QS_n(n, F(n), \frac{O(1)}{F(n)})$ equation system to a $QS_{poly(n)}(O(1), F(n), \frac{O(1)}{F(n)})$ equation system.

Section 1 describes equation systems and the general manipulations performed throughout the proof. We introduce a new notion of a *gap under assumptions*, that is, a guarantee that if we select only assignments from a specified subset of assignments, the maximal fraction of satisfied equations must be on either sides of the gap. We use this notion to separate the proof into two major components. One component (section 4) transforms the

system into one of constant depend, however one which possesses a gap property under some assumptions (gap-under-assumption). The other component (section 3) eliminates those assumptions – using the aforementioned low-degree-test – resulting in a system that possesses the gap property without any assumptions. The proof of the main theorem appears in section 2, and accordingly, is broken into two main lemmas.

Many of the techniques used in this paper are similar to those used in previous proofs for PCP characterizations of NP, however some additional ideas and amendments had to be incorporated.

1 Equation Systems

1.1 Basic Definitions

The proof of our main theorem proceeds by a reduction from $QS_n(n, F(n), \frac{O(1)}{F(n)})$, shown NP-hard in [HPS93]. We thus begin with a set of quadratic equations that have a gap-property. Our final goal is to transform it into another set of equations that possesses both the gap and constant depend parameter. Let us define an *equation system*. An equation system is a system $\Psi = \{\psi_1, ..., \psi_m\}$ of Boolean functions over variables $\mathcal{V}$ such that every $\psi \in \Psi$ is a *conjunction of equations*. We call such a Boolean function a *conjunction*.

We say that $\psi \in \Psi$ is a degree-d conjunction if it is a conjunction of degree-d equations. We say that Ψ is a *degree-d* (*esp. linear or quadratic*) equation system if every $\psi \in \Psi$ is a degree-d conjunction. The range of the variables $\mathcal{V}$ is a finite field $\mathcal{F}$.

The property of an equation system that we will be most interested in is the gap of the system.

Definition 2 ($(\epsilon, 1)$-Gap of an equation system) *Let $\Psi = \{\psi_1, .., \psi_m\}$ be an equation system over variables $\mathcal{V}$. Ψ is said to have an $(\epsilon, 1)$-gap if either of the following holds.*

- *There exists an assignment to $\mathcal{V}$ satisfying all ψ_i's.*
- *Every assignment to $\mathcal{V}$ satisfies no more than an ϵ fraction of the ψ_i's.*

ϵ is also called the error probability of the system, a term inspired by the following procedure. Suppose we are given an assignment for Ψ. We would like to verify that the assignment satisfies Ψ. We pick a random conjunction $\psi \in_R \Psi$ and return its value. If the assignment satisfies the entire system, ψ is surely satisfied, and we always obtain the right answer. If the system is not satisfied, we obtain TRUE with probability at most ϵ. ϵ is therefore an upper bound for the error probability of the procedure. An instance of $QS_n(D, F, \epsilon)$, in these terms, is an equation system with an $(\epsilon, 1)$-gap. We will prove that distinguishing between the two sides of the gap is NP-hard (with values for D, F, ϵ as stated).

The [HPS93] NP-hardness result for $QS_n(n, F(n), \frac{O(1)}{F(n)})$ appears very similar to our end goal, theorem 1; yet there is one parameter that makes all the difference: the *depend* (D) of the equation system.

Definition 3 (Depend) *For a conjunction $\psi \in \Psi$, we denote by $depend(\psi)$ the number of variables ψ depends on. The depend of an equation system Ψ is the maximum depend over $\psi \in \Psi$:*

$$depend(\Psi) = \max_{\psi \in \Psi} (depend(\psi))$$

The goal of this paper is to lower the depend parameter of the [HPS93] equation system – namely $O(n)$ – to constant.

1.2 Assumptions

Let Ψ be an instance of $QS_n(n, F(n), \frac{O(1)}{F(n)})$ with $F(n) = 2^{\Theta((\log n)^{1-\delta})}$; i.e. Ψ is a quadratic equation system with variables $\mathcal{V}$ ranging over a field $\mathcal{F}$. Ψ has a $(\frac{O(1)}{|\mathcal{F}|}, 1)$-gap, yet $depend(\Psi) = O(n)$.
Denote by $\mathrm{Assign}(\mathcal{V})$ the set of all possible assignments to $\mathcal{V}$,

$$\mathrm{Assign}(\mathcal{V}) = \{A \mid \ A : \mathcal{V} \to \mathcal{F}\}$$

For every assignment $A \in \mathrm{Assign}(\mathcal{V})$ define the weight function

$$\omega(\Psi, A) \stackrel{def}{=} \Pr_{\psi \in \Psi}[\psi \text{ is satisfied by } A]$$

Saying Ψ has an $(\epsilon, 1)$-gap, is equivalent to $\omega(\Psi) \notin (\epsilon, 1)$ for

$$\omega(\Psi) \stackrel{def}{=} \max_{A \in \mathrm{Assign}(\mathcal{V})}(\omega(\Psi, A))$$

Consider the following modification of the weight function. Instead of taking the maximum over *all* the assignments in $\mathrm{Assign}(\mathcal{V})$, we narrow the scope to assignments in a subset $\mathsf{A} \subset \mathrm{Assign}(\mathcal{V})$, viewing A as a constraint on the assignments. We thus define

$$\omega_{\mathsf{A}}(\Psi) \stackrel{def}{=} \max_{A \in \mathsf{A}}(\omega(\Psi, A))$$

and require that $\omega_{\mathsf{A}}(\Psi) \notin (\epsilon, 1)$.

We will be considering equation systems that when restricting their assignments to a specified set A, display a gap (measured by the new weight function ω_{A}) between 'yes' and 'no' instances. We slightly strengthen this notion by widening the gap between the 'yes' and the 'no' cases. This is done by introducing a narrower constraint $\mathsf{A}_{\mathtt{yes}} \subseteq \mathsf{A}$, and requiring that if $\omega_{\mathsf{A}}(\Psi) = 1$ (a 'yes' instance), there must also be an $A' \in \mathsf{A}_{\mathtt{yes}}$, such that $\omega(\Psi, A') = 1$.

The pair A and $\mathsf{A}_{\mathtt{yes}}$ are an *assumption-pair*, or an *assumption*, on the equation system. We sometimes denote the larger set in the pair (A) by $\mathsf{A}_{\mathtt{no}}$. This notation emphasizes the fact that for a 'no' instance, any assignment $A \in \mathsf{A}_{\mathtt{no}}$ has $\omega(\Psi, A) \leq \epsilon$; and for a 'yes' instance, there exists an $A \in \mathsf{A}_{\mathtt{yes}}$ such that $\omega(\Psi, A) = 1$. Let us state the definition of a *gap under an assumption*,

Definition 4 ($(\epsilon, 1)$-Gap Under Assumption) *Let $\Psi = \{\psi_1, ..., \psi_m\}$ be an equation system over variables $\mathcal{V}$. Let $(\mathsf{A}_{\mathtt{yes}}, \mathsf{A}_{\mathtt{no}})$ be an assumption-pair (for $\mathsf{A}_{\mathtt{yes}} \subseteq \mathsf{A}_{\mathtt{no}} \subseteq \mathrm{Assign}(\mathcal{V})$). Ψ is said to have an $(\epsilon, 1)$-gap under $(\mathsf{A}_{\mathtt{yes}}, \mathsf{A}_{\mathtt{no}})$ if exactly one of the following holds:*

Yes: There exists an assignment to $\mathcal{V}$, $A \in \mathsf{A}_{\mathtt{yes}}$, that satisfies all of the ψ_i's, i.e. $\omega(\Psi, A) = 1$.

No: Every assignment to $\mathcal{V}$, $A \in \mathsf{A}_{\mathtt{no}}$, satisfies no more than an ϵ fraction of the ψ_i's, i.e. $\omega(\Psi, A) \leq \epsilon$.

Assumptions can be joined by *intersection*:

$$(\mathsf{A}_{\mathtt{yes}}, \mathsf{A}_{\mathtt{no}}) \cap (\mathsf{A}'_{\mathtt{yes}}, \mathsf{A}'_{\mathtt{no}}) \stackrel{def}{=} (\mathsf{A}_{\mathtt{yes}} \cap \mathsf{A}'_{\mathtt{yes}}, \mathsf{A}_{\mathtt{no}} \cap \mathsf{A}'_{\mathtt{no}})$$

When we say that an equation system has a gap under several assumptions, it is implied that the system has a gap under the intersection of the assumptions.

We use assumptions to separate the proof of theorem 1 into two main components. We are able to reduce the depend of an equation system to constant while maintaining the gap by introducing assumptions on the assignments to the variables of the system. This is done by adding new variables that *supposedly* encode the original ones, and *assuming* consistency between the original and the new variables. We attend to the consistency issue later, when getting rid of the assumptions.

LDF Assumptions

Low-degree functions are used throughout the proof for encoding variables, in a way that enables easy consistency verification.

Definition 5 (Low Degree Function - $[r,d]$-LDF) *Let $\mathcal{F}$ be a field. A function $f:\mathcal{F}^d\to\mathcal{F}$ is said to have degree r, $deg(f)=r$, if its values are the point evaluation of an r degree polynomial on $\mathcal{F}^d$. We say that f is an $[\mathcal{F},r,d]$-LDF if $deg(f)\leq r$.*

Since we use one single field throughout the proof, we abbreviate $[\mathcal{F},r,d]$-LDF to $[r,d]$-LDF (sometimes even LDF). We also use the notation $f\in LDF_{r,d}$.

Given a set of variables $\mathcal{V}$ containing some information, we wish to extend $\mathcal{V}$ by adding auxiliary variables that encode the information in $\mathcal{V}$, in a way that enables checking consistency. We next present an extension method used throughout the paper, called the low-degree-extension.

Definition 6 (Low Degree Extension) *Let $h<|\mathcal{F}|^c$ for some $0<c<\frac{1}{2}$, and define $\mathcal{H}=\{0,1,\ldots,h\}\subseteq\mathcal{F}$. Let $\mathcal{V}$ be an arbitrary set of variables, and take $d\stackrel{def}{=}\log_{h+1}(|\mathcal{V}|)$.*[1] *Identify every point $\boldsymbol{x}\in\mathcal{H}^d$ with a distinct variable, denoted $\mathcal{V}[\boldsymbol{x}]$.*

Extend the domain $\mathcal{H}^d$ to $\mathcal{F}^d$, and add variables for the new points. Denote the entire set of variables $\widehat{\mathcal{V}}$, denoting by $\widehat{\mathcal{V}}[\boldsymbol{x}]$ the variable corresponding to $\boldsymbol{x}\in\mathcal{F}^d$. $\widehat{\mathcal{V}}$ is called the low-degree extension of $\mathcal{V}$ with parameter h.

Given an arbitrary assignment $A:\mathcal{V}\to\mathcal{F}$, we say that the assignment $\widehat{A}:\widehat{\mathcal{V}}\to\mathcal{F}$ is a **proper-extension** *of A of degree r if*

$$\forall\boldsymbol{x}\in\mathcal{H}^d\qquad\widehat{A}(\mathcal{V}[\boldsymbol{x}])=A(\mathcal{V}[\boldsymbol{x}])$$

and if $f:\mathcal{F}^d\to\mathcal{F}$ defined by $\forall\boldsymbol{x}\in\mathcal{F}^d\quad f(\boldsymbol{x})\stackrel{def}{=}\widehat{A}(\mathcal{V}[\boldsymbol{x}])$ is an $[r,d]$-LDF.

Note that for an arbitrary assignment $A:\mathcal{V}\to\mathcal{F}$ there is exactly one proper-extension of degree $r=hd$.

We restrict the assignments to $\widehat{\mathcal{V}}$ to proper-extensions of assignments to $\mathcal{V}$, by using *LDF-assumptions*.

Definition 7 ($[r_{yes},r_{\mathtt{no}},d]$-LDF-Assumption) *Let $\mathcal{F}$ be a field, and assume $r_{yes}\leq r_{\mathtt{no}}\ll|\mathcal{F}|$. Let $\mathcal{D}\stackrel{def}{=}\{v_1,..,v_{|\mathcal{F}|^d}\}$ be a subset of variables, each corresponding to a distinct point in $\mathcal{F}^d$. An $[r_{yes},r_{\mathtt{no}},d]$-assumption $\mathcal{L}_{\mathcal{D}}[r_{yes},r_{\mathtt{no}},d]$ is a pair $(\mathsf{A}_{\mathtt{no}},\mathsf{A}_{\mathtt{yes}})$ such that*

$$\mathsf{A}_{\mathtt{no}}=\{A\mid\ \exists f\in LDF_{r_{\mathtt{no}},d},\ \forall\boldsymbol{x}\in\mathcal{F}^d\ \ A[\mathcal{D}[\boldsymbol{x}]]=f(\boldsymbol{x})\}$$

$$\mathsf{A}_{\mathtt{yes}}=\{A\mid\ \exists f\in LDF_{r_{yes},d},\ \forall\boldsymbol{x}\in\mathcal{F}^d\ \ A[\mathcal{D}[\boldsymbol{x}]]=f(\boldsymbol{x})\}$$

$\mathcal{D}$ is called the domain of the assumption (assumption-domain).

We accompany every low-degree-extension, with an LDF-assumption on its domain.

LDF-assumptions can be joined by intersection, like every assumption. However, we allow joining of LDF-assumptions *only* if their respective domains are *pairwise disjoint*, i.e. a variable is allowed to participate in no more than one LDF-assumption. The reason for this is simple. At a later stage in the proof, we will get rid of the LDF-assumptions by verifying, for every conjunction, that the LDF-assumptions over the conjunction's variables are obeyed. If a variable were to participate in *many* LDF-assumptions, it would over-increase the depend parameter of the final equation system.

[1] w.l.o.g. we assume that every degree, dimension, etc. is a natural number.

A *compound* LDF-assumption $\mathcal{L}_{\mathcal{D}_1,...,\mathcal{D}_s}[r_{yes}, r_{no}, d]$ is an assumption that is the intersection of the LDF-assumptions $\mathcal{L}_{\mathcal{D}_1}[r_{yes}, r_{no}, d]$, ..., $\mathcal{L}_{\mathcal{D}_s}[r_{yes}, r_{no}, d]$; (all with the same r_{yes}, r_{no}, d) i.e. $\mathcal{L}_{\mathcal{D}_1,...,\mathcal{D}_s}[r_{yes}, r_{no}, d] = (\mathsf{A}_{no}, \mathsf{A}_{yes})$ is

$$\begin{aligned} \mathsf{A}_{no} &= \mathsf{A}^1_{no} \cap \mathsf{A}^2_{no} \cap \cdots \cap \mathsf{A}^s_{no} \\ \mathsf{A}_{yes} &= \mathsf{A}^1_{yes} \cap \mathsf{A}^2_{yes} \cap \cdots \cap \mathsf{A}^s_{yes} \end{aligned}$$

where $\mathcal{L}_{\mathcal{D}_i}[r_{yes}, r_{no}, d] = (\mathsf{A}^i_{no}, \mathsf{A}^i_{yes})$ and where $\mathcal{D}_1, .., \mathcal{D}_s \subset \mathcal{V}$ are pairwise-disjoint subsets.

We sometimes omit the list of domains and simply write $\mathcal{L}[r_{yes}, r_{no}, d]$.

The Multiplicative-Assumption

In the proof of theorem 1, we will transform an equation system from quadratic to linear, by replacing quadratic terms with new variables. The multiplicative assumption, defined herein, will then be used to ensure *consistency* of the replacement. Note that the multiplicative assumption must be defined in a way that enables getting rid of it at a later stage.

Definition 8 (Multiplicative-Extension) *Let $h < |\mathcal{F}|^c$ for some $0 < c < \frac{1}{2}$, and let $\mathcal{V}$ be an arbitrary set of variables. Let v_e be a special variable that will represent the value 1, and take $\widehat{\mathcal{V}}$ to be the low-degree-extension of $\mathcal{V} \cup \{v_e\}$, with parameter h (we denote as before, $\mathcal{H} = \{0, .., h\} \subset \mathcal{F}$ and $d \stackrel{def}{=} \log_{h+1}(|\mathcal{V}| + 1)$, so $\left|\widehat{\mathcal{V}}\right| = |\mathcal{F}|^d$). Every point $\boldsymbol{x} \in \mathcal{H}^d$ corresponds to a variable in $\mathcal{V} \cup \{v_e\}$, denoted $\mathcal{V}[\boldsymbol{x}]$, and denote by $e \in \mathcal{H}^d$ the point that corresponds to the variable $v_e = \mathcal{V}[e]$.*

Extend the domain $\mathcal{F}^d$ to $\mathcal{F}^{2d}$ by identifying a point $(\boldsymbol{x}, e) \in \mathcal{F}^{2d}$ with the point $\boldsymbol{x} \in \mathcal{F}^d$, and add variables for the new points. Denote the entire variable-set $\mathcal{V}_\times \supset \widehat{\mathcal{V}}$. $\mathcal{V}_\times$ is called the multiplicative-extension of $\mathcal{V}$, with parameter h.

Given an arbitrary assignment $A : \mathcal{V} \to \mathcal{F}$, we say that the assignment $A_\times : \mathcal{V}_\times \to \mathcal{F}$ is a proper extension of A if

$$\forall \boldsymbol{x} \in \mathcal{F}^{2d} \qquad A_\times(\mathcal{V}_\times[\boldsymbol{x}]) = \boldsymbol{f}_\times(\boldsymbol{x})$$

where $\boldsymbol{f}_\times : \mathcal{F}^{2d} \to \mathcal{F}$ is the unique $[2hd, 2d]$-LDF that obeys

1. $\forall \boldsymbol{x} \in \mathcal{H}^d \quad \boldsymbol{f}_\times(\boldsymbol{x}, e) = A(\mathcal{V}[\boldsymbol{x}])$
2. $\forall \boldsymbol{x}, \boldsymbol{y} \in \mathcal{F}^d \quad \boldsymbol{f}_\times(\boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{f}_\times(\boldsymbol{x}, e) \cdot \boldsymbol{f}_\times(\boldsymbol{y}, e)$.
   Unless A is everywhere zero, this also implies that $\boldsymbol{f}_\times(e, e) = 1$.

We restrict assignments to $\mathcal{V}_\times$ to proper-extensions of assignments to $\mathcal{V}$. This is done by introducing a multiplicative-assumption over the variables.

Definition 9 (Multiplicative-Assumption) *Let $\mathcal{V}_\times$ be a multiplicative extension with parameter h of a variable-set $\mathcal{V}$. A multiplicative-assumption χ over $\mathcal{V}_\times$ is defined by $\chi \stackrel{def}{=} (\mathsf{A}, \mathsf{A})$ where*

$$\mathsf{A} = \{A : \mathcal{V}_\times \to \mathcal{F} : \quad \forall (\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{F}^{2d}$$

$$A(\mathcal{V}_\times[(\boldsymbol{x}, \boldsymbol{y})]) = A(\mathcal{V}_\times[(\boldsymbol{x}, e)]) \cdot A(\mathcal{V}_\times[(\boldsymbol{y}, e)])\}$$

This assumption is always accompanied by an LDF-assumption over $\mathcal{V}_\times$.

1.3 Gap-Maintaining Reduction

Our manipulations of equation systems gradually modify the parameters of the system. Let us introduce a notation for an equation system, that contains all its relevant parameters.

Let $\mathcal{V}$ be a set of variables, let $\mathcal{A}$ be a set of assumptions on the assignments to $\mathcal{V}$, and let $\epsilon > 0$ and $d, D \geq 1$. Denote by $\mathbf{EQ}^d_{\epsilon,D}(\mathcal{A}|\mathcal{V})$, *the set of all degree-d-equation-systems over variables $\mathcal{V}$ with depend D, that have an $(\epsilon, 1)$-gap under the assumptions in $\mathcal{A}$.* Implicit in this notation is n, the size of the system. Throughout the proof we only use $\mathbf{EQ}^d_{\ldots}(\ldots)$ for $d = 1, 2$.

Definition 10 (Gap-Maintaining Reduction) *We say that there is a gap maintaining reduction from a subset $\mathcal{E} \subseteq \mathbf{EQ}^d_{\epsilon,D}(\mathcal{A}|\mathcal{V})$ to $\mathbf{EQ}^{d'}_{\epsilon',D'}(\mathcal{A}'|\mathcal{V}')$ and write $\mathcal{E} \Rightarrow \mathbf{EQ}^{d'}_{\epsilon',D'}(\mathcal{A}'|\mathcal{V}')$ if there is a constant $c > 0$ such that $\epsilon' \leq \epsilon + |\mathcal{F}|^{-c}$ and a general algorithm that on input $\Psi \in \mathcal{E}$ constructs, in polynomial time, an equation system $\Psi' \in \mathbf{EQ}^{d'}_{\epsilon',D'}(\mathcal{A}'|\mathcal{V}')$, such that,*

$$(\omega_{\mathsf{A}_{yes}(\mathcal{A})}(\Psi) = 1) \rightarrow (\omega_{\mathsf{A}_{yes}(\mathcal{A}')}(\Psi') = 1) \tag{1}$$

and

$$(\omega_{\mathsf{A}_{no}(\mathcal{A})}(\Psi) \leq \epsilon) \rightarrow (\omega_{\mathsf{A}_{no}(\mathcal{A}')}(\Psi') \leq \epsilon') \tag{2}$$

In that case we also write $\Psi \Rightarrow \Psi'$ meaning that (1) and (2) hold for Ψ and Ψ'. Note that the size of Ψ' is necessarily polynomial in n.

2 Proof of the Main Theorem

In this section we prove the main theorem, relying on two lemmas (the *sum-check* lemma and the *LDF-reader* lemma). The rest of the paper consists of two additional sections, one for each lemma. Let us restate the main theorem:

Theorem 1 (Main Theorem) *$\forall \delta, c > 0$, $\exists$ $c_1, c_2 > 0$ (constants) such that $QS_n(c_1, F(n), \frac{c_2}{F(n)})$ is NP-hard, for $F(n) = 2^{c(\log n)^{1-\delta}}$.*

Proof: Let $F(n) = 2^{c(\log n)^{1-\delta}}$, and fix for the rest of the proof, a field $\mathcal{F} = \mathbb{Z}_p$, for a prime p, $F(n) \leq p \leq 2F(n)$. We prove NP hardness of $QS_{poly(n)}(O(1), F(n), \frac{O(1)}{F(n)})$ by a reduction from the $QS_n(n, F(n), \frac{O(1)}{F(n)})$ (known to be NP-hard, [HPS93]). Let Ψ_0 be an instance of $QS_n(n, F(n), \frac{O(1)}{F(n)})$ with variables $\mathcal{V}$.

Let us first give a skeleton of the proof and proceed with the necessary details. We construct a sequence of equation systems, $\Psi_0 \Rightarrow \Psi_1 \Rightarrow .. \Rightarrow \Psi_6$. We wish to apply lemma 1 thereby reducing the depend of the equation-system to constant. This lemma works only on linear equations, so we first transform the equations from quadratic (Ψ_0) to linear (Ψ_1). This transformation is done by substituting each multiplicative term with a distinct new variable, and introducing a multiplicative-assumption, over these variables, that ensures each such new variable is the product of the corresponding two variables. We then apply the sum-check lemma (lemma 1), obtaining an equation-system Ψ_2 with a constant depend parameter, and some additional LDF-assumptions.

We now turn to eliminate the assumptions of the equation-system. First we eliminate the multiplicative-assumption by substituting back each variable for its multiplicative-term (Ψ_3). We then use the LDF-reader (lemma 2) to eliminate the LDF-assumptions (Ψ_4). Finally, we amplify the error probability (Ψ_5) and use a simple technique to transform conjunctions of equations back into equations (Ψ_6).

We now proceed with the details of the proof.

Eliminating Multiplicative Terms. We construct a linear equation system Ψ_1 from the quadratic Ψ_0. Let $\mathcal{V}_\times$ be the multiplicative-extension (see definition 8) of $\mathcal{V}$ of degree h, where $h = |\mathcal{F}|^{c'}$ for some $0 < c' < \frac{1}{2}$ (e.g. $c' = \frac{1}{4}$). For every $\psi \in \Psi_0$ we construct ψ' by substituting each quadratic term $\mathcal{V}[\boldsymbol{x}] \cdot \mathcal{V}[\boldsymbol{y}]$ in ψ with the new variable $\mathcal{V}_\times[(\boldsymbol{x}, \boldsymbol{y})]$. We also multiply constants in the equations by the variable $\mathcal{V}_\times[(e, e)]$ in order to make the equations homogeneous. We introduce χ, defined to be the multiplicative-assumption (see definition 9), and an LDF-assumption $\mathcal{L}_{\mathcal{V}_\times}[2hd, 2hd(\log n)^2, 2d]$ where $d \stackrel{def}{=} \log_{h+1}(|\mathcal{V}|+1)$ (the degrees of the LDF-assumption are chosen so that in the 'yes' case, every assignment to $\mathcal{V}_\times$ must be a proper extension of an assignment to $\mathcal{V}$, while in the no case, there is a $(\log n)^2$ factor of additional slackness).

We now have a *homogeneous linear* equation system

$$\Psi_1 \in \mathbf{EQ}^1_{\epsilon_1, |\mathcal{V}_\times|}(\chi, \mathcal{L}_{\mathcal{V}_\times}[2hd, 2hd(\log n)^2, 2d] \mid \mathcal{V}_\times)$$

where $\epsilon_1 = \frac{O(1)}{|\mathcal{F}|}$ and $\Psi_0 \Rightarrow \Psi_1$.

These linear equations are in fact sums of elements with coefficients, such that the sum is supposed to equal zero (note that the Boolean functions of Ψ_0, and hence of Ψ_1, are actually *equations* rather than conjunctions of equations). The next lemma shows how to replace such a sum-check with a system of equations, each depending on a constant number of variables, maintaining the gap under some additional LDF-assumptions.

Lemma 1 (Sum-Check) *Let $0 < c_0 < \frac{1}{2}$ be a constant. There exist constants $c_1, c_2, c_3 > 0$, and a polynomial-time algorithm, that, given an equation-system*

$$\Psi \in \mathbf{EQ}^1_{\epsilon, |\mathcal{V}|}(\chi, \mathcal{L}_{\mathcal{V}}[r_{yes}, r_{\mathtt{no}}, d] \mid \mathcal{V})$$

where χ is an arbitrary assumption over $\mathcal{V}$, $r_{yes} \leq |\mathcal{F}|^{c_0}$ and $r_{\mathtt{no}} = r_{yes} \cdot (\log n)^2$ and where every conjunction $\psi \in \Psi$ is singleton (i.e. is actually an equation*); constructs an equation-system*

$$\widetilde{\Psi} \in \mathbf{EQ}^1_{\widetilde{\epsilon}, c_2}\left(\chi, \mathcal{L}_{\mathcal{D}_1, .., \mathcal{D}_s}[\widetilde{r_{yes}}, r_{\mathtt{no}}, c_3 d]\right)$$

where

- $\Psi \Rightarrow \widetilde{\Psi}$
- $\widetilde{\epsilon} = \epsilon + |\mathcal{F}|^{-c_1}$, *and* $\widetilde{r_{yes}} = 2d(r_{yes} + 1)$.

The proof of this lemma appears in section 4.

The sum-check lemma is applied to Ψ_1 to obtain a new equation system Ψ_2 whose depend parameter is a constant c_2.

Removing the multiplicative-assumption. We will next 'substitute back' the multiplicative terms for the corresponding variables. We discard the multiplicative-assumption χ on $\mathcal{V}_\times$ by replacing, in every conjunction $\psi \in \Psi_2$, every occurrence of a variable $\mathcal{V}_\times[(\boldsymbol{x}, \boldsymbol{y})]$ (for $\boldsymbol{y} \neq e$) with the product $\mathcal{V}_\times[(\boldsymbol{x}, e)] \cdot \mathcal{V}_\times[(\boldsymbol{y}, e)]$. We also add to every conjunction the equation $(\mathcal{V}_\times[(e, e)] = 1)$. This yields a *quadratic* equation system Ψ_3. Define $\widehat{\mathcal{V}} = \{\mathcal{V}_\times[(\boldsymbol{x}, e)] \mid \boldsymbol{x} \in \mathcal{F}^d\}$, renaming $\mathcal{V}_\times[(\boldsymbol{x}, e)]$ by $\widehat{\mathcal{V}}[\boldsymbol{x}]$. The equations in Ψ_3 depend only on variables from $\widehat{\mathcal{V}}$, and the $[2hd, 2hd(\log n)^2, 2d]$-assumption on $\mathcal{V}_\times$ can be exchanged for an LDF-assumption over $\widehat{\mathcal{V}}$. It is easy to see that

$$\Psi_3 \in \mathbf{EQ}^2_{\epsilon_3, 2c_2+1}\left(\mathcal{L}_{\widehat{\mathcal{V}}, \mathcal{D}_1, .., \mathcal{D}_s}[\cdots] \mid \mathcal{V}'\right)$$

where $\epsilon_3 = \epsilon_2$.

Eliminating the LDF-assumptions. We next state the LDF-reader lemma that eliminates all the LDF-assumptions, and yields an equation-system without any assumptions (i.e. with a gap in the conventional sense).

Lemma 2 (LDF-Reader) *Let $c_1, c_2 > 0$ and $0 < c_3, c_4 < 1$ be arbitrary constants, and let $\epsilon = |\mathcal{F}|^{-c_4}$. There exists $g < (\log n)^2$ and constants $c_5, c_6 > 0$ such that,*

$$\mathbf{EQ}^2_{\epsilon,c_1}\left(\{\mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[|\mathcal{F}|^{c_3}, g\cdot|\mathcal{F}|^{c_3}, c_2\cdot\log_{|\mathcal{F}|} n]\}\right)$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon^{c_5},c_6}(\phi)$$

where ϕ denotes the empty set.

We apply this lemma to Ψ_3, but first recall that

$$\Psi_3 \in \mathbf{EQ}^2_{\epsilon,O(1)}\left(\mathcal{L}[|\mathcal{F}|^{c'}, g\cdot|\mathcal{F}|^{c'}, O(\log_{|\mathcal{F}|} n)]\right)$$

for some $0 < c' < 1$, $g = (\log n)^2$, and $\epsilon = |\mathcal{F}|^{-c}$ for some $c > 0$. The above also holds if g is reduced to any value $1 < g' \leq g$, since an $[rg', d]$-LDF is always an $[rg, d]$-LDF, and we have only narrowed the constraints. Taking g' to be the $g < (\log n)^2$ promised by lemma 2, we apply this lemma to Ψ_3 and obtain an equation-system Ψ_4 which is the desired QS instance, except for two small differences: It is a system of conjunctions rather than equations, and its error-probability is $|\mathcal{F}|^{-c}$ rather than $\frac{O(1)}{|\mathcal{F}|}$.

We amplify the error-probability by transforming $\Psi_4 = \{\psi_1, .., \psi_m\}$ into $\Psi_5 = \{(\psi_{i_1}) \wedge \cdots \wedge (\psi_{i_k}) \mid 1 \leq i_1, \ldots, i_k \leq m\}$ The error probability is raised to the power k ($k = \lceil \frac{1}{c} \rceil$ constant), thus obtaining the desired $(\frac{1}{|\mathcal{F}|}, 1)$-gap.

Ψ_5 is a system of *conjunctions* of quadratic equations, rather than a system of quadratic equations. This is solved by the following (last) transformation. We replace every conjunction $\psi = \varphi_1 \wedge \cdots \wedge \varphi_k \in \Psi_5$ with all possible ($|\mathcal{F}|^k$) linear combinations of the equations φ_i:

$$\Psi_\psi \stackrel{def}{=} \{\sum_{i=1}^{k} \alpha_i \cdot \varphi_i \ : \ \alpha_1, .., \alpha_k \in \mathcal{F}\}$$

where a linear combination of equations is defined in the obvious way.

It is easy to see that the system of equations $\Psi_6 = \cup_{\psi\in\Psi_5}\Psi_\psi$ obeys $\Psi_5 \Rightarrow \Psi_6$, since any assignment that did not satisfy a conjunction ψ, cannot satisfy more than $\frac{1}{|\mathcal{F}|}$ of the equations in Ψ_ψ. The error probability increases by no more than $\frac{1}{|\mathcal{F}|}$, hence $\Psi_6 \in QS_{poly(n)}(O(1), F(n), \frac{O(1)}{F(n)})$.

The reduction we have described implies that deciding whether $\omega(\Psi_6) = 1$ or $\omega(\Psi_6) \leq \frac{O(1)}{|\mathcal{F}|}$ enables deciding whether $\omega(\Psi_0) = 1$ or $\omega(\Psi_0) \leq \frac{O(1)}{|\mathcal{F}|}$; hence $QS_n(O(1), F(n), \frac{O(1)}{F(n)})$ is NP-hard. ■

3 Reading LDFs

In this section we prove the following lemma, showing how to eliminate LDF-assumptions, and remain with a gap in the conventional sense. This lemma can be alternatively viewed as a low-degree-test whose error-probability is exponentially small in the number of bits accessed.

Lemma 2 (LDF-Reader) *Let $c_1, c_2 > 0$ and $0 < c_3, c_4 < 1$ be arbitrary constants, and let $\epsilon = |\mathcal{F}|^{-c_4}$. There exists $g < (\log n)^2$ and constants $c_5, c_6 > 0$ such that,*

$$\mathbf{EQ}^2_{\epsilon,c_1}\left(\{\mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[|\mathcal{F}|^{c_3}, g\cdot|\mathcal{F}|^{c_3}, c_2\cdot\log_{|\mathcal{F}|} n]\}\right)$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon^{c_5},c_6}(\phi)$$

where ϕ denotes the empty set.

We prove this lemma via three sub-lemmas. Let us begin by sketching the proof and stating the lemmas. We then proceed to the proof itself.

The main idea is to take Ψ and repeatedly replace its assumptions, until ultimately they can be easily verified (i.e. by accessing a constant number of additional variables). Then, by adding the verification tests to the conjunctions, we eliminate the need for assumptions altogether.

We employ two main techniques for achieving such equation-system reductions. One technique reduces the dimension of the LDFs in the LDF-assumptions, and the other reduces the degree ($r_{\mathtt{no}}$ and r_{yes}). Then, when the LDF-assumptions are of very low degree and dimension, (namely, they are $[O(\log\log n), O(\log\log n), O(1)]$-assumptions) we use a third linearization technique (similar to [ALM+92]) that eliminates the assumptions altogether (by verifying them in the body of the conjunctions).

The first technique, named *cube representation*, uses geometric properties of domains, and their sub-cubes, to replace $[r_{yes}, r_{\mathtt{no}}, d]$-assumptions by $[r_{yes}, r_{\mathtt{no}}, O(1)]$-assumptions hence the dimension of the LDF becomes constant, while the total degree remains the same:

Lemma 3 (Cube-Representation) *Let $r_{yes} \leq |\mathcal{F}|^c$ for some $0 < c < 1$, and $r_{\mathtt{no}} = g \cdot r_{yes}$ where $1 \leq g < (\log n)^2$. Also assume $\epsilon = |\mathcal{F}|^{-c_2}$ for some $0 < c_2 < 1$, $d \leq \log_{|\mathcal{F}|} n$ and $D = O(1)$. There exists $\alpha = |\mathcal{F}|^{-c_3}$ for some $0 < c_3 < 1$, such that,*

$$\mathbf{EQ}^2_{\epsilon,D}(\{\mathcal{L}[r_{yes}, r_{\mathtt{no}}, d]\})$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon+\alpha,D+2K}(\{\mathcal{L}[r_{yes}, r_{\mathtt{no}}, D+3]\})$$

where $K \leq D = O(1)$ bounds, for the first $\mathbf{EQ}_{\ldots}(\cdots)$ system, the number of different assumption domains whose variables participate in a conjunction.

Moreover, there is a reduction between the above, that maintains the bound (K) on the number of different assumption-domains that appear in each conjunction.

The second technique, named *embedding extension*, replaces $[r_{yes}, r_{\mathtt{no}}, O(1)]$ assumptions by $[\widetilde{r_{yes}}, \widetilde{r_{\mathtt{no}}}, d]$-assumptions where $\widetilde{r_{\mathtt{no}}}$ and $\widetilde{r_{yes}}$ are significantly smaller than r_{yes} and $r_{\mathtt{no}}$, while the dimension d is slightly more than constant.

In essence r_{yes} and $r_{\mathtt{no}}$ are of approximately the same size. For technical reasons, the embedding extension 'works' only if $\frac{r_{\mathtt{no}}}{r_{yes}} \geq d$. Every iterative application of the embedding extension technique 'uses up' some of the gap between r_{yes} and $r_{\mathtt{no}}$. For this reason we needed the sum-check lemma and the multiplication lemma to provide a large enough gap $\frac{r_{\mathtt{no}}}{r_{yes}} = (\log n)^{O(1)}$.

Lemma 4 (Embedding Extension) *Let $t = O(1)$. For any $k = O(\log_{|\mathcal{F}|} n)$, if $\frac{r_{\mathtt{no}}}{r_{yes}} \geq kt$ then*

$$\mathbf{EQ}^2_{\epsilon,D}(\{\mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[r_{yes}, r_{\mathtt{no}}, t]\})$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon,D}\left(\{\mathcal{L}_{\widetilde{\mathcal{D}}_1,..,\widetilde{\mathcal{D}}_s}[kt \cdot \sqrt[k]{r_{yes}},\ \frac{r_{\mathtt{no}}}{r_{yes}} \cdot \sqrt[k]{r_{yes}},\ kt]\}\right)$$

Moreover, there is a linear reduction between the above, that maintains the bound on the number of different assumption-domains that appear in each conjunction.

This lemma holds, in fact, for any $t = O(\log_{|\mathcal{F}|} n)$. We state it for $t = O(1)$ since this is how it will be used.

We will apply this lemma for $k \ll r_{yes}$, hence the dominating term in the new degree is $\sqrt[k]{r_{yes}}$. The k parameter determines how fast the degree decreases. Therefore, choosing k as large as possible gives the greatest reduction of the degree, well worth the slight increase in the dimension.

We alternate the use of these two techniques a constant $(\lfloor \frac{1}{\delta} \rfloor + 2)$ number of iterations until we reach an equation system with $[O(\log\log n), O(\log\log n), O(1)]$-assumptions. We then use a linearization technique that by transforming the LDF-assumptions into *linear* assumptions, and then actually checking the linear assumptions by interpolation, eliminates the assumptions altogether:

Lemma 5 (Linearization) *Let $\epsilon = |\mathcal{F}|^{-c}$ for some $c > 0$.*

$$\mathbf{EQ}^2_{\epsilon, O(1)} \left(\{ \mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[O(\log\log n), O(\log\log n), O(1)] \} \right)$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon^{O(1)}, O(1)}(\phi)$$

Having stated the lemmas, we turn to prove the LDF-reader lemma (lemma 2). The proofs of the three techniques (sub-lemmas 3, 4 and 5) follow in the next subsections.

Proof: (**of lemma 2)** The proof begins with an iterative application of lemmas 3 and 4, and ends with one application of lemma 5. We rely on the fact that the reduction relation '$\Rightarrow$' is a transitive relation (as long as it is composed a constant number of times). Let

$$\Psi_0 \in \mathbf{EQ}^2_{\epsilon, O(1)} \left(\{ \mathcal{L}[|\mathcal{F}|^{c_3} \ , \ g \cdot |\mathcal{F}|^{c_3} \ , \ O(\log_{|\mathcal{F}|} n)] \} \right)$$

We shall show a sequence of equation-reductions

$$\Psi_0 \Rightarrow \Psi_1 \Rightarrow \ldots \Rightarrow \Psi_l$$

such that $l = O(\frac{1}{\delta})$ and

$$\Psi_l \in \mathbf{EQ}^2_{\epsilon^{O(1)}, O(1)}(\phi)$$

Note that by combining lemma 3 and then lemma 4 with $k = \log_{|\mathcal{F}|} n \approx (\log n)^{\delta}$, we obtain,

$$\mathbf{EQ}^2_{\epsilon, D}(\{ \mathcal{L}[r_{yes}, r_{\mathtt{no}}, d] \})$$

by lemma 3, $\Downarrow$

$$\mathbf{EQ}^2_{\epsilon+\alpha, D+K}(\{ \mathcal{L}[r_{yes}, r_{\mathtt{no}}, D+3] \})$$

by lemma 4, $\Downarrow$

$$\mathbf{EQ}^2_{\epsilon+\alpha, D+K} \left(\mathcal{L}[k(D+3) \cdot \sqrt[k]{r_{yes}}, \frac{r_{\mathtt{no}}}{r_{yes}} \cdot \sqrt[k]{r_{yes}}, k(D+3)] \right)$$

The first transformation is legal (by lemma 3) as long as $\frac{r_{\mathtt{no}}}{r_{yes}} \leq (\log n)^2$. The second transformation is legal (by lemma 4) as long as $\frac{r_{\mathtt{no}}}{r_{yes}} \geq k(D+3)$. We apply this *double transformation* to Ψ_0 (denoting $r_{yes,0} \stackrel{def}{=} r_{yes}$ and $r_{\mathtt{no},0} \stackrel{def}{=} r_{\mathtt{no}}$) and obtain $\Psi_0 \Rightarrow \Psi_1$; and in the same manner for a general $i > 0$,

$$\Psi_i \in \mathbf{EQ}^2_{\epsilon+i\alpha, D+iK} \left(\mathcal{L}[r_{yes,i}, r_{\mathtt{no},i}, d_i] \right)$$

using the same value for k, and where

$$\begin{aligned} d_i &= (D + K(i-1) + 3)k = O(i \cdot k) \\ r_{\mathtt{no},i} &= \sqrt[k]{r_{yes,i-1}} \cdot \frac{r_{\mathtt{no},i-1}}{r_{yes,i-1}} \\ r_{yes,i} &= \sqrt[k]{r_{yes,i-1}} \cdot d_i \end{aligned}$$

We choose the following values for $g_i \stackrel{def}{=} \frac{r_{\mathtt{no},i}}{r_{yes,i}}$ so as to ensure that it is legal to apply the double transformation. In order to legally apply lemma 4 we must have

$$\frac{r_{\mathtt{no},i-1}}{r_{yes,i-1}} \geq \frac{r_{\mathtt{no},i}}{r_{yes,i}} \cdot d_i$$

In order to legally apply lemma 3 we must have

$$\frac{r_{\mathtt{no},i}}{r_{yes,i}} < (\log n)^2$$

Let $\Delta \stackrel{def}{=} \lfloor \frac{1}{\delta} \rfloor$. We wish to have $\Delta + 2$ iterations of the 'double transformation' with different values for k, all of which obey $k \leq O((\log n)^\delta)$. Hence we choose $g_{\Delta+2} \stackrel{def}{=} \frac{r_{\mathtt{no},\Delta+2}}{r_{yes,\Delta+2}} = 1$, and for $i < \Delta + 2$,

$$\begin{aligned} g_i \stackrel{def}{=} \frac{r_{\mathtt{no},i}}{r_{yes,i}} = \prod_{j=i+1}^{\Delta+2} (d_j) &= \prod_{j=i+1}^{\Delta+2} O(j \cdot k) \\ &\leq (O(\Delta) \cdot (\log n)^\delta)^\Delta \\ &< (\log n)^2 \end{aligned}$$

so it is 'legal' to apply lemma 3 and lemma 4 for all $0 \leq i < \Delta + 2$. The size of $r_{\mathtt{no},i}, r_{yes,i}$, is by induction, $\forall 0 \leq i < \Delta$,

$$2^{\Omega((\log n)^{1-(i+1)\delta})} \leq r_{yes,i}, r_{\mathtt{no},i} \leq 2^{O((\log n)^{1-(i+1)\delta})}$$

and specifically for $i = \Delta$,

$$r_{yes,\Delta} \leq 2^{O((\log n)^\delta)}$$

$r_{yes,\Delta}$ is significantly smaller than $r_{yes,0}$, but not small enough. We will repeat the double transformation two additional times with smaller values for k so as to obtain the desired $O(\log \log n)$ degrees. We now choose $k = \log_2(r_{yes,\Delta})$ which is $\leq O((\log n)^\delta)$ because of the above, making $\sqrt[k]{r_{yes,\Delta}} = 2$, and obtain $\Psi_\Delta \Rightarrow \Psi_{\Delta+1}$, $\Psi_{\Delta+1} \in \mathbf{EQ}^2_{\epsilon_{\Delta+1}, D+(\Delta+1)K}(\mathcal{L}[r_{yes,\Delta+1}, r_{\mathtt{no},\Delta+1}, d_{\Delta+1}])$, with

$$\begin{aligned} \epsilon_{\Delta+1} &= \epsilon_0 + (\Delta+1)\alpha \\ d_{\Delta+1} &= (D + K\Delta + 3) \cdot k \\ &< O((\Delta+1) \cdot (\log n)^\delta) \\ r_{\mathtt{no},\Delta+1} &= \sqrt[k]{r_{yes,\Delta}} \cdot g_\Delta = 2 \cdot g_\Delta, \text{ where } g_\Delta = \frac{r_{\mathtt{no},\Delta}}{r_{yes,\Delta}} \\ r_{yes,\Delta+1} &= \sqrt[k]{r_{yes,\Delta}} \cdot d_{\Delta+1} = 2 \cdot d_{\Delta+1} \end{aligned}$$

Note that at this stage, $r_{yes,\Delta+1}$ is already *logarithmic*, as opposed to the almost polynomial $r_{yes,0}$ we began with in Ψ_0. We need to perform this transformation one more time so that $r_{yes}, r_{\mathtt{no}} = O(\log \log n)$, and we can apply

the linearization technique. We again choose $k = \log_2(r_{yes,\Delta+1}) = O(\delta \log\log n)$, making $\sqrt[k]{r_{yes,\Delta+1}} = 2$, and finally obtain $\Psi_{\Delta+1} \Rightarrow \Psi_{\Delta+2}, \Psi_{\Delta+2} \in \mathbf{EQ}^2_{\epsilon_{\Delta+2},D+(\Delta+2)K}\left(\mathcal{L}[r_{yes,\Delta+2}, r_{\mathtt{no},\Delta+2}, d_{\Delta+2}]\right)$ with

$$\begin{aligned}
\epsilon_{\Delta+2} &= \epsilon_0 + (\Delta+2)\alpha \\
d_{\Delta+2} &= (D + K(\Delta+1) + 3) \cdot k \\
&< O(\Delta+2) \cdot O(\delta \log\log n) = O(\log\log n) \\
r_{\mathtt{no},\Delta+2} &= \sqrt[k]{r_{yes,\Delta+1}} \cdot g_{\Delta+1} = 2 \cdot g_{\Delta+1} \\
r_{yes,\Delta+2} &= \sqrt[k]{r_{yes,\Delta+1}} \cdot d_{\Delta+2} = 2 \cdot d_{\Delta+2}
\end{aligned}$$

We have chosen $g_{\Delta+2} = 1$ and thus $r_{\mathtt{no},\Delta+2} = r_{yes,\Delta+2} = O(\log\log n)$. With a final application of lemma 3 we obtain an equation system $\Psi_{\Delta+2\frac{1}{2}}$ with the same degrees $r_{yes,\Delta+2\frac{1}{2}} = r_{yes,\Delta+2}$, and $r_{\mathtt{no},\Delta+2\frac{1}{2}} = r_{\mathtt{no},\Delta+2}$ but with $d_{\Delta+2\frac{1}{2}} = O(1)$. We are now ready to apply lemma 5. We thus obtain an equation system Ψ_3, $\Psi_{\Delta+2\frac{1}{2}} \Rightarrow \Psi_{\Delta+3}$, where

$$\Psi_{\Delta+3} \in \mathbf{EQ}^2_{\epsilon^*,O(\Delta)}(\phi)$$

for $\epsilon^* = \epsilon_0 + (\Delta+3) \cdot \alpha + \frac{1}{|\mathcal{F}|} = |\mathcal{F}|^{-c^*}$ for some $0 < c^* < 1$. ■

3.1 Cube Representation

In this subsection we show a general algorithm that, given a system with a constant depend and a gap under $[r_{yes}, r_{\mathtt{no}}, d]$-LDF-assumptions, generates a system with LDF-assumptions of the same degree, but of *constant* dimension. The generated system will maintain the gap of the original system. The depend parameter will not be increased by much.

Lemma 3 (Cube-Representation) *Let $r_{yes} \leq |\mathcal{F}|^c$ for some $0 < c < 1$, and $r_{\mathtt{no}} = g \cdot r_{yes}$ where $1 \leq g < (\log n)^2$. Also assume $\epsilon = |\mathcal{F}|^{-c_2}$ for some $0 < c_2 < 1$, $d \leq \log_{|\mathcal{F}|} n$ and $D = O(1)$. There exists $\alpha = |\mathcal{F}|^{-c_3}$ for some $0 < c_3 < 1$, such that,*

$$\mathbf{EQ}^2_{\epsilon,D}(\{\mathcal{L}[r_{yes}, r_{\mathtt{no}}, d]\})$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon+\alpha,D+2K}\left(\{\mathcal{L}[r_{yes}, r_{\mathtt{no}}, D+3]\}\right)$$

where $K \leq D = O(1)$ bounds, for the first $\mathbf{EQ}_{\ldots}(\cdots)$ system, the number of different assumption domains whose variables participate in a conjunction.

Moreover, there is a reduction between the above, that maintains the bound (K) on the number of different assumption-domains that appear in each conjunction.

We begin with an equation system $\Psi \in \mathbf{EQ}^2_{\epsilon,D}(\{\mathcal{L}[r_{yes}, r_{\mathtt{no}}, d]\})$. Ψ has a polynomial number of LDF-assumptions, each over a variable-set corresponding to $\mathcal{F}^d$ (these variable-sets are pairwise disjoint). We would like to substitute these d-dimensional LDF-assumptions with LDF-assumptions of constant dimension, namely $(D+3)$.

Suppose we substitute an assumption $\mathcal{L}$ over a variable-set $\mathcal{D}$, with many LDF-assumptions, one for each $(D+3)$-dimensional affine subspace of $\mathcal{D}$ (we identify each point $\boldsymbol{x} \in \mathcal{F}^d$ with a variable $\mathcal{D}[\boldsymbol{x}] \in \mathcal{D}$ and regard $\mathcal{D}$ both as a variable-set and as a geometric domain). The intersection of these new assumptions is exactly equal to $\mathcal{L}$.

However, the $(D+3)$-dimensional affine subspaces of $\mathcal{D}$ ($(D+3)$-cubes for short) are far from being disjoint, and we cannot have LDF-assumptions over intersecting domains (see explanation following definition 7 in subsection 1.2). We overcome this problem by duplicating the variables: we add a *new* variable-set for each $(D+3)$-dimensional affine subspace of $\mathcal{D}$.

Then, for each $\psi \in \Psi$, we construct a set Ψ'_ψ of conjunctions that simulate ψ by replacing the original variables with duplicates, and then somehow verifying consistency between the duplicates and the original variables.

The Construction of Ψ'

We will describe the construction of an equation system $\Psi' \in \mathbf{EQ}^2_{\epsilon+\alpha,D+2K}(\{\mathcal{L}[r_{yes}, r_{\mathtt{no}}, D+3]\})$ from Ψ, such that $\Psi \Rightarrow \Psi'$.

Variables. Let $\mathcal{D}$ be an assumption-domain in Ψ. For each $(D+3)$-cube $\mathcal{C} \subset \mathcal{D}$, we add a new set $\mathcal{C}[\cdot]$ of cube-variables, with a distinct variable $\mathcal{C}[\boldsymbol{x}]$ for each point $\boldsymbol{x} \in \mathcal{C} \subset \mathcal{D}$. Note that now each point $\boldsymbol{x} \in \mathcal{D}$ has many variables representing it, one for each $(D+3)$-cube containing $\boldsymbol{x}$.

Assumptions. Let $\mathcal{D}$ be an assumption-domain in Ψ. For each cube $\mathcal{C} \subset \mathcal{D}$, Ψ' will have an $[r_{yes}, r_{\mathtt{no}}, D+3]$-assumption on $\mathcal{C}[\,\cdot\,]$. These assumptions replace the $[r_{yes}, r_{\mathtt{no}}, d]$-assumption on $\mathcal{D}$. The assumption replacement can be summarized by

$$\mathcal{L}_{\mathcal{D}}[r_{yes}, r_{\mathtt{no}}, d] \longmapsto \{\mathcal{L}_{\mathcal{C}[\cdot]}[r_{yes}, r_{\mathtt{no}}, D+3] \ : \ \mathcal{C} \text{ is a } (D+3)\text{-cube in } \mathcal{D}\}$$

The Conjunctions. For every $\psi \in \Psi$ we shall have a set of conjunctions Ψ'_ψ, and set

$$\Psi' \stackrel{def}{=} \bigcup_{\psi \in \Psi} \Psi'_\psi$$

Each conjunction in Ψ'_ψ will simulate ψ by referring not to the original variables, but to cube-variables that duplicate them. In addition, the conjunctions in Ψ'_ψ will have consistency equations verifying the consistency of the cube-variables with the original variables.

ψ depends on variables from K assumption-domains[2], denoted $\mathcal{D}_1, \ldots, \mathcal{D}_K$ and called the assumption domains of ψ, with at most D variables from each $\mathcal{D}_i$, since the total depend is D. For each i, take $\mathbf{x}_i \subseteq \mathcal{D}_i$ to be an arbitrary set of D variables, that contains the variables of ψ in $\mathcal{D}_i$. The set $\mathbf{x}_i$ is called the *base* of ψ in $\mathcal{D}_i$.

Definition 11 *Let $\mathcal{D} = \mathcal{F}^d$ and let $\mathbf{x}$ be a subset $\mathbf{x} = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_p\} \subseteq \mathcal{D}$. We denote by $\mathcal{S}_{\mathbf{x}}$ the set of $(D+3)$-cubes in $\mathcal{D}$ that contain $\mathbf{x}$.*

For every choice of points $\boldsymbol{y}_i \in \mathcal{D}$ and cubes $\mathcal{C}_i \in \mathcal{S}_{\mathbf{x}_i \cup \{\boldsymbol{y}_i\}}$, $i = 1, .., K$ (where $\mathbf{x}_i$ is, again, the base of ψ in $\mathcal{D}_i$), we have a conjunction $\psi[\boldsymbol{y}_1, \ldots, \boldsymbol{y}_K, \mathcal{C}_1, \ldots, \mathcal{C}_K] \in \Psi'_\psi$ that reads the variables of ψ from the $\mathcal{C}_i$s.

$$\Psi'_\psi = \{\psi[\boldsymbol{y}_1, \ldots, \boldsymbol{y}_K, \mathcal{C}_1, \ldots, \mathcal{C}_K] \ : \ \forall\ i \ \ \boldsymbol{y}_i \in \mathcal{D}, \ \ \mathcal{C}_i \in \mathcal{S}_{\mathbf{x}_i \cup \{\boldsymbol{y}_i\}}\}$$

where $\psi[\boldsymbol{y}_1, \ldots, \boldsymbol{y}_K, \mathcal{C}_1, \ldots, \mathcal{C}_K]$ is the conjunction of the following:

- The conjunction ψ, where every variable of the form $\mathcal{D}_i[\boldsymbol{x}]$ is substituted by $\mathcal{C}_i[\boldsymbol{x}]$ ($\mathcal{C}_i$ contains the point $\boldsymbol{x}$ since it contains the base of ψ in $\mathcal{D}_i$).
- For each $1 \leq i \leq K$, we have the equation $(\mathcal{C}_i[\boldsymbol{y}_i] = \mathcal{D}_i[\boldsymbol{y}_i])$. These equations are called the *consistency test equations*.

The variables of $\psi[\boldsymbol{y}_1, \ldots, \boldsymbol{y}_K, \mathcal{C}_1, \ldots, \mathcal{C}_K] \in \Psi'$, unlike those of ψ in Ψ, are not guaranteed to have values of global LDFs on the $\mathcal{D}_i$'s. However, we will show that the gap is maintained in Ψ', because in almost all the conjunctions of Ψ'_ψ, *if the consistency tests succeed*, the variables have the values of *permissible* LDFs on the $\mathcal{D}_i$'s. We will elaborate on this when we prove that Ψ' maintains the gap.

[2] ψ may depend on variables from less than K assumption-domains, in which case we associate it with arbitrary domains.

Adjusting the Construction. For technical reasons, we need Ψ'_ψ to have the same size for each $\psi \in \Psi$. We therefore assume that the base $\mathbf{x}_i$ of ψ in an assumption-domain $\mathcal{D}_i$, always spans a D-dimensional subspace in $\mathcal{D}_i$. Otherwise, we add arbitrary points to $\mathbf{x}_i$. It is easy to see that this ensures that Ψ'_ψ always has the same size.

In addition we get rid of unused variables and assumptions. If no conjunction of Ψ' uses variables from a new domain $\mathcal{C}[\,\cdot\,]$, we discard both $\mathcal{C}[\cdot]$ and the LDF-assumption over it.

The Parameters of Ψ'

The assumptions of Ψ' are $[r_{yes}, r_{\mathtt{no}}, D+3]$-LDF assumptions. The depend of each conjunction equals the depend of ψ, plus the depend of the K consistency test equations. The depend is thus bounded above by $D + 2K$.

We therefore have $\Psi' \in \mathbf{EQ}^2_{\epsilon',D+2K}(\{\mathcal{L}[r_{yes}, r_{\mathtt{no}}, D+3]\})$. Ψ' has the properties claimed in the lemma, and we have yet to show that the gap is maintained, $\Psi \Rightarrow \Psi'$.

The Gap

Denote by $(\mathsf{A}_{yes}, \mathsf{A}_{no})$ the intersection of all the assumptions of Ψ, and by $(\mathsf{A}'_{yes}, \mathsf{A}'_{no})$ the intersection of the assumptions of Ψ'.

The Yes Case. Suppose $\omega_{\mathsf{A}_{yes}}(\Psi) = 1$. Then it is easy to construct a satisfying assignment for Ψ'. Let $\mathcal{A} \in \mathsf{A}_{yes}$ be an assignment such that $\omega(\Psi, \mathcal{A}) = 1$. We defi ne an assignment $\mathcal{A}'$ for Ψ', by taking $\mathcal{A}' = \mathcal{A}$ on the variables of Ψ that remain in Ψ'. The assignment to the new variables is defi ned as follows. For a domain $\mathcal{D}$ of an assumption in Ψ, and a $(D+3)$-cube $\mathcal{C} \subset \mathcal{D}$, we defi ne $\mathcal{A}'(\mathcal{C}[\boldsymbol{x}]) \stackrel{def}{=} \mathcal{A}(\mathcal{D}[\boldsymbol{x}])$ for each new variable $\mathcal{C}[\boldsymbol{x}]$. $\mathcal{A}' \in \mathsf{A}'_{yes}$, because the new cubes are assigned restrictions of $[r_{yes}, r_{\mathtt{no}}, d]$-LDFs, i.e. $[r_{yes}, r_{\mathtt{no}}, D+3]$-LDFs. It is clear from the construction that $\omega(\Psi', \mathcal{A}') = 1$.

The No Case. Now suppose $\omega_{\mathsf{A}_{yes}}(\Psi) \leq \epsilon$, and let $\mathcal{A}' \in \mathsf{A}'_{no}$ be an arbitrary assignment for Ψ'. We need to show that $\omega(\Psi', \mathcal{A}') \leq \epsilon' \stackrel{def}{=} \epsilon + \alpha$, for $\alpha = |\mathcal{F}|^{-c_3}$ (we will take $c_3 = \frac{1-c}{25}$, recalling that c is the constant in $r_{yes} \leq |\mathcal{F}|^c$).

We begin by showing that the consistency equations guarantee that the LDFs on the cube-variables $\mathcal{C}$ are related to the values on the original variables $\mathcal{D}$ in a *permissibility* sense. We will then show how to deduce an assignment $\mathcal{A}$ to the original system from the given assignment $\mathcal{A}'$, such that if $\omega(\Psi', \mathcal{A}') > \epsilon'$ then $\omega(\Psi, \mathcal{A}) > \epsilon$, in contradiction.

Definition 12 (Permissibility) *Let $\mathcal{D} = \mathcal{F}^d$ and let $A : \mathcal{D} \to \mathcal{F}$ be an arbitrary function. An LDF $P : \mathcal{D} \to \mathcal{F}$ is called* **ρ-permissible** *with respect to A (or with respect to $\mathcal{D}$, when A is clear from the context), if it agrees with A on at least a ρ-fraction of the points in $\mathcal{D}$.*

A point $\boldsymbol{x} \in \mathcal{D}$ is **$[r, \rho]$-permissible** *(with respect to A), if there exists a ρ-permissible $[r, d]$-LDF on $\mathcal{D}$, that agrees with A on $\boldsymbol{x}$.*

In the discussion that follows, $\rho \stackrel{def}{=} \left(\frac{4d \cdot r_{\mathtt{no}}}{\mathcal{F}}\right)^{\frac{1}{4}}$ and $r_{\mathtt{no}}$ will serve as the implicit permissibility and degree parameters respectively. Given the assignment $\mathcal{A}'$ and an assumption-domain $\mathcal{D}$, we are interested in the relation between $\mathcal{A}'$'s values on the original variables $\mathcal{D}[\cdot]$, and its values on the new cube-variables $\mathcal{C}[\cdot]$. The function $\mathcal{A}'(\mathcal{C}[\cdot])$ is an $[r_{\mathtt{no}}, D+3]$-LDF (this is guaranteed by the LDF-assumption on $\mathcal{C}$), and we would like to test if it is a restriction of some permissible LDF on $\mathcal{D}$, with respect to $\mathcal{A}'|_{\mathcal{D}}$. From here on, we consider permissibility always with respect to $\mathcal{A}'|_{\mathcal{D}}$, the restriction of the assignment $\mathcal{A}'$ to an assumption-domain $\mathcal{D}$ in Ψ.

Let $\psi \in \Psi$, and let $\mathcal{D}_1, .., \mathcal{D}_K$ be the assumption-domains of ψ. Choose K random points $\boldsymbol{y}_i \in_R \mathcal{D}_i$, and then K random cubes $\mathcal{C}_i$ that contain both the random point $\boldsymbol{y}_i$ and the base of ψ in $\mathcal{D}_i$. This is equivalent to choosing

a random conjunction in Ψ'_ψ. The following lemma shows that if at least one of the cubes $\mathcal{C}_1, .., \mathcal{C}_K$ is assigned an impermissible LDF, it is with very low probability that the consistency equations of $\psi[y_1, \ldots, y_K, \mathcal{C}_1, \ldots, \mathcal{C}_K] \in \Psi'_\psi$ evaluate to TRUE.

Lemma 6 (Cube-vs.-Point) *Let $\mathcal{D} = \mathcal{F}^d$ and $f : \mathcal{D} \to \mathcal{F}$. Let $\mathbf{x} = \{x_1, \ldots, x_D\} \subseteq \mathcal{D}$, and suppose we have an $[r, D+3]$-LDF, $F_\mathcal{C} : \mathcal{C} \to \mathcal{F}$, for each cube $\mathcal{C} \in \mathcal{S}_\mathbf{x}$.*

Choose a random point $x \in_R \mathcal{D}$ and a random cube $\mathcal{C} \in_R \mathcal{S}_{\mathbf{x} \cup \{x\}}$. The probability that $f(x) = F_\mathcal{C}(x)$, yet $F_\mathcal{C}$ is not $[r, \rho]$-permissible, is $O(\rho^{\frac{1}{6}})$. This is true for every choice of ρ which satisfies $\rho > 2\sqrt{\frac{rd}{|\mathcal{F}|}}$

The proof of this lemma is a consequence of the low-degree-test of [RS97], and appears in the appendix (section 5). A conjunction $\psi[y_1, \ldots, y_K, \mathcal{C}_1, \ldots, \mathcal{C}_K]$, one of whose cubes is assigned an impermissible LDF, is called an *impermissible conjunction*.

Corollary 1 *For all $\psi \in \Psi$, the fraction of satisfied impermissible conjunctions in Ψ'_ψ, is*

$$\Pr_{\psi[y_1,\ldots,y_K,\mathcal{C}_1,\ldots,\mathcal{C}_K] \in \Psi'_\psi} (\exists i \;\; \mathcal{C}_i \text{ is assigned an impermissible LDF}) < O(K\rho^{\frac{1}{6}})$$

Proof: Let $\mathcal{D}_1, .., \mathcal{D}_K$ be the assumption domains of ψ. Choose a random point $y_1 \in \mathcal{D}_1$, and a random cube $\mathcal{C}_1 \in \mathcal{S}_{\mathbf{x}_1 \cup y_1}$. If $\mathcal{C}_1$ is assigned a restriction of an impermissible LDF, then the consistency equation $(\mathcal{C}_1[y_1] = \mathcal{D}_1[y_1])$ is satisfied with probability $\leq O(\rho^{\frac{1}{6}})$ by the above lemma. Hence the fraction of satisfied equations $\psi[y_1, \ldots, \mathcal{C}_1, \ldots] \in \Psi'_\psi$ such that $\mathcal{C}_1$ is assigned an impermissible LDF is $\leq O(\rho^{\frac{1}{6}})$. We multiply this fraction by K allowing any of the cubes $\mathcal{C}_1, \mathcal{C}_2, .., \mathcal{C}_K$ to be impermissible. ■

We now know that other than an $O(K\rho^{\frac{1}{6}})$ fraction, a satisfied conjunction in Ψ' reads from cube-variables that are assigned an LDF that is a restriction of a permissible LDF (i.e. reads from a permissible cube).

Proposition 7 *The fraction of conjunctions that are satisfied by $\mathcal{A}'$ is bounded by $\epsilon + O(K\rho^{\frac{1}{6}}) + K\rho^2$. (In other words $\omega(\Psi', \mathcal{A}') \leq \epsilon + O(K\rho^{\frac{1}{6}}) + K\rho^2$).*

Proof: We construct an assignment $\mathcal{A} \in \mathsf{A}_{no}$ for Ψ from $\mathcal{A}'$ such that the bound $\omega(\Psi, \mathcal{A}) \leq \epsilon$ implies the above bound on $\omega(\Psi', \mathcal{A}')$.

Let $\mathcal{D}$ be an assumption domain in Ψ. For a permissible LDF $P : \mathcal{D} \to \mathcal{F}$, denote by $\mathrm{Supp}_\mathcal{D}(P)$ the points in $\mathcal{D}$ that support P and do not support any other permissible LDF (these points are called *non-ambiguous*)

$$\mathrm{Supp}_\mathcal{D}(P) \stackrel{def}{=} \{x \in \mathcal{D} \text{ non-ambiguous} \mid \mathcal{A}'(\mathcal{D}[x]) = P(x)\}$$

The assignment $\mathcal{A}$ is constructed via a random procedure. Let $\mathcal{A}$ be identical to $\mathcal{A}'$ outside assumption domains of Ψ. The assumption-domains $\mathcal{D}$ of Ψ are assigned by the following random procedure.

1. For every assumption domain $\mathcal{D}$ in Ψ, choose a random point $y \in_R \mathcal{D}$.
2. If there is a permissible LDF P on $\mathcal{D}$ such that $y \in \mathrm{Supp}_\mathcal{D}(P)$, assign $\mathcal{A}(\mathcal{D}) \stackrel{def}{=} P$. Otherwise assign $\mathcal{A}(\mathcal{D}) \equiv 0$.
3. Choose a random conjunction $\psi \in_R \Psi$. Output TRUE - $\mathcal{A}$ if ψ is satisfied by $\mathcal{A}$.
4. For the assumption domains $\mathcal{D}_1, .., \mathcal{D}_K$ of ψ choose a random $(D+3)$-cube $\mathcal{C}_i \in_R \mathcal{S}_{\mathbf{x}_i \cup \{y_i\}}$, $i = 1, .., K$, where $\mathbf{x}_i$ is the base of ψ in $\mathcal{D}_i$, and $y_i \in \mathcal{D}_i$ is the point randomly chosen in step 1.

   If $\psi[y_1, .., y_K, \mathcal{C}_1, .., \mathcal{C}_K]$ is satisfied by $\mathcal{A}'$, output TRUE - $\mathcal{A}'$.

Denote by T the event 'the procedure outputs TRUE-$\mathcal{A}$ in step 3', and by T' the event 'the procedure outputs TRUE-$\mathcal{A}'$ in step 4'. We claim that these two events are roughly overlapping.

Given the choice of $\mathcal{A}$, $\Pr(T|\mathcal{A}) = \omega(\Psi, \mathcal{A}) \leq \epsilon$, hence $\Pr(T) \leq \epsilon$. On the other hand, noting that the conjunction $\psi[y_1, .., y_K, \mathcal{C}_1, .., \mathcal{C}_K]$ chosen in step 4 is a uniformly random conjunction in Ψ', $\Pr(T') = \omega(\Psi', \mathcal{A}')$ by definition.

We need to bound $\Pr(T' \setminus T)$ in order to establish a bound on $\omega(\Psi', \mathcal{A}') = \Pr(T')$. We will show that when the event $T' \setminus T$ occurs, the conjunction $\psi' \stackrel{def}{=} \psi[y_1, .., y_K, \mathcal{C}_1, .., \mathcal{C}_K]$ selected in step 4 is either impermissible or *ambiguous* (defined below) and we will bound the probability of both of these cases. By corollary 1 we know that the probability that the conjunction $\psi' \stackrel{def}{=} \psi[y_1, .., y_K, \mathcal{C}_1, .., \mathcal{C}_K]$ selected in step 4 is impermissible yet T' occurs (i.e. ψ' is both impermissible and satisfied by $\mathcal{A}'$), is bounded by $O(K\rho^{\frac{1}{6}})$. We can therefore assume ψ' is permissible and focus on the event $\Pr((T' \cap \{\psi' \text{ is permissible}\}) \setminus T)$.

Recall that ψ' is defined as the simulation of ψ where the original variables are replaced by cube-variables, in conjunction with consistency test equations. When the event T' occurs, $\mathcal{A}'$ satisfies ψ and in particular the consistency test equations hold $\mathcal{A}'(\mathcal{D}_i[y_i]) = \mathcal{A}'(\mathcal{C}_i[y_i])$. This means that $\mathcal{A}'(\mathcal{D}_i[y_i])$ supports the permissible LDF assigned to $\mathcal{C}_i$ (Since we assume that ψ' is permissible, the values assigned by $\mathcal{A}'$ to $\mathcal{C}_i$ are permissible). The only case in which ψ can be unsatisfied, while ψ is satisfied, is if for some $1 \leq i \leq K$ in step 2 we defined $\mathcal{A}(\mathcal{D}_i[\cdot])$ to equal a permissible LDF whose restriction to $\mathcal{C}_i$ differs from $\mathcal{A}'(\mathcal{C}_i[\cdot])$. This can happen only if y_i is *ambiguous*, i.e. $\mathcal{A}'(\mathcal{D}[y_i])$ agrees with more than one permissible LDF (in this case we refer to ψ' as ambiguous).

We claim that for most choices of $y_i \in_R \mathcal{D}_i$, there is *at most* one permissible LDF that agrees with $\mathcal{A}'(\mathcal{D}[y_i])$. This is established by showing that there are relatively few permissible LDFs, and follows since every two distinct $[r_{\mathtt{no}}, d]$-LDFs agree on a very small fraction $(\frac{r_{\mathtt{no}}d}{|\mathcal{F}|})$ of their domain.

Claim 8 *There are no more than $2\rho^{-1}$ ρ-permissible LDFs. (As long as $\rho > 2\sqrt{\frac{r_{\mathtt{no}}t}{|\mathcal{F}|}}$)*

This claim appears in the appendix (section 5) as claim 25, along with its simple proof.

Altogether at most $\binom{2\rho^{-1}}{2}\frac{r_{\mathtt{no}}d}{|\mathcal{F}|}$ fraction of the points $y \in \mathcal{D}$, are *ambiguous*, i.e. agree with more than one permissible LDF. Putting $\rho = \left(\frac{4r_{\mathtt{no}}d}{|\mathcal{F}|}\right)^{\frac{1}{4}}$, the fraction of ambiguous points in a domain is bounded by

$$\binom{2\rho^{-1}}{2}\frac{r_{\mathtt{no}}d}{|\mathcal{F}|} \leq 2\sqrt{\frac{|\mathcal{F}|}{4r_{\mathtt{no}}d}} \cdot \frac{r_{\mathtt{no}}d}{|\mathcal{F}|} = \sqrt{\frac{r_{\mathtt{no}}d}{|\mathcal{F}|}} < \rho^2$$

We call a conjunction $\psi[y_1, .., y_K, \mathcal{C}_1, .., \mathcal{C}_K]$, one of whose points y_i is ambiguous, an *ambiguous conjunction*. We have shown that there are no more than $K \cdot \rho^2$ ambiguous conjunctions.

All in all we have bounded the probability $\Pr(T' \setminus T)$ by showing that

$$T' \setminus T \subseteq \{\psi' \text{ is impermissible}\} \cup \{\psi' \text{ is ambiguous}\}$$

Since $\Pr(\psi' \text{ is impermissible}) < O(K\rho^{\frac{1}{6}})$ and $\Pr(\psi' \text{ is ambiguous}) < \rho^2$ we have

$$\Pr(T') \leq \Pr(T) + \Pr(T' \setminus T) \leq \epsilon + O(K\rho^{\frac{1}{6}}) + K\rho^2$$

■

This completes the proof of lemma 3.

3.2 The Embedding Extension

We state and prove lemma 4, showing how to replace $[r_{yes}, r_{\mathtt{no}}, t]$-assumptions with considerably lower-degree assumptions, slightly increasing the dimension.

Lemma 4 (Embedding Extension) *Let $t = O(1)$. For any $k = O(\log_{|\mathcal{F}|} n)$, if $\frac{r_{\mathtt{no}}}{r_{yes}} \geq kt$ then*

$$\mathbf{EQ}^2_{\epsilon,D}(\{\mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[r_{yes}, r_{\mathtt{no}}, t]\})$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon,D}\left(\{\mathcal{L}_{\widetilde{\mathcal{D}}_1,..,\widetilde{\mathcal{D}}_s}[kt \cdot \sqrt[k]{r_{yes}},\ \frac{r_{\mathtt{no}}}{r_{yes}} \cdot \sqrt[k]{r_{yes}},\ kt]\}\right)$$

Moreover, there is a reduction between the above, that maintains the bound on the number of different assumption-domains that appear in each conjunction.

Proof: Let $\Psi \in \mathbf{EQ}^2_{\epsilon,D}(\{\mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[r_{yes}, r_{\mathtt{no}}, t]\})$. We will show a general algorithm to construct $\widetilde{\Psi}$, such that $\widetilde{\Psi} \in \mathbf{EQ}^2_{\epsilon,D}\left(\{\mathcal{L}_{\widetilde{\mathcal{D}}_1,..,\widetilde{\mathcal{D}}_s}[kt \cdot \sqrt[k]{r_{yes}},\ \frac{r_{\mathtt{no}}}{r_{yes}} \cdot \sqrt[k]{r_{yes}},\ kt]\}\right)$. Recall that the $\mathcal{D}_i$s are pairwise disjoint, and $\mathcal{D}_i$ is identified with a copy of $\mathcal{F}^t$. For a point $\boldsymbol{x} \in \mathcal{F}^t$, we denote by $\mathcal{D}_i[\boldsymbol{x}]$ the variable that corresponds to $\boldsymbol{x}$. We extend $\mathcal{F}^t$ to a domain $\mathcal{F}^{tk}$, i.e. we add a set of new variables $\widetilde{\mathcal{D}}_i$ that 'extends' $\mathcal{D}_i$. $\widetilde{\mathcal{D}}_i$ is identified with the extension domain $\mathcal{F}^{tk}$. We will use the following mapping proposition:

Proposition 9 *Let $b = \sqrt[k]{r_{yes}}$. There is an injective map $\mathcal{M} : \mathcal{F}^t \to \mathcal{F}^{tk}$ with the following properties:*

- *For any $[r_{yes}, t]$-LDF, f on $\mathcal{F}^t$, there is a $[btk, tk]$-LDF f_e on $\mathcal{F}^{tk}$ such that*

$$\forall \boldsymbol{x} \in \mathcal{F}^t \ : \quad f(\boldsymbol{x}) = f_e(\mathcal{M}(\boldsymbol{x}))$$

- *$\mathcal{M}$ is a manifold of degree $b^{k-1} = r_{yes}/b$ in $\mathcal{F}^{tk}$.*
  (Hence for any $[\widetilde{r}, tk]$-LDF f_e on $\mathcal{F}^{tk}$, its restriction to $\mathcal{M}$, $f_e \circ \mathcal{M} : \mathcal{F}^t \to \mathcal{F}$ is an $[\widetilde{r} \cdot b^{k-1}, t]$-LDF over $\mathcal{F}^t$.)

Proof: The embedding map $\mathcal{M}$ is defined as follows. A point $\bar{\boldsymbol{x}} = (\boldsymbol{x}_1, .., \boldsymbol{x}_t) \in \mathcal{F}^t$ is mapped to $\bar{y} = \mathcal{M}(\bar{\boldsymbol{x}})$ by replacing each coordinate $\boldsymbol{x}_i$ with k coordinates $y_{i,0}, .., y_{i,k-1}$ such that the following equations hold,

$$\forall i, m \ \ y_{i,m} = \boldsymbol{x}_i^{b^m}$$

In fact, the points of $\mathcal{F}^t$ are mapped to a manifold in $\mathcal{F}^{tk}$ that is defined by the following equations:

$$\begin{aligned} y_{i,1} &= y_{i,0}^b \\ y_{i,2} &= y_{i,1}^b \\ &\vdots \\ y_{i,k-1} &= y_{i,k-2}^b \end{aligned}$$

It is now clear that $\mathcal{M}$ is an injective map and a manifold of degree b^{k-1}. This gives the second statement in the proposition. For the first statement, consider an $[r_{yes}, t]$-LDF f over $\mathcal{F}^t$. We construct the corresponding LDF f_e over $\mathcal{F}^{t \cdot k}$ by replacing the monomials of f, $\boldsymbol{x}_i^j$, with $(y_{i,0})^{b_0}(y_{i,1})^{b_1}...(y_{i,k-1})^{b_{k-1}}$ where $b_0 b_1 ... b_{k-1}$ is the base b representation of j, i.e. $j = \sum_{p=0}^{k-1} b_p \cdot b^p$. For example if $b = 3$, $f(\bar{\boldsymbol{x}}) = \boldsymbol{x}_0 \boldsymbol{x}_1^5 + 3$ becomes $f_e(\bar{y}) = y_{0,0} \cdot (y_{i,0})^2 (y_{i,1})^1 + 3$.

It is easy to see that $\forall \boldsymbol{x},\ \boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{f}_e(\mathcal{M}(\boldsymbol{x}))$ and evaluating $\boldsymbol{f}$ on any point $\boldsymbol{x} \in \mathcal{F}^t$ can be done in polynomial time by evaluating $\boldsymbol{f}_e$ on $\mathcal{M}(\boldsymbol{x}) \in \mathcal{F}^{tk}$. One may think of the embedding as adding redundancy in the representation of a point, in order to make the computation of the LDF 'easier'.

The dimension of $\boldsymbol{f}_e$ is the dimension of $\boldsymbol{f}$ times $\boldsymbol{k}$, yet the degree in each variable is $< \boldsymbol{b}$, hence the total degree is $\boldsymbol{b} \cdot \boldsymbol{tk}$. This completes the proof of the proposition. ■

Constructing $\widetilde{\Psi}$. We extend every $\mathcal{D}_i$ (viewed as a copy of $\mathcal{F}^t$) by adding a new subset $\widetilde{\mathcal{D}}_i$ of variables that corresponds to $\mathcal{F}^{kt}$. This accounts for the restriction $\boldsymbol{k} = O(\frac{1}{t} \cdot \log_{|\mathcal{F}|} n)$, since we want the size of $\widetilde{\Psi}$ to be polynomial in the size of Ψ.

We 'switch' the assumption $\mathcal{L}_{\mathcal{D}_i}[r_{yes}, r_{\mathtt{no}}, t]$ for $\mathcal{L}_{\widetilde{\mathcal{D}}_i}[\widetilde{r_{yes}}, \widetilde{r_{\mathtt{no}}}, kt]$, where

$$\begin{aligned} \widetilde{r_{\mathtt{no}}} &= g \cdot b \\ \widetilde{r_{yes}} &= kt \cdot b \end{aligned}$$

where $b = \sqrt[k]{r_{yes}}$ and $kt \leq g \leq r_{\mathtt{no}}/r_{yes}$ is arbitrary. Every occurrence of the variable $\mathcal{D}_i[\boldsymbol{x}]$ in a conjunction is *replaced* by the new variable $\widetilde{\mathcal{D}}_i[\mathcal{M}(\boldsymbol{x})]$, hence $depend(\widetilde{\Psi}) = depend(\Psi)$.

The variables $\mathcal{D}_i$ no longer appear in any of the conjunctions and are discarded. Note that the variables from $\widetilde{\mathcal{D}}_i$ that appear in the conjunctions are only those on the manifold $\mathcal{M}(\mathcal{D}_i) \subset \widetilde{\mathcal{D}}_i$. The rest of the variables seem to have been artifi cially added, for the sake of the assumption. In fact, there will be many conjunctions that depend on these variables later when we eliminate the assumptions $\mathcal{L}_{\widetilde{\mathcal{D}}_i}[\cdots]$.

Maintaining the Gap. Denote $\mathcal{L}_{\mathcal{D}_1,\ldots,\mathcal{D}_s}[r_{yes}, r_{\mathtt{no}}, t] = (\mathsf{A}_{\mathtt{yes}}, \mathsf{A}_{\mathtt{no}})$, and $\mathcal{L}_{\widetilde{\mathcal{D}}_1,\ldots,\widetilde{\mathcal{D}}_s}[\widetilde{r_{yes}}, \widetilde{r_{\mathtt{no}}}, kt] = (\widetilde{\mathsf{A}}_{\mathtt{yes}}, \widetilde{\mathsf{A}}_{\mathtt{no}})$.

Yes: If $\omega_{\mathsf{A}_{yes}}(\Psi) = 1$, then $\omega_{\widetilde{\mathsf{A}}_{yes}}(\widetilde{\Psi}) = 1$: Let $A \in \mathsf{A}_{\mathtt{yes}}$ be an assignment such that $\omega(\Psi, A) = 1$. $A \in \mathsf{A}_{\mathtt{yes}}$ means that A assigns $\mathcal{D}_i$'s variables values of an $[r_{yes}, t]$-LDF, say $\boldsymbol{f}_i$. It now follows that the assignment $\widetilde{A}$ defi ned by

$$\forall i \forall \boldsymbol{x} \in \mathcal{F}^{kt} \qquad \widetilde{A}(\widetilde{\mathcal{D}}_i[\boldsymbol{x}]) \stackrel{def}{=} (\boldsymbol{f}_i)_e(\boldsymbol{x})$$

where $(\boldsymbol{f}_i)_e$, is as defi ned in proposition 9, and equaling A everywhere else will satisfy $\widetilde{\Psi}$ since $((\boldsymbol{f}_i)_e \circ \mathcal{M})(\boldsymbol{x}) = \boldsymbol{f}_i(\boldsymbol{x})$. The existence of a $[\boldsymbol{kt} \cdot \boldsymbol{b}, \boldsymbol{kt}]$-LDF $(\boldsymbol{f}_i)_e$ is guaranteed by the fi rst half of proposition 9, hence $\widetilde{A} \in \widetilde{\mathsf{A}}_{\mathtt{yes}}$.

No: If $\omega_{\mathsf{A}_{\mathtt{no}}}(\Psi) \leq \epsilon$, then $\omega_{\widetilde{\mathsf{A}}_{\mathtt{no}}}(\widetilde{\Psi}) \leq \epsilon$. Let $\widetilde{A} \in \widetilde{\mathsf{A}}_{\mathtt{no}}$ be an arbitrary assignment. $\widetilde{A} \in \widetilde{\mathsf{A}}_{\mathtt{no}}$ means that A assigns $\widetilde{\mathcal{D}}_i$'s variables values of an $[\boldsymbol{g} \cdot \boldsymbol{b}, \boldsymbol{kt}]$-LDF. The function $\widetilde{A}(\widetilde{\mathcal{D}}_i[\mathcal{M}(\cdot)])$ is an $[\boldsymbol{g} \cdot \boldsymbol{b} \cdot \boldsymbol{b}^{k-1}, t]$-LDF because of proposition 9. Note that since $\boldsymbol{g} \leq r_{\mathtt{no}}/r_{yes}$,

$$\boldsymbol{g} \cdot \boldsymbol{b} \cdot \boldsymbol{b}^{k-1} = \boldsymbol{g} \cdot r_{yes} \leq r_{\mathtt{no}}$$

hence $\widetilde{A}$ defi nes an assignment $A \in \mathsf{A}_{\mathtt{no}}$ for Ψ with $\omega(\widetilde{A}) = \omega(A) \leq \epsilon$ (by the gap-under-assumption property of Ψ).

Note that we are restricted to a polynomial-size construction (in n, the size of the original input). The extension domain, $\mathcal{F}^{kt}$, must have a polynomial number of points. This means that its dimension cannot exceed $\boldsymbol{kt} \leq O(\log_{|\mathcal{F}|} n) = O((\log n)^{\delta})$. ■

3.3 Linearization

In this subsection we prove lemma 5, showing how to eliminate certain LDF-assumptions altogether. Note that $r_{yes} = r_{\mathtt{no}}$, hence for brevity we shall from now on refer to the assumptions as $[r, d]$-assumptions implying that $r_{yes} = r = r_{\mathtt{no}}$.

Lemma 5 (Linearization) *Let $\epsilon = |\mathcal{F}|^{-c}$ for some $c > 0$.*

$$\mathbf{EQ}^2_{\epsilon,O(1)}\left(\{\mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[O(\log\log n), O(\log\log n), O(1)]\}\right)$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon^{O(1)},O(1)}(\phi)$$

Proof: The proof has two steps. First, in proposition 10, we show an embedding technique which is very similar to the one in lemma 4 and which replaces the LDF-assumptions with linear assumptions. We then use a linear interpolation technique that is similar to [ALM+92], and eliminates the linear assumptions.

Proposition 10 *Let $r = O(\log\log n)$, and $c = O(1)$,*

$$\mathbf{EQ}^2_{\epsilon,O(1)}\left(\{\mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[r, r, c]\}\right)$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon,O(1)}\left(\{\mathcal{L}_{\widetilde{\mathcal{D}}_1,..,\widetilde{\mathcal{D}}_s}[1, 1, \binom{r+c}{c}]\}\right)$$

Proof: Let $\Psi \in \mathbf{EQ}^2_{\epsilon,O(1)}\left(\{\mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[r, r, O(1)]\}\right)$. We show a general algorithm that constructs $\widetilde{\Psi} \in \mathbf{EQ}^2_{\epsilon,O(1)}\left(\{\mathcal{L}_{\widetilde{\mathcal{D}}_1,..,\widetilde{\mathcal{D}}}$ by embedding the variables $\mathcal{D}_j$ ($1 \le j \le s$) as a manifold in a larger domain ($\widetilde{\mathcal{D}}_j$); such that all $[r, c]$-LDFs on $\mathcal{D}_j$ are mapped to $[1, \binom{r+c}{c}]$-LDFs on the new domain $\widetilde{\mathcal{D}}_j$.

Let $f : \mathcal{F}^c \to \mathcal{F}$ be an $[r, c]$-LDF. f can be described by an equation of the form

$$\forall \bar{x} \in \mathcal{F}^c \qquad f(\bar{x}) = \sum_{i=1}^{M} \gamma_i m_i(\bar{x})$$

where $M \stackrel{def}{=} \binom{r+c}{c}$, $\gamma_i \in \mathcal{F}$ and the m_i's are all the monomials of degree at most r. (There are exactly $M = \binom{r+c}{c}$ such monomials).

We map f to the $[1, M]$-LDF $f_L : \mathcal{F}^M \to \mathcal{F}$ defined by

$$\forall \bar{y} \in \mathcal{F}^M \qquad f_L(\bar{y}) = \sum_{i=1}^{M} \gamma_i y_i$$

Constructing $\widetilde{\Psi}$. We embed the domain $\mathcal{F}^c$, identified with the variables $\mathcal{D}_j$, in a larger domain $\mathcal{F}^M$. We take a new set of variables, $\widetilde{\mathcal{D}}_j$, which we identify with the points in $\mathcal{F}^M$. $\mathcal{F}^c$ is mapped to a manifold in $\mathcal{F}^M$ by the map $\mathcal{M}$, defined by

$$\forall \bar{x} \in \mathcal{F}^c \qquad \mathcal{M}(\bar{x}) \stackrel{def}{=} (m_1(\bar{x}), m_2(\bar{x}), \ldots, m_M(\bar{x}))$$

Note that for all $\bar{x}$,

$$f(\mathcal{M}(\bar{x})) = f_L(\bar{x})$$

so f_L is, in a sense, an extension of f.

We remain with the same conjunctions of Ψ, replacing every variable $\mathcal{D}_j[x]$ by $\widetilde{\mathcal{D}}_j[\mathcal{M}(x)]$, hence $depend(\Psi) = depend(\widetilde{\Psi})$. The variables $\mathcal{D}_j$ are discarded, and the assumption $\mathcal{L}_{\mathcal{D}_j}[r, r, c]$ is replaced by $\mathcal{L}_{\widetilde{\mathcal{D}}_j}[1, 1, M]$.

Maintaining the Gap. Denote $\mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[r, r, c] = (\mathsf{A}_{yes}, \mathsf{A}_{no})$ and $\mathcal{L}_{\widetilde{\mathcal{D}}_1,..,\widetilde{\mathcal{D}}_s}[1, 1, M] = (\widetilde{\mathsf{A}}_{yes}, \widetilde{\mathsf{A}}_{no})$.

Yes: If $\omega_{\mathsf{A}_{\mathtt{yes}}}(\Psi) = 1$, then $\omega_{\widetilde{\mathsf{A}}_{\mathtt{yes}}}(\widetilde{\Psi}) = 1$: Let $A \in \mathsf{A}_{\mathtt{yes}}$ be an assignment such that $\omega(\Psi, A) = 1$. $A \in \mathsf{A}_{\mathtt{yes}}$ means that the values assigned by A to the variables $\mathcal{D}_j$ make up an $[r, c]$-LDF f_j. It follows that the assignment $\widetilde{A} \in \widetilde{\mathsf{A}}_{\mathtt{yes}}$ defined by assigning the corresponding LDF $(f_j)_L$ to $\widetilde{\mathcal{D}}_j$ has $\omega(\widetilde{A}, \widetilde{\Psi}) = 1$.

No: If $\omega_{\mathsf{A}_{\mathtt{no}}}(\Psi) \leq \epsilon$, then $\omega_{\widetilde{\mathsf{A}}_{\mathtt{no}}}(\widetilde{\Psi}) \leq \epsilon$: Suppose there is an assignment $\widetilde{A} \in \widetilde{\mathsf{A}}_{\mathtt{no}}$ such that $\omega(\widetilde{\Psi}, \widetilde{A}) > \epsilon$, and reach a contradiction. $\widetilde{A} \in \widetilde{\mathsf{A}}_{\mathtt{no}}$ means that the values assigned by $\widetilde{A}$ to the variables $\widetilde{\mathcal{D}}_j$ make up a $[1, M]$-LDF $\widetilde{f}_j$. Define an assignment A as follows: $\forall x \in \mathcal{F}, \ A[\mathcal{D}_j[x]] \stackrel{def}{=} \widetilde{f}_j(\mathcal{M}(x))$. When restricted to $\mathcal{M}(\mathcal{D}_j)$, $\widetilde{f}_j|_{\mathcal{M}(\mathcal{D}_j)}$ is a $[1, M]$-LDF, hence $A \in \mathsf{A}_{\mathtt{no}}$. Obviously $\omega(\Psi, A) = \omega(\widetilde{\Psi}, \widetilde{A}) > \epsilon$, a contradiction.

This completes the proof of proposition 10. ■

Note that the construction in proposition 10 is polynomial, since the number of points in the extension domain, $|\mathcal{F}|^M = \left(2^{\Theta((\log n)^{1-\delta})}\right)^{(O(\log\log n))^c} < n$. Applying the linear embedding sooner (i.e. on assumptions of larger degrees) would have lead to a super-polynomial construction. We now have to show

$$\mathbf{EQ}^2_{\epsilon, O(1)}\left(\{\mathcal{L}_{\widetilde{\mathcal{D}}_1,..,\widetilde{\mathcal{D}}_s}[1, 1, \binom{r+c}{c}]\}\right)$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon^{O(1)}, O(1)}(\phi)$$

By lemma 3, we have

$$\mathbf{EQ}^2_{\epsilon, O(1)}\left(\{\mathcal{L}_{\widetilde{\mathcal{D}}_1,..,\widetilde{\mathcal{D}}_s}[1, 1, \binom{r+c}{c}]\}\right)$$

$$\Downarrow$$

$$\mathbf{EQ}^2_{\epsilon^{O(1)}, O(1)}\left(\{\mathcal{L}_{\widetilde{\mathcal{D}}'_1,..,\widetilde{\mathcal{D}}'_{s'}}[1, 1, O(1)]\}\right)$$

Note that the LDF-assumption is of a *constant* degree and dimension. The final step of eliminating the assumptions altogether is self-evident: *interpolation.* Instead of assuming that an assignment is a linear function on a set of variables $\mathcal{D}_i$ that correspond to a domain $\mathcal{F}^t$ (where $t = O(1)$); we arbitrarily choose $t + 1$ points $x_0, .., x_t \in \mathcal{F}^t$ that span $\mathcal{F}^t$, and interpolate. More elaborately, let $\Psi \in \mathbf{EQ}^2_{\epsilon^{O(1)}, O(1)}\left(\{\mathcal{L}_{\widetilde{\mathcal{D}}'_1,..,\widetilde{\mathcal{D}}'_{m'}}[1, 1, O(1)]\}\right)$. For every $x \in \mathcal{F}^t$,

$$x = \sum_{j=0}^{t} \alpha_j x_j, \quad for\ some\ \alpha_j \in \mathcal{F}$$

since the points $\boldsymbol{x}_0, .., \boldsymbol{x}_t$ span $\mathcal{F}^t$. We construct $\boldsymbol{\Psi}'$ in the following manner. Every occurrence of the variable $\mathcal{D}_i[\boldsymbol{x}]$, is replaced by

$$\sum_{j=0}^{t} \alpha_j \mathcal{D}_i[\boldsymbol{x}_j]$$

The pairwise disjointness of $\mathcal{D}_i$ assures that there is only one way to replace each variable $\mathcal{D}_i[\boldsymbol{x}]$. We narrowed the variable set $\mathcal{V}$ to $\mathcal{V}' \subset \mathcal{V}$, containing $t+1$ arbitrary spanning variables from each domain, plus any other variables that did not participate in the LDF-assumption. We have constructed an equation system $\boldsymbol{\Psi}' \in \mathbf{EQ}^2_{\epsilon^{O(1)}, O(1)}(\phi)$.

Denoting $\mathcal{L}_{\widetilde{\mathcal{D}}'_1,..,\widetilde{\mathcal{D}}'_{m'}}[1,1,O(1)] = (\mathsf{A_{yes}}, \mathsf{A_{no}})$, where $\mathsf{A_{yes}} = \mathsf{A_{no}}$, we have

$$\forall A \in \mathsf{A_{yes}} = \mathsf{A_{no}}, \qquad \omega(\boldsymbol{\Psi}', A|_{\mathcal{V}'}) = \omega(\boldsymbol{\Psi}, A)$$

hence altogether we have that the gap is maintained. ■

This completes the proof of the linearization technique.

4 Sum-Check

In this section we give a proof of lemma 1. We show a general algorithm that takes as input a system of linear equations with polynomial depend, which under an LDF-assumption and an arbitrary assumption χ, has an $(\epsilon, 1)$-gap. The algorithm outputs a linear equation system with a *constant* depend and an $(\epsilon', 1)$-gap under some additional LDF-assumptions.

All equation systems that are discussed in this section are systems of *actual linear equations*. Conjunctions of linear equations are only formed as an intermediate product in the proofs of the sub-lemmas.

Lemma 1 (Sum-Check) *Let $0 < c_0 < \frac{1}{2}$ be a constant. There exist constants $c_1, c_2, c_3 > 0$, and a polynomial-time algorithm, that, given an equation-system*

$$\boldsymbol{\Psi} \in \mathbf{EQ}^1_{\epsilon,|\mathcal{V}|}\left(\chi, \mathcal{L}_{\mathcal{V}}[r_{yes}, r_{\mathtt{no}}, d] \mid \mathcal{V}\right)$$

where χ is an arbitrary assumption over $\mathcal{V}$, $r_{yes} \leq |\mathcal{F}|^{c_0}$ and $r_{\mathtt{no}} = r_{yes} \cdot (\log n)^2$ and where every conjunction $\psi \in \boldsymbol{\Psi}$ is singleton (i.e. is actually an equation*); constructs an equation-system*

$$\widetilde{\boldsymbol{\Psi}} \in \mathbf{EQ}^1_{\widetilde{\epsilon},c_2}\left(\chi, \mathcal{L}_{\mathcal{D}_1,..,\mathcal{D}_s}[\widetilde{r_{yes}}, r_{\mathtt{no}}, c_3 d]\right)$$

where

- $\boldsymbol{\Psi} \Rightarrow \widetilde{\boldsymbol{\Psi}}$
- $\widetilde{\epsilon} = \epsilon + |\mathcal{F}|^{-c_1}$, *and* $\widetilde{r_{yes}} = 2d(r_{yes}+1)$.

4.1 Proof of Lemma 1

The proof will show how to transform an equation system $\boldsymbol{\Psi}$ as described above, to an equation system $\widetilde{\boldsymbol{\Psi}}$ with the desired properties.

We use two techniques iteratively in the construction of $\widetilde{\boldsymbol{\Psi}}$. The *arithmetization technique* from [BFL91], summarized in lemma 13, and the *curve-extension* technique, which is similar to [ALM+92, LS91, FL92], and is formalized in lemma 16. We also use the *linearization technique*, described in lemma 17.

Aside from having different parameters, these techniques have a similar structure. They receive an equation system as input, and substitute each equation in it with a set of new equations, that depend mostly on freshly

added variables. We regard each new set of equations as a distinct equation system, although they are defined over the same variable set and have the same assumptions.

The construction of $\widetilde{\Psi}$ takes the form of a tree of equation systems. At the root of the tree lies Ψ_0, which is a slight modification of the initial equation system, Ψ. We apply the arithmetization technique to Ψ_0, and obtain a set of equation systems, which we place as the offsprings of Ψ_0 in the tree. The next level of the tree is generated by applying the *curve extension* technique to every offspring, and hanging the resulting systems on it.

We continue to add levels to the tree, by alternately applying either the arithmetization or the curve extension techniques to each leaf. The last level of the tree, reached after $O(\frac{1}{\delta})$ iterations, is produced using the linearization technique, which generates equation systems of constant depend.

We denote by Ψ_i the union of all equation systems in the i'th level of the tree. The depend of Ψ_i, D_i, decreases as i increases, until it finally becomes constant in the system, Ψ_b, that corresponds to the bottom level of the tree. We will show that Ψ_0 also maintains the gap property of Ψ, and therefore taking $\widetilde{\Psi} \stackrel{def}{=} \Psi_b$ concludes the proof.

The Structure of this Section. Each technique used in the tree construction, is described in a sub-lemma. These lemmas, apart from some different parameters, share the same (somewhat abstruse) structure. We first present a so called *lemma-template* for the sub-lemmas, that captures their common properties. We then state the parameters that give each of the three sub-lemmas.

After stating the sub-lemmas, we utilize them to construct the equation-system tree. We then define $\widetilde{\Psi}$ as the system corresponding to the bottom level of the tree, and prove that it has the desired properties, thus concluding the proof of lemma 1. The remaining three subsections (4.2, 4.3 and 4.4) are dedicated to the proofs of the three sub-lemmas.

The Lemma Template

The lemma states the existence of an algorithm that takes as input an equation system Ψ of a specific form, and transforms it into an equation system $\widetilde{\Psi}$ with a special structure. $\widetilde{\Psi}$ keeps the variables of Ψ, together with variables, and assignments the for $\widetilde{\Psi}$ are mapped to assignments for Ψ. We state the lemma template, and then elaborate on its statement.

As parameters, the lemma receives functions $\mathcal{C}$, $\widetilde{r_{yes}}$, $\widetilde{D}$, and $\widetilde{d}$, of four variables - $r_{yes}, r_{\mathtt{no}}, D$ and d, which appear in the statement of the lemma.

Lemma-Template $(\mathcal{C}, \widetilde{r_{yes}}, \widetilde{D}, \widetilde{d})$ *Let $c > 0$, and let $r_{yes}, r_{\mathtt{no}}, D, d$ be parameters which satisfy the condition $\mathcal{C}$. There exists a transformation $\mathcal{T}$, and a general algorithm, that, given an equation system*

$$\Psi \in \mathbf{EQ}^1_{1,D}\left(\mathcal{L}_{\mathcal{V}}[r_{yes}, r_{\mathtt{no}}, d], \mathsf{A}_{ext} \mid \mathcal{V} \cup \mathcal{V}_{ext}\right)$$

where

- *Every equation in Ψ depends on $\leq c$ variables from $\mathcal{V}_{ext}$ (i.e. depends mostly on $\mathcal{V}$).*
- *A_{ext} is an assumption over $\mathcal{V}_{ext}$.*

constructs an equation-system

$$\widetilde{\Psi} \in \mathbf{EQ}^1_{1,\widetilde{D}+c+1}\left(\mathcal{L}_{\mathcal{D}}, \mathcal{L}_{\mathcal{V}}[r_{yes}, r_{\mathtt{no}}, d], \mathsf{A}_{ext} \mid \mathcal{D} \cup \mathcal{V} \cup \mathcal{V}_{ext}\right)$$

with the following structure:

- *The new variables are partitioned to variable-sets $\mathcal{D} = \cup_{\psi \in \Psi} \mathcal{D}_\psi$ with an assumption $\mathcal{L}_{\mathcal{D}} = \cap_{\psi \in \Psi} \mathcal{L}_{\mathcal{D}_\psi}[\widetilde{r_{yes}}, r_{\mathtt{no}}, \widetilde{d}]$*

- *$\widetilde{\Psi}$ is partitioned to subsystems $\widetilde{\Psi} = \cup_{\psi\in\Psi}\widetilde{\Psi}_\psi$, where the equations of $\widetilde{\Psi}_\psi$ depend on $\leq c+1$ variables from $\mathcal{V}\cup\mathcal{V}_{ext}$, and their other variables are from $\mathcal{D}_\psi$.*
- *The size of each subsystem $\widetilde{\Psi}_\psi$ is determined only by r_{yes}, $r_{\mathtt{no}}$, and d (and does not depend on Ψ or ψ). This is called the* uniformity property *of the output system.*

Denote the assumptions of Ψ by $(\mathsf{A}_{yes},\mathsf{A}_{no})$, and the assumptions of $\widetilde{\Psi}$ by $(\widetilde{\mathsf{A}}_{yes},\widetilde{\mathsf{A}}_{no})$. The algorithm ensures the following properties of assignments for Ψ and $\widetilde{\Psi}$:

Extension: *Each assignment $\mathcal{A}\in\mathsf{A}_{yes}$ satisfying Ψ, can be extended to an assignment $\widetilde{\mathcal{A}}\in\widetilde{\mathsf{A}}_{yes}$, which satisfies $\widetilde{\Psi}$.*

Restriction: *$\mathcal{T}$ is a transformations of assignments for Ψ with the following property. For an assignment $\widetilde{\mathcal{A}}\in\widetilde{\mathsf{A}}_{no}$ of $\widetilde{\Psi}$, take $\mathcal{A}$ to be the assignment for Ψ given by*

$$\mathcal{A} \stackrel{def}{=} \mathcal{T}(\widetilde{\mathcal{A}}|_{\mathcal{V}\cup\mathcal{V}_{ext}})$$

Then $\mathcal{A}\in\mathsf{A}_{no}$, and $\omega(\Psi,\mathcal{A}) > \omega(\widetilde{\Psi},\widetilde{\mathcal{A}}) - |\mathcal{F}|^{-c_1}$, where c_1 is some constant.

The algorithm described by the lemma is used later, when we construct the equation-system tree, to take an equation system in the tree and generate its offsprings. Given an input system Ψ, it constructs an equation system $\widetilde{\Psi}$, composed of one sub-system, $\widetilde{\Psi}_\psi$, per equation $\psi\in\Psi$ (we use the sub-systems as the offsprings of Ψ).

For the algorithm to be applicable to Ψ, its equations must depend mostly on variables from an LDF-assumption domain, denoted $\mathcal{V}$, with parameters that satisfy the condition $\mathcal{C}$. Apart from that, each equation may depend on at most c of the rest of the variables, grouped in $\mathcal{V}_{ext}$ (we always take c to be a constant). Ψ may have any assumption A_{ext} over $\mathcal{V}_{ext}$.

The systems $\widetilde{\Psi}_\psi$, generated by the algorithm, are all defined over the same variables and assumptions. These include the variables and assumption of Ψ, and an additional variable set, $\mathcal{D}_\psi$, for each equation $\psi\in\Psi$, with an $[\widetilde{r_{yes}}, r_{\mathtt{no}}, \widetilde{d}]$ assumption over it.

The equations of each subsystem $\widetilde{\Psi}_\psi$, depend mostly on variables from $\mathcal{D}_\psi$, called its *principal domain*. The parameters that determine the LDF-assumption over $\mathcal{D}_\psi$ are $\widetilde{r_{yes}}$, called the *principal degree* of $\widetilde{\Psi}_\psi$, and $\widetilde{d}$, called the *principal dimension* (note that the $r_{\mathtt{no}}$ parameter is the same as in Ψ). $\widetilde{D}$, that determines (up to a constant) the depend of $\widetilde{\Psi}_\psi$, is called its *principal depend*. The uniformity property implies that the sizes of all the $\widetilde{\Psi}_\psi$'s are equal.

The assignments for Ψ and $\widetilde{\Psi}$ are closely related. A satisfying assignment for Ψ, if properly extended to the new variables, would also satisfy $\widetilde{\Psi}$. On the other hand, each assignment $\widetilde{\mathcal{A}}$ for $\widetilde{\Psi}$, can be transformed to an assignment $\mathcal{A}$ for Ψ which satisfies almost the same fraction. Apart from the arithmetization lemma, the transformation $\mathcal{T}$ is just the identity, so $\mathcal{A}$ is obtained from $\widetilde{\mathcal{A}}$ by restriction to the variables of Ψ.

Note that the since sizes of all the $\widetilde{\Psi}_\psi$'s are the same, the fraction of equations that $\widetilde{\mathcal{A}}$ satisfies, $\omega(\widetilde{\Psi},\widetilde{\mathcal{A}})$, equals the average $\mathtt{avg}_{\psi\in\widetilde{\Psi}}\omega(\widetilde{\Psi}_\psi,\widetilde{\mathcal{A}})$.

The Parameters of the Sub-Lemmas

The three lemmas that we use in the construction, are instances of the lemma-template with the following parameters.

The Arithmetization Lemma. This lemma, when applied to a system with small principal degree and dimension, will produce systems with small depend. However, the principal degree and dimension of the output systems are not decreased. Therefore an iterative application of the arithmetization lemma alone, would not further reduce the depend. The arithmetization lemma is obtained by setting the parameters of the lemma-template to:

- $\mathcal{C}$ - The conditions on the parameters of the input system are $r_{yes} \leq |\mathcal{F}|^{c_0}$ and $2d(r_{yes}+1)r_{\mathtt{no}} \leq |\mathcal{F}|^{c_1}$, for some constants $0 < c_0 < c_1 < \frac{1}{2}$.
- $\widetilde{r_{yes}} = 2d(r_{yes}+1)$
- $\widetilde{D} = 2d \cdot r_{yes}$
- $\widetilde{d} = d+1$

The Curve Extension Lemma. This lemma is used to reduce the principal degree and dimension of the input system. It is used alternately with the arithmetization lemma. The curve extension lemma is given by setting the parameters to:

- $\mathcal{C}$ - The parameters of the input system should satisfy $r_{yes}, r_{\mathtt{no}}, d, D < |\mathcal{F}|^{c'}$ for some constant $0 < c_1 < \frac{1}{2}$ such that $d \cdot \sqrt[d]{r_{yes}D} < r_{\mathtt{no}}$.
- $\widetilde{r_{yes}} = \min\{d, 2\log_2 D\} \cdot \max\{(r_{yes}D)^{\frac{1}{d}}, 2\}$
- $\widetilde{D} = r_{yes}D$
- $\widetilde{d} = \min\{d, \log_2 D\}$

The Constant Depend (Linearization) Lemma. This lemma drops the depend of the output systems to constant as we desire. However, due to the strong condition on the input system, we must reach very small principal degree and dimension before the lemma is applicable. The constant depend lemma is obtained by setting the parameters of the lemma-template to:

- $\mathcal{C}$ - The parameters of the input system should satisfy $r_{yes}D < O(\log^{\delta} n)$, and $r_{\mathtt{no}}, d < |\mathcal{F}|^{c'}$ for some constant $0 < c' < \frac{1}{2}$.
- $\widetilde{r_{yes}} = 1$
- $\widetilde{D} = 2$
- $\widetilde{d} = r_{yes}D$

The Equation-System Tree Construction

Let Ψ be the input equation system, as described in the sum-check lemma. We now commence with the reduction of Ψ to $\widetilde{\Psi}$.

The construction of $\widetilde{\Psi}$ from Ψ has the form of a tree. At the root of the tree lies Ψ_0 - a slight alteration of Ψ. The levels below the root are added, one at a time, using the lemmas mentioned above.

The Root of the Tree. Ψ has an assumption χ over $\mathcal{V}$, in addition to the LDF-assumption. Therefore the arithmetization lemma cannot be applied to it directly. So we construct a system Ψ_0 , a technical alteration of Ψ to which the arithmetization lemma is applicable, and place it as the root of the tree. The (boring) details of the construction of Ψ_0 are encapsulated in the following claim.

Claim 11 *There exists a polynomial time algorithm that, given Ψ, produces an equation system*

$$\Psi_0 \in \mathbf{EQ}^{1}_{\widetilde{\epsilon},|\mathcal{V}|+1}\left(\mathcal{L}_{\mathcal{V}}[r_{yes}, r_{\mathtt{no}}, d], \mathsf{A}_{ext} \mid \mathcal{V} \cup \mathcal{V}_{ext}\right)$$

such that $\Psi \Rightarrow \widetilde{\Psi}$ and

- $\widetilde{\epsilon} = \epsilon + |\mathcal{F}|^{-c}$ *for some positive constant c (ϵ is the gap parameter of Ψ).*
- A_{ext} *is an assumption over $\mathcal{V}_{ext}$. A_{ext} is composed of an $[r_{yes}, r_{\mathtt{no}}, d]$-LDF assumption over $\mathcal{V}_{ext}$, and of the assumption χ (to be exact, this is a duplicate of χ that is introduced over $\mathcal{V}_{ext}$ instead of over $\mathcal{V}$).*
- *Every equation of Ψ_0 depends on at most one variable outside $\mathcal{V}$.*

The proof proceeds by duplicating the variable set $\mathcal{V}$ of Ψ, and applying the assumption χ over the duplicate, instead of over $\mathcal{V}$. Then the equations are changed to verify consistency between $\mathcal{V}$ and the duplicate. The full proof is deferred to the end of this section.

Having placed the root, we proceed to construct the levels of the tree below the root. First, we describe how one level is added to the tree, using one of the three techniques.

Adding One Level. We go over the systems in the leafs of the tree. To each leaf system Φ, we apply the chosen technique, obtaining a set $\{\widetilde{\Phi}_{\psi}\}_{\psi \in \Phi}$ of equation systems, that share their variables and assumptions. We place these equation systems as the offsprings of Φ in the tree.

Generating the Whole Tree. To generate the tree of equation systems, we start with the tree containing only the root - Ψ_0. We then generate $\frac{2}{\delta}$ new levels[3], by alternately applying the arithmetization and the curve-extension techniques (the arithmetization technique is applied first). Then we apply the arithmetization method once more, and finally, we add the bottom level using the constant depend technique. Recall that for each technique to be applicable, the input system must satisfy a certain condition. We verify that the techniques are applicable throughout the construction, when we compute the parameters of the equation-system tree.

The Parameters of the Equation-System Tree

Note that equation-systems that lie on the same level of the tree, have the same parameters. Let us compute the parameters of the equation-systems in the i'th level of the tree, where i goes from zero - the root, to $b \stackrel{def}{=} \frac{2}{\delta} + 2$ - the bottom level. Let D_i, r_i and d_i denote the principal depend, degree, and dimension respectively, of the systems in the i'th level. Note that the $r_{\mathtt{no}}$ parameter remains the same as it was on the original system thourghout the construction.

[3]We assume $\frac{1}{\delta}$ is whole, without loss of generality.

The Parameters of Level 0. The parameters of Ψ_0, as given in claim 11 are

- $r_0 = |\mathcal{F}|^{c_0} = 2^{O(\log^{1-\delta} n)}$
- $D_0 = O(n)$
- $d_0 = O(\log^{\delta} n)$ (this is obtained from the fact that the size of Ψ_0, and therefore of $\mathcal{V}$, is polynomial in n).

Note that the parameters of Ψ_0 satisfy the condition of the arithmetization lemma, and therefore the construction of the second level is valid. The parameters of levels $1, 2, \ldots, (\frac{2}{\delta} - 3)$ in the tree are given by the following proposition.

Proposition 12 *For $i = 1, \ldots, (\frac{1}{\delta} - 1)$, the parameters of the systems in the $(2i-1)$-level of tree, generated using the arithmetization technique, are*

- $r_{2i-1} = 2^{O(\log^{1-i\delta} n)}$
- $D_{2i-1} = 2^{O(\log^{1-i\delta} n)}$
- $d_{2i-1} = O(\log^{\delta} n)$

and for $i = 1, \ldots, (\frac{1}{\delta} - 2)$, the parameters for the $2i$-level (obtained by the curve extension technique) are

- $r_{2i} = 2^{O(\log^{1-(i+1)\delta} n)}$
- $D_{2i} = 2^{O(\log^{1-i\delta} n)}$
- $d_{2i} = O(\log^{\delta} n)$

Note that the condition of the curve-extension lemma holds for the parameters of the $(2i-1)$ levels, and the condition of the arithmetization lemma holds for the $2i$ levels, and therefore the construction is valid.

Proof: The proposition is obtained by simple induction on the levels. The parameters of Ψ_1, the base of the induction, are calculated by applying the arithmetization lemma to the parameters of Ψ_0, thus obtaining

- $r_1 = 2d_0(r_0 + 1) = 2^{O(\log^{1-\delta} n)}$
- $D_1 = 2d_0 r_0 = 2^{O(\log^{1-\delta} n)}$
- $d_1 = d_0 + 1 = O(\log^{\delta} n)$

which complies with the proposition. The proceeding steps of the induction are obtained by similar calculations, which we omit. ■

The $b-4 = \frac{2}{\delta} - 2$ Level. This level is obtained using the curve-extension technique. We compute the parameters of the previous level by setting $2i - 1 = 2(\frac{1}{\delta} - 1) - 1$ in the above proposition, and then conpute parameters of this level using the curve-extension lemma:

- $r_{b-4} = O(\log^{\delta} n)O(1) = O(\log^{\delta} n)$
- $D_{b-4} = 2^{O(\log^{\delta} n)} \cdot 2^{O(\log^{\delta} n)} = 2^{O(\log^{\delta} n)}$
- $d_{b-4} = O(\log^{\delta} n)$

The $b-3$ Level. This level is obtained using the arithmetization technique. Its parameters are

- $r_{b-3} = O(\log^{2\delta} n)$
- $D_{b-3} = O(\log^{2\delta} n)$
- $d_{b-3} = O(\log^{\delta} n)$

The $b-2$ Level. The parameters of this level, generated using the curve-extension technique, are

- $r_{b-2} = O(\log\log n)$
- $D_{b-2} = O(\log^{4\delta} n)$
- $d_{b-2} = O(\log\log n)$

The $b-1$ Level. The parameters of this level, the last level before applying the constant-depend technique, are

- $r_{b-1} = O(\log\log^2 n)$
- $D_{b-1} = O(\log\log^2 n)$
- $d_{b-1} = O(\log\log n)$

The b Level. Finally we have systems that satisfy the condition of the constant-depend lemma. We apply the lemma and get

- $r_b = 1$
- $D_b = 2$
- $d_b = O(\log\log^4 n)$

The system $\widetilde{\Psi}$. We are now ready to define $\widetilde{\Psi}$. Let $\widetilde{\Psi}$ be the union of all systems in the bottom (b) level of the tree.

Maintaining the Gap

We need to verify that $\widetilde{\Psi}$ maintains the gap of Ψ_0 (which, in turn, maintains the gap of Ψ, as shown in claim 11). Denote the assumptions of Ψ_0 by $(\mathsf{A}_{yes}, \mathsf{A}_{no})$, and the assumptions of $\widetilde{\Psi}$ by $(\widetilde{\mathsf{A}}_{yes}, \widetilde{\mathsf{A}}_{no})$.

We define intermediate stages between Ψ_0 and $\widetilde{\Psi}$. For $i = 1, \ldots, b$, let Ψ_i be the union of equation-systems in the i'th level of the tree (then $\widetilde{\Psi}$ is just Ψ_b).

Yes: Let $\mathcal{A}$ be a satisfying legal assignment for Ψ_0. We will use the extension property iteratively to extend $\mathcal{A}$ to the variables of systems in deeper levels of the tree.

$\mathcal{A}$ can be extended to an assignment $\mathcal{A}_1$, satisfying Ψ_1, and therefore all of Ψ_1's components, which inhabit level 1 in the tree.. This is by the extension property of the arithmetization technique, that was used to construct level 1.

In a similar way, we extend the assignment of each system in level 1, to a satisfying assignment for its offsprings. We may regard the different assignment extensions, of systems from level 1, as one assignment which satisfies system 2 (since there is no contradiction between the different extensions).

By induction, we continue to extend the assignment down the tree, until we obtain a satisfying assignment for level b, that is, for $\widetilde{\Psi}$.

No: Let us give a randomized interpretation of the restriction property (see the properties of the output system in the lemma-template). Let Φ be an equation system, and let $\widetilde{\Phi}$ be a system obtained from Φ using one of the three techniques. Fix an assignment $\widetilde{\mathcal{A}}$ for $\widetilde{\Phi}$, and let $\mathcal{A}$ be the assignment for Φ that is obtained from the restriction property. If we pick random equations $\psi \in \Phi$ and $\widetilde{\psi} \in \widetilde{\Phi}$, the restriction property implies that

$$\Pr\left[\mathcal{A} \text{ satisfies } \psi\right] > \Pr\left[\widetilde{\mathcal{A}} \text{ satisfies } \widetilde{\psi}\right] - |\mathcal{F}|^{-c} \qquad (*)$$

where c is some constant (without loss of generality, we assume all the lemmas have the same constant).

Now fix an assignment $\mathcal{A}_b$ for Ψ_b (which is, in fact, $\widetilde{\Psi}$). We will show that $\omega(\Psi_b, \mathcal{A}_b)$ is small. First, we define an assignment $\mathcal{A}_i$ for each Ψ_i, using the restriction property iteratively.

$\mathcal{A}_b$ is already defined. Having defined $\mathcal{A}_{i+1}$, we define $\mathcal{A}_i$ as follows. Consider an equation system Φ on the i'th level of the tree (recall that Ψ_i is the union of all systems in the i's level of the tree). The offsprings of Φ in the tree are obtained by applying one of the three techniques on Φ. Since $\mathcal{A}_{i+1}$ assigns, among others, the variables of the offsprings of Φ in the tree, we can use the restriction property on $\mathcal{A}_{i+1}$ to get an assignment for Φ.

We obtain, in a similar way, assignments for all the systems in the i'th level of the tree. We unite these assignments to an assignment for all of Ψ_i - by observing the definition of the restriction property, it is easy to verify that no contradictions exist between the assignments for different systems in the i's level.

Consider the following random process. Denote Ψ_0 by Φ_0 and pick a random equation $\varphi_0 \in \Phi_0$. Then pick a random offspring Φ_1 of Φ_0 (not necessarily the one associated with φ_0), and a random equation $\varphi_1 \in \Phi_1$. Continue to pick a random offspring of Φ_1, denoted Φ_2, and a random equation $\varphi_2 \in \Phi_2$. We proceed by induction until we have picked a path in the tree, $\Phi_0, \ldots, \Phi_b$, and the equations $\varphi_0, \ldots, \varphi_b$.

φ_i is a random equation in Φi, for $i = 0, 1, \ldots, b$. Note that as a consequence of the uniformity property (see the lemma-template), φ_i is also chosen randomly from the set of all equations of the offspring of $\Phi i - 1$, for $i = 1, 2, \ldots, b$. Another consequence of the uniformity property is that φ_b is a random equation in Ψ_b. Putting these observations together with an iterative use of equation $(*)$, gives

$$\begin{aligned}\omega(\Psi_0, \mathcal{A}_0) &= \Pr\left[\mathcal{A}_0 \text{ satisfies } \varphi_0\right] > \Pr\left[\mathcal{A}_1 \text{ satisfies } \varphi_1\right] - |\mathcal{F}|^{-c} > \\ &> \Pr\left[\mathcal{A}_2 \text{ satisfies } \varphi_2\right] - 2|\mathcal{F}|^{-c} > \cdots > \\ &> \Pr\left[\mathcal{A}_b \text{ satisfies } \varphi_b\right] - b|\mathcal{F}|^{-c} = \omega(\Psi_b, \mathcal{A}_b) - b|\mathcal{F}|^{-c}\end{aligned}$$

Suppose $\omega(\Psi) < \epsilon$. The above inequality implies that $\omega(\widetilde{\Psi}) < \epsilon + O(|\mathcal{F}|^{-c})$, thus implying the no case.

We modify $\widetilde{\Psi}$ slightly so it has the required parameters. The changes we make will not damage the gap property of $\widetilde{\Psi}$. The depend of $\widetilde{\Psi}$ is constant, by the above calculation. The $r_{\mathtt{no}}$ parameter of the LDF-assumptions of $\widetilde{\Psi}$ is the same as that of Ψ, as required. The r_{yes} parameter is not the same for all the assumptions we constructed, but from the parameter calculation it is clear that the variable sets added in level 1 of the tree, have the biggest r_{yes} parameter - $2d_0(r_0 + 1) = 2d(r_{yes} + 1)$. We therefore change the LDF assumptions of $\widetilde{\Psi}$, setting the r_{yes} parameter to $2d(r_{yes} + 1)$ on all of them.

The dimension of the LDF assumptions of $\widetilde{\Psi}$ is always $\leq O(\log^{\delta} n) = O(d)$, but it varies. So we add dummy variables to the assumption-domains of Ψ, such that the depend of all assumption-domains is the same, and modify the LDF assumptions accordingly. The parameters of $\widetilde{\Psi}$ are now as the lemma requires.

There is one more thing left to handle. The sum-check lemma requires the assumption χ to be over $\mathcal{V}$. Instead, claim 11 introduces a duplicate of χ over a variable-set that is a duplicate of $\mathcal{V}$. This is of course only a technical problem, that can be fixed by switching names between $\mathcal{V}$ and its duplicate.

Proof of claim 11

We have yet to prove claim 11, that is used to make Ψ suitable for application of the arithmetization lemma. The proof proceeds by duplicating the variable set $\mathcal{V}$ of Ψ, and applying the assumption χ over the duplicate, instead of over $\mathcal{V}$. Then the equations are changed to verify consistency between $\mathcal{V}$ and the duplicate.

The Variables and Assumptions of Ψ_0. Ψ_0 keeps $\mathcal{V}$ and its LDF-assumption. In addition, it has a duplicate $\mathcal{V}_{ext}$ of $\mathcal{V}$, with the same LDF-assumption as of $\mathcal{V}$ over it. Instead of the assumption χ, we employ the *same* assumption, only over $\mathcal{V}_{ext}$. The intersection of the latter assumption and the LDF-assumption over $\mathcal{V}_{ext}$ is denoted by A_{ext}.

The Equations of Ψ_0. For each equation $\psi \in \Psi$, we construct a set Ψ_ψ of $|\mathcal{V}|$ conjunctions as follows: For every point $\boldsymbol{x} \in \mathcal{F}^d$, Ψ_ψ will have the conjunction $\psi \wedge (\mathcal{V}[\boldsymbol{x}] = \mathcal{V}_{ext}[\boldsymbol{x}])$, composed of the equation ψ, and a consistency test. We set $\Psi_0 \stackrel{def}{=} \cup_{\psi\in\Psi} \Psi_\psi$ (actually we transform the conjunctions into equations using proposition 28, but we prefer to think about conjunctions when we discuss the gap).

Let us show that the gap is maintained. Denote the assumptions of Ψ by $(\mathsf{A}_{yes}, \mathsf{A}_{no})$, and the assumptions of Ψ_0 by $(\widetilde{\mathsf{A}}_{yes}, \widetilde{\mathsf{A}}_{no})$.

Yes: Assume $\mathcal{A} \in \mathsf{A}_{yes}$ satisfied $\omega(\Psi, \mathcal{A}) = 1$. We extend $\mathcal{A}$ to an assignment $\widetilde{\mathcal{A}}$ for Ψ_0, by setting the assignment of $\mathcal{V}_{ext}$ to be the same as of $\mathcal{V}$. Obviously, $\widetilde{\mathcal{A}} \in \widetilde{\mathsf{A}}_{yes}$, and $\omega(\Psi_0, \widetilde{\mathcal{A}}) = 1$.

No: Suppose $\omega_{\mathsf{A}_{no}}(\Psi) < \epsilon$, and let $\widetilde{\mathcal{A}} \in \mathsf{A}_{no}$ be an assignment for Ψ_0. We will show that $\omega(\Psi_0, \widetilde{\mathcal{A}})$ is small.

We consider two cases. One, is that $\widetilde{\mathcal{A}}$ gives $\mathcal{V}$ and $\mathcal{V}_{ext}$ the same assignments. In this case, taking $\mathcal{A} \stackrel{def}{=} \widetilde{\mathcal{A}}|_{\mathcal{V}}$ gives an assignment $\mathcal{A} \in \mathsf{A}_{no}$, which obviously satisfies $\omega(\Psi_0, \widetilde{\mathcal{A}}) = \omega(\Psi, \mathcal{A}) < \epsilon$.

If, on the other hand, $\widetilde{\mathcal{A}}$ assigns $\mathcal{V}$ and $\mathcal{V}_{ext}$ different LDFs, then these LDFs are equal on at most an $\frac{r_{\mathtt{no}}d}{|\mathcal{F}|}$ of the points. Therefore out of each Ψ_ψ, at most an $\frac{r_{\mathtt{no}}d}{|\mathcal{F}|}$ fraction of the conjunctions are satisfied, because of the consistency equations, and thus $\omega(\Psi_0, \widetilde{\mathcal{A}}) < \frac{r_{\mathtt{no}}d}{|\mathcal{F}|}$.

In both cases, if we take c_1 such that $\frac{r_{\mathtt{no}}d}{|\mathcal{F}|} < |\mathcal{F}|^{-c_1}$, we get $\omega(\Psi_0, \widetilde{\mathcal{A}}) < \epsilon + \frac{r_{\mathtt{no}}d}{|\mathcal{F}|}$, and conclude the proof of the claim.

■

4.2 The Arithmetization Lemma

In essence, the following arithmetization lemma states the ψsum-check technique, well-known from previous PCP proofs (see [BFL91]). This technique reduces the depend of the system considerably, while maintaining the gap.

Lemma 13 (Arithmetization) *The arithmetization lemma is obtained from the lemma-template by setting its parameters to*

$\mathcal{C}$ *- The conditions on the parameters of the input system are* $r_{yes} \leq |\mathcal{F}|^{c_0}$ *and* $r_{\mathtt{no}} = |\mathcal{F}|^{c_1}$*, for some constants* $0 < c_0 < c_1 < \frac{1}{2}$.

$\widetilde{r_{yes}} = 2d(r_{yes} + 1)$

$\widetilde{D} = 2d \cdot r_{yes}$

$\widetilde{d} = d + 1$

Proof: The proof follows from the arithmetization technique of [BFL91]. We start with an equation system Ψ as in the claim, and construct $\widetilde{\Psi}$.

Constructing $\widetilde{\Psi}$

The variables. For each $\psi \in \Psi$, we shall add a *new* set of $|\mathcal{F}|^{d+1}$ variables, $\mathcal{D}_\psi$, to represent the *sum-check-LDF* (defined below) related to ψ. We also introduce an LDF-assumption $\mathcal{L}_{\mathcal{D}_\psi}[\widetilde{r_{yes}}, \widetilde{r_{\mathtt{no}}}, d+1]$ over $\mathcal{D}_\psi$. Let $\mathcal{D} \stackrel{def}{=} \bigcup_{\psi\in\Psi} \mathcal{D}_\psi$. The variable-set for the new equation system $\widetilde{\Psi}$ is $\mathcal{D} \cup \mathcal{V} \cup \mathcal{V}_{ext}$.

The Equations. We proceed to define the equations of $\widetilde{\Psi}$. For each $\psi \in \Psi$, we construct a set of equations $\widetilde{\Psi}_\psi$, such that $\widetilde{\Psi} \stackrel{def}{=} \cup_{\psi\in\Psi} \widetilde{\Psi}_{eq}$.
Let $\psi \in \Psi$ be an equation. ψ is of the form

$$\sum_{x\in\mathcal{F}^d} \kappa_\psi(x)\mathcal{V}[x] = s_\psi$$

where $\kappa_\psi : \mathcal{F}^d \to \mathcal{F}$ is the coefficient function of ψ, and $s_\psi \stackrel{def}{=} \sum_{j=1}^{c} \beta_j v_j$, is a linear combination of variables outside $\mathcal{V}$. The following claim will be used to change κ_ψ so it is supported by $\mathcal{H}^d \subseteq \mathcal{F}^d$, where $\mathcal{H} = \{0, 1, \ldots, r_{yes}\}$.

Claim 14 *There exists a function $\kappa'_\psi; \mathcal{H}^d \to \mathcal{F}^d$ such that all LDFs $p : \mathcal{F}^d \to \mathcal{F}$ of degree r_{yes} in each variable, satisfy*

$$\sum_{x\in\mathcal{H}^d} \kappa'_\psi(x)p(x) = \sum_{x\in\mathcal{F}^d} \kappa_\psi(x)p(x)$$

Proof: This claim follows easily, if we recall that for each $x \in \mathcal{F}^d$, the value at x of an LDF p as above, can be interpolated by a linear combination of its values on $\mathcal{H}^d$ (where the coefficients of the combination are independent of p). Since $\sum_{x\in\mathcal{F}^d} \kappa_\psi(x)p(x)$ is a linear combination of the values $\{p(x)\}_{x\in\mathcal{F}^d}$, it also can be interpolated. ■

Remark 1 *Note that in case $\mathcal{V}$ is assigned an LDF of degree r_{yes} in each variable, then ψ is satisfied iff it is still satisfied when we substitute κ_ψ with κ'_ψ (and pad it with zeros where it is undefined).*

We define $\widehat{\kappa}_\psi$ to be the proper extension of κ'_ψ (see definition8). Then $\widehat{\kappa}_\psi$ is an $[r_{yes}d, d]$-LDF. $\widehat{\kappa}_\psi$ is used to define the set of *path checks* of ψ, which are then used to form $\widetilde{\Psi}_\psi$.

Definition 13 (Path Check) *Let $\psi \in \Psi$ be an equation, and let $x = (x_1, \ldots, x_d) \in \mathcal{F}^d$. The path-check $path[\psi, x]$ is defined as the conjunction of the following equations.*

- *The top equation -*

$$\sum_{i=0}^{r_{yes}} \mathcal{D}_\psi[(1, i, \bar{0})] = \sum_{j=1}^{c} \beta_j v_j$$

- *$d-1$ path equations, for $k = 1, 2, \ldots, (d-1)$ -*

$$\mathcal{D}_\psi[(k, x_1, \ldots, x_k, \bar{0})] = \sum_{i=0}^{r_{yes}} \mathcal{D}_\psi[(k+1, x_1, \ldots, x_k, i, \bar{0})]$$

- *The base equation -*

$$\mathcal{D}_\psi[(d, x_1, \ldots, x_d)] = \widehat{\kappa}_\psi(x_1, \ldots, x_d)\mathcal{V}[(x_1, \ldots, x_d)]$$

We define $\widetilde{\Psi}_\psi = \{path[\psi, x] \mid x \in \mathcal{F}^d\}$. It is easy to see that $depend(\widetilde{\Psi}_\psi) < c + 2r_{yes}d$, and that every equation $\psi \in \widetilde{\Psi}_\psi$ depends on c variables outside $\mathcal{D}_\psi$, namely $v_1, .., v_c$. It is also obvious that the sizes of all $\widetilde{\Psi}_\psi$'s are equal. We define $\widetilde{\Psi} = \cup_{\psi\in\Psi}\widetilde{\Psi}_{eq}$.

The Extension and Restriction Properties

Denote the assumptions of Ψ by $(\mathsf{A}_{yes}, \mathsf{A}_{no})$, and the assumptions of $\widetilde{\Psi}$ by $(\widetilde{\mathsf{A}}_{yes}, \widetilde{\mathsf{A}}_{no})$.

Extension: Suppose $\omega(\Psi) = 1$, i.e. Ψ can be satisfied by an $\mathcal{A} : \mathcal{V} \cup \mathcal{V}_{ext} \to \mathcal{F}$, $\mathcal{A} \in \mathsf{A}_{yes}$. We shall extend $\mathcal{A}$ to a legal satisfying assignment $\widetilde{\mathcal{A}} : \mathcal{D} \cup \mathcal{V} \cup \mathcal{V}_{ext} \to \mathcal{F}$ for $\widetilde{\Psi}$. We first define the sum-check tree, that is used to construct $\widetilde{\mathcal{A}}$. The reader is referred to [BFL91] for a more detailed explanation of this object.

Definition 14 (The Sum-Check tree) *Let $f : \mathcal{F}^d \to \mathcal{F}$ be an $[r, d]$-LDF, and let $\mathcal{H} \stackrel{def}{=} \{1, 2, \ldots, h\}$ for some parameter $h < |\mathcal{F}|$. For $k = 1, 2, \ldots, d$, we define an $[r, k]$-LDF $g_k : \mathcal{F}^k \to \mathcal{F}$ by*

$$\forall x \in \mathcal{F}^k \quad g_k(x) \stackrel{def}{=} \sum_{y\in\mathcal{H}^{(d-k)}} f(x, y)$$

The sum-check tree *of f with parameter h is the vector of polynomials $(g_1, g_2, \ldots, g_d)$. We also define the* sum-check LDF *of f with parameter h, as the $[d(r+1), d+1]$-LDF $g : \mathcal{F}^{d+1} \to \mathcal{F}$ defined by*

$$g(x_0, x_1, .., x_d) \stackrel{def}{=} \sum_{k=1}^{d}\left(\prod_{i=1, i\neq k}^{d}\frac{x_0 - i}{k - i}\right)\cdot g_k(x_1, .., x_k)$$

The sum-check tree is embeded in the sum-check LDF by the identity

$$g(k, x_1, \ldots, x_d) = g_k(x_1, \ldots, x_k) \quad k = 1, 2, \ldots, d$$

We define $\widetilde{\mathcal{A}} : \mathcal{D} \cup \mathcal{V} \cup \mathcal{V}_{ext} \to \mathcal{F}$ to equal $\mathcal{A}$ on $\mathcal{V} \cup \mathcal{V}_{ext}$, and on each $\mathcal{D}_\psi$ we define $\widetilde{\mathcal{A}}$ to be the sum-check LDF of $\widehat{\kappa}_\psi \cdot \mathcal{A}|_\mathcal{V}$ with parameter r_{yes}. It is clear that $\widetilde{\mathcal{A}} \in \widetilde{\mathsf{A}}_{yes}$ (because the sum-check LDF is a $[2d(r_{yes}+1), d+1]$-LDF). The path and base equations in each path check, are obsviously satisfied by $\widetilde{\mathcal{A}}$, as is directly implied from the definition of the sum-check tree, so we only need to check that the top equations are satified.

Focus on one equation $\psi \in \Psi$, and let $(g_1, \ldots, g_d)$ be the sum-check tree of $\widehat{\kappa}_\psi \cdot \mathcal{A}|_\mathcal{V}$ with paramter r_{yes}. The top equation in all the path-checks of $\widetilde{\Psi}_\psi$, is

$$\sum_{i=0}^{r_{yes}}\mathcal{D}_\psi[(1, i, \bar{0})] = s_\psi$$

$\mathcal{D}_\psi[(1, \cdot, \bar{0})]$ is assigned g_1, so the top equation is satisfied iff

$$\sum_{i=0}^{r_{yes}} g_1(i) = s_\psi$$

By the definition of the sum-check tree,

$$\sum_{i=0}^{r_{yes}} g_1(i) = \sum_{i=0}^{r_{yes}} \sum_{y \in \mathcal{H}^{d-1}} \widehat{\kappa}(i,y)\mathcal{A}(\mathcal{V}[i,y]) = \sum_{x \in \mathcal{H}^d} \widehat{\kappa}(x)\mathcal{A}(\mathcal{V}[x])$$

so the top equation is satisfied iff

$$\sum_{x \in \mathcal{H}^d} \widehat{\kappa}(x)\mathcal{A}(\mathcal{V}[x]) = s_\psi$$

Recall that $\mathcal{A} \in \mathsf{A}_{yes}$, and in particular, it assigns $\mathcal{V}$ an r_{yes}-degree polynomial. Therefore, by the definition of $\widehat{\kappa}_\psi$,

$$\sum_{x \in \mathcal{H}^d} \widehat{\kappa}(x)\mathcal{A}(\mathcal{V}[x]) = \sum_{x \in \mathcal{F}^d} \kappa(x)\mathcal{A}(\mathcal{V}[x])$$

so the top equation is satisfied iff

$$\sum_{x \in \mathcal{F}^d} \kappa(x)\mathcal{A}(\mathcal{V}[x]) = s_\psi$$

that is, iff ψ is satisfied. Since $\mathcal{A}$ is a satisfying assignment, ψ is indeed satisfied, and therefore so is the top equation.

Restriction: Let $\widetilde{\mathcal{A}} \in \widetilde{\mathsf{A}}_{no}$ be an assignment for $\widetilde{\Psi}$. We shall define an assignment $\mathcal{A} \in \mathsf{A}_{no}$, that satisfies almost the same fraction as $\widetilde{\mathcal{A}}$. As required by the restriction property, $\mathcal{A}$ will be defined using only the restriction of $\widetilde{\mathcal{A}}$ to $\mathcal{V} \cup \mathcal{V}_{ext}$.

On $\mathcal{V}_{ext}$, we define $\mathcal{A}$ to be the restriction of $\widetilde{\mathcal{A}}$ to $\mathcal{V}ext$. As for $\mathcal{V}$, we will define $\mathcal{A}$ as a transformation of $\widetilde{\mathcal{A}}|_{\mathcal{V}}$ - $\mathcal{A}|_{\mathcal{V}}$ is defined to be the proper extension of $\widetilde{\mathcal{A}}|_{\mathcal{H}^d}$, with parameter r_{yes}. $\mathcal{A}_{\mathcal{V}}$ is therefore of degree r_{yes} in each variable.

We will show that for each $\psi \in \Psi$, the system $\widetilde{\Psi}_\psi \subseteq \widetilde{\Psi}$ *verifies* the equation ψ, in the following sense.

Claim 15 *if $\mathcal{A}$ does not satisfy ψ, then $\omega(\widetilde{\Psi}_\psi, \widetilde{\mathcal{A}}) < \frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$.*

Note that $\left|\widetilde{\Psi}_\psi\right| = |\mathcal{F}|^d$, is the same for each ψ, and therefore $\omega(\widetilde{\Psi}, \widetilde{\mathcal{A}}) = \mathbf{avg}_{\psi \in \Psi}(\widetilde{\Psi}_\psi, \widetilde{\mathcal{A}})$. We thus derive from the above claim that

$$\omega(\Psi, \mathcal{A}) > \omega(\widetilde{\Psi}, \widetilde{\mathcal{A}}) - \frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$$

Since $d \leq O((\log n)^\delta)$, and $r_{\mathtt{no}} = |\mathcal{F}|^{c_1}$ for some $c_1 < 1$, there exists $c_2 > 0$ such that $\frac{2r_{\mathtt{no}}d^2}{|\mathcal{F}|} < |\mathcal{F}|^{-c_2}$. Therefore

$$\omega(\Psi, \mathcal{A}) > \omega(\widetilde{\Psi}, \widetilde{\mathcal{A}}) - |\mathcal{F}|^{-c_2}$$

which is what we needed to show. It is only left to prove claim 15.

Proof of Claim 15. Suppose ψ is the equation

$$\sum_{x \in \mathcal{F}^d} \kappa_\psi(x)\mathcal{V}[x] = s_\psi \qquad (\psi)$$

We will prove the claim in three steps, that correspond to the three types of equations in each path-check.

The Base. Define $f : \mathcal{F}^d \to \mathcal{F}$ by

$$\forall\, x \in \mathcal{F}^d \quad f(x) = \widetilde{\mathcal{A}}(\mathcal{V}[x])\widehat{\kappa}_\psi(x)$$

and define f' as the restriction of $\widetilde{\mathcal{A}}$ to $\mathcal{D}_\psi[(d, \cdot)]$. We will now show that if $f \neq f'$, then $\omega(\widetilde{\Psi}, \widetilde{\mathcal{A}}) < \frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$.

$\widehat{\kappa}_\psi$ is an $r_{yes}d$-degree LDF (as we noted before), and $\widetilde{\mathcal{A}}(\mathcal{V}[\,\cdot\,])$ is an $r_{\mathtt{no}}$-degree LDF because of the LDF-assumption over $\mathcal{V}$. Therefore f is an LDF of degree $r_{\mathtt{no}} + d(r_{yes}+1) \leq 2r_{\mathtt{no}}$. As for f', it is an $[r_{\mathtt{no}}, d]$-LDF due to the assumption over $\mathcal{D}$.

We have exactly one path check $path[\psi, x]$ for each point $x \in \mathcal{F}^d$, and the base equation of this path-check verifies that $f(x) = f'(x)$. Since both f and f' are $[2r_{\mathtt{no}}, d]$-LDFs, if f does not equal f', the base equation of $path[\psi, x]$ fails to be satisfied for all x's but a $\frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$-fraction. Therefore $\omega(\widetilde{\Psi}_\psi, \widetilde{\mathcal{A}}) < \frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$, as needed.

The Path. Let f be defined as above, and let $(g_1, \ldots, g_d)$ be the sum-check tree of f with parameter r_{yes}. The functions $g_1, \ldots, g_k$ are $2r_{\mathtt{no}}$-degree LDFs. For $k = 1, 2, \ldots, d$ we define an $r_{\mathtt{no}}$-degree LDF $g'_k : \mathcal{F}^k \to \mathcal{F}$ by

$$\forall\, x = (x_1, \ldots, x_k) \in \mathcal{F}^k \quad g'_k(x) = \widetilde{\mathcal{A}}(\mathcal{D}_\psi[k, x_1, \ldots, x_k, \bar{0}])$$

We will show that if $(g'_1, \ldots, g'_d) \neq (g_1, \ldots, g_d)$, then $\omega(\widetilde{\Psi}, \widetilde{\mathcal{A}}) < \frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$.

First, suppose $g'_d \neq g_d$. Note that $g'_d = f'$ and $g_d = f$, so in this case the inequality follows from the base step above. In any other case, let $k < d$ be the largest which satisfies $g'_k \neq g_k$. Note that by definition,

$$g_k(\,\cdot\,) = \sum_{i=0}^{r_{yes}} g_{k+1}(\,\cdot\,, i) \tag{1}$$

Now choose a random point $x = (x_1, \ldots, x_d) \in \mathcal{F}^d$. Observe that $(x_1, \ldots, x_k)$ is a random point in $\mathcal{F}^k$, and also that $path[\psi, x]$ is a random conjunction in $\widetilde{\Psi}_\psi$.

Consider the path equation of $path[\psi, x]$ that corresponds to k. It is satisfied iff $g'_k(x_1, \ldots, x_k) = \sum_{i=0}^{r_{yes}} g'_{k+1}(x_1, \ldots, x$
By the choice of k, $g'_{k+1} = g_{k+1}$, and together with equation 1, we deduce that this path equation is satisfied iff $g'_k(x_1, \ldots, x_k) = g_k(x_1, \ldots, x_k)$.

Since $g'_k \neq g_k$, the probability of the equation (and therefore of $path[\psi, x]$) to be satisfied is bounded by $\frac{2r_{\mathtt{no}}k}{|\mathcal{F}|} < \frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$. Since we chose a random path-check in $\widetilde{\Psi}_\psi$, we gather that $\omega(\widetilde{\Psi}_\psi, \widetilde{\mathcal{A}}) < \frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$, as needed.

The Top. We now show the claim, i.e. that if $\mathcal{A}$ does not satisfy ψ, then $\omega(\widetilde{\Psi}_\psi, \widetilde{\mathcal{A}}) < \frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$. Under the notation as above, we may assume without loss of generality that $(g'_1, \ldots, g'_d) = (g_1, \ldots, g_d)$ - otherwise $\omega(\widetilde{\Psi}_\psi, \widetilde{\mathcal{A}}) < \frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$, by the previous sections.

Therefore $\mathcal{D}_\psi(d, \cdot, \bar{0})$ equals g_1, so the top equation of *all* path-checks $path[\psi, \cdot\,]$ (all of them are the same), is satisfied iff

$$\sum_{i=0}^{r_{yes}} g_1(i) = s_\psi$$

By definition,

$$\sum_{i=0}^{r_{yes}} g_1(i) = \sum_{x \in \mathcal{H}^d} g_d(x) = \sum_{x \in \mathcal{H}^d} \widehat{\kappa}_\psi(x)\widetilde{\mathcal{A}}(\mathcal{V}[x])$$

so the path-checks are satisfied iff

$$\sum_{x\in\mathcal{H}^d}\kappa'_\psi(x)\mathcal{A}(\mathcal{V}[x])=\sum_{x\in\mathcal{H}^d}\widehat{\kappa}_\psi(x)\widetilde{\mathcal{A}}(\mathcal{V}[x])=\sum_{x\in\mathcal{H}^d}g_d(x)=s_\psi$$

According to remark 1, the above equation is satisfied iff $\mathcal{A}$ satisfies ψ. By our assumption, $\mathcal{A}$ does not satisfy ψ, hence all the top equations are not satisfied, and therefore so are all the path-checks in $\widetilde{\Psi}_\psi$. This implies that $\omega(\widetilde{\Psi}_\psi,\widetilde{\mathcal{A}})=0<\frac{2r_{\mathtt{no}}d}{|\mathcal{F}|}$, and we conclude the proof of claim 15

Finally, to end up with systems of equations instead of systems of conjunctions, we replace (as in proposition 28), every path conjunction $\varphi_1\wedge\cdots\wedge\varphi_{d+1}$ in $\widetilde{\Psi}$ with a set (of size $|\mathcal{F}|^{d+1}$) of linear equations $\{\sum_{i=0}^{d+1}\alpha_i\varphi_i \mid \alpha_i\in\mathcal{F}\}$ (where linear combinations of equations are defined in the obvious manner). The properties of $\widetilde{\Psi}$ are maintained (the weight of the system relative to any assignment may change by at most a $\frac{1}{|\mathcal{F}|}$ factor by this transition, as shown in the proof of proposition 28). ■

4.3 The Curve-Extension Lemma

This lemma uses a technique called curve extension, similar to [ALM+92, LS91, FL92], to reduce the principal degree and dimension of the input system.

Lemma 16 (Curve-Extension) *The curve-extension lemma is obtained from the lemma-template by setting its parameters to*

$\mathcal{C}$ - *The parameters of the input system should satisfy* $r_{yes}, r_{\mathtt{no}}, d, D<|\mathcal{F}|^{c_1}$ *for some constant* $0<c_1<\frac{1}{2}$ *such that* $d\cdot\sqrt[d]{r_{yes}D}<r_{\mathtt{no}}$.

$\widetilde{r_{yes}}=\min\{d,2\log_2 D\}\cdot\max\{(r_{yes}D)^{\frac{1}{d}},2\}$

$\widetilde{D}=r_{yes}D$

$\widetilde{d}=\min\{d,\log_2 D\}$

Proof: For each equation of Ψ, we create a domain that contains a 'copy' of its variables. The degree of the LDF-assumption on this domain is significantly smaller than that of $\mathcal{V}$. Then we substitute each equation by a set of conjunctions, that depend mostly on the new variables but also verify consistency with the old variables.
Focus on one equation $\psi\in\Psi$. ψ is of the form

$$\sum_{i=1}^{D}\alpha_i\mathcal{V}[x_i]=\sum_{j=1}^{c}\beta_j v_j\qquad(\alpha_i,\beta_j\in\mathcal{F})\qquad(\psi)$$

where $x_i\in\mathcal{F}^d$ and $v_j\notin\mathcal{V}$.

Let $\Gamma_\psi:\mathcal{F}\to\mathcal{F}^d$ be the D-degree curve satisfying

$$\forall 1\le i\le D\quad\Gamma_\psi(i)=x_i$$

We add a set $\mathcal{D}_\psi$ of $r_{yes}D$ *new* variables - a variable $\mathcal{D}_\psi[\Gamma_\psi(i)]$, for each point $\Gamma_\psi(i),\quad i=1,2,\ldots,r_{yes}D$.

Since we assume the assignment of $\mathcal{V}$ makes an $[r_{yes},r_{\mathtt{no}},d]$-LDF, its composition with $\Gamma_\psi(i)$ would be an $r_{yes}D$-degree LDF. Given the values of such an LDF on an arbitrary set of $r_{yes}D$ points, its value at *any* other

point can be linearly interpolated. We can thus define ψ_a for every $a \in \mathcal{F}$, to be the conjunction of the following two equations,

$$\sum_{i=1}^{D} \alpha_i \cdot \mathcal{D}_\psi[\Gamma_\psi(i)] = \sum_{j=1}^{c} \beta_j v_j \qquad (\psi|_{\mathcal{D}_\psi})$$

$$\mathcal{V}[\Gamma_\psi(a)] = \sum_{i=1}^{r_{yes}D} \gamma_i \cdot \mathcal{D}_\psi[\Gamma_\psi(i)] \qquad (a)$$

where $\gamma_1, \ldots, \gamma_{r_{yes}D} \in \mathcal{F}$ are the interpolation coefficients, such that every $[r_{yes}D, 1]$-LDF, p, satisfies

$$p(a) = \sum_{i=1}^{r_{yes}D} \gamma_i \cdot p(i)$$

(we leave it to the reader to verify the existence of such coefficients).

Define the equation set for ψ as $\widetilde{\Psi}_\psi = \{\psi_a \mid a \in \mathcal{F}\}$, and let $\widetilde{\Psi} = \cup_{\psi \in \Psi} \widetilde{\Psi}_\psi$. Note that $\left|\widetilde{\Psi}_\psi\right|$ is the same for every $\psi \in \Psi$, as required. It is easy to see that $depend(\widetilde{\Psi}) = r_{yes}D + 1 + c$, and that for every equation $\psi \in \widetilde{\Psi}_\psi$, ψ depends on $\leq c + 1$ variables outside $\mathcal{D}_\psi$.

For each $\psi \in \Psi$ we add new variables to $\mathcal{D}_\psi$ - that the extended $\mathcal{D}_\psi$ will have a variable for each point in the proper-extension of the old variables of $\mathcal{D}_\psi$, with parameter $h \stackrel{def}{=} \max\{\sqrt[d]{D}, 2\}$. The dimension of $\mathcal{D}_\psi$ would therefore be $\widetilde{d} \stackrel{def}{=} \min\{d, \log_2 D\}$. We introduce an $[h\widetilde{d}, r_{\mathtt{no}}, \widetilde{d}]$ LDF-assumption over each $\mathcal{D}_\psi$, so the assumptions of $\widetilde{\Psi}$ comply with the terms of the lemma. We have yet to verify that the gap is maintained.

The Extension and Restriction Properties.

Denote $(\mathsf{A}_{yes}, \mathsf{A}_{no}) \stackrel{def}{=} \mathcal{L}_{\mathcal{V}}[r_{yes}, r_{\mathtt{no}}, d] \cap \mathsf{A}_{ext}$.

Extension: Let $A \in \mathsf{A}_{yes}$ be an assignment with $\omega(\Psi, A) = 1$. We construct an assignment $\widetilde{\mathcal{A}}$, that equals A on $\mathcal{V} \cup \mathcal{V}_{ext}$. For each $\psi \in \Psi$, we assign $\mathcal{D}_\psi$ the proper-extension of $\{A(\mathcal{V}[\Gamma_\psi(i)])\}_{i=1}^{r_{yes}D}$.

It is obvious that $\widetilde{\mathcal{A}} \in \mathsf{A}_{yes}$ and that $\omega(\widetilde{\Psi}, \widetilde{\mathcal{A}}) = 1$.

Restriction: Let $\widetilde{\mathcal{A}} \in \mathsf{A}_{no}$ be an arbitrary assignment, and let A be its restriction to $\mathcal{V} \cup \mathcal{V}_{ext}$ (so in this case, the transformation $\mathcal{T}$ mentioned in the lemma-template is just the identity).

We define g to be the $[r_{\mathtt{no}}, d]$-LDF defined by,

$$g(x) \stackrel{def}{=} \widetilde{\mathcal{A}}(\mathcal{V}[x])$$

It follows that $g \circ \Gamma_\psi : \mathcal{F} \to \mathcal{F}$ is an $[r_{\mathtt{no}}D, 1]$-LDF.

Let $f : \mathcal{F} \to \mathcal{F}$ be the $[r_{yes}D, 1]$-LDF defined by

$$\forall 1 \leq i \leq r_{yes}D \qquad f(i) \stackrel{def}{=} \widetilde{\mathcal{A}}(\mathcal{D}_\psi[\Gamma_\psi(i)])$$

The second equation (labeled (a)) in each conjunction of $\widetilde{\Psi}_\psi$, checks whether $g \circ \Gamma_\psi(a) = f(a)$ for a point $a \in \mathcal{F}$. If these LDFs are different, they agree on at most a $\frac{r_{\mathtt{no}}D}{|\mathcal{F}|}$ fraction of their domain. If for some

$1 \leq i \leq r_{yes}D$, $\widetilde{\mathcal{A}}(\mathcal{V}[\boldsymbol{x}]) \neq \widetilde{\mathcal{A}}(\mathcal{D}_\psi[\boldsymbol{x}])$ on the point $\boldsymbol{x} = \Gamma_\psi(i)$, the LDFs are bound to differ. This bounds by $\frac{r_{\mathtt{no}}D}{|\mathcal{F}|}$ the fraction of conjunctions in $\widetilde{\Psi}$ originating from $\psi \in \Psi$ where

$$\exists 0 \leq i \leq r_{yes}D \qquad \widetilde{\mathcal{A}}(\mathcal{D}_\psi[\Gamma_\psi(i)]) \neq \widetilde{\mathcal{A}}(\mathcal{V}[\Gamma_\psi(i)])$$

The equation $(\psi|_{\mathcal{D}_\psi})$ that appears in every conjunction in $\widetilde{\Psi}_\psi$, implies that the fraction of conjunctions in $\widetilde{\Psi}$ that evaluate to `true` and that originated from $\psi \in \Psi$ for which

$$\forall 0 \leq i \leq r_{yes}D \qquad \widetilde{\mathcal{A}}(\mathcal{D}_\psi[\Gamma_\psi(i)]) = \widetilde{\mathcal{A}}(\mathcal{V}[\Gamma_\psi(i)])$$

is bounded by $\omega(\Psi, \mathcal{A})$.

Altogether, we get that $\omega(\widetilde{\Psi}, \widetilde{\mathcal{A}}) < \omega(\Psi, \mathcal{A}) + \frac{r_{\mathtt{no}}D}{|\mathcal{F}|}$. Obviously there exists a constant $c_2 > 0$, such that $\frac{r_{\mathtt{no}}D}{|\mathcal{F}|} < |\mathcal{F}|^{-c_2}$, which implies $\omega(\widetilde{\Psi}, \widetilde{\mathcal{A}}) < \omega(\Psi, \mathcal{A}) + |\mathcal{F}|^{-c_2}$.

Finally, we apply the linear combinations technique (see proposition 28) to obtain a system of equations (rather than conjunctions) with the same properties (up to changes in weight of up to $\frac{1}{\mathcal{F}}$). ■

4.4 The Constant-Depend Lemma

We now show a technique that is similar to the curve-extension technique, except that it uses linearization to obtain *constant* depend.

Lemma 17 (Constant Depend) *The constant-depend lemma is obtained from the lemma-template by setting its parameters to*

- $\mathcal{C}$ - *The parameters of the input system should satisfy* $r_{yes}D < O(\log^\delta n)$, *and* $r_{\mathtt{no}}, d < |\mathcal{F}|^{c'}$ *for some* $0 < c' < \frac{1}{2}$.
- $\widetilde{r_{yes}} = 1$
- $\widetilde{D} = 2$
- $\widetilde{d} = r_{yes}D$

Proof: We create for each equation of Ψ, a domain that contains not only a 'copy' of its variables, but all of their linear combinations. We then substitute the equation by a set of conjunctions, that depend on the new variables and verify consistency with the old variables.
Focus on one $\psi \in \Psi$. It is of the form

$$\alpha_i, \beta_j \in \mathcal{F}, \qquad \sum_{i=1}^{D} \alpha_i \mathcal{V}[\boldsymbol{x}_i] = \sum_{j=1}^{c} \beta_j v_j \qquad (\psi)$$

where $\boldsymbol{x}_i \in \mathcal{F}^d$ and $v_j \notin \mathcal{V}$.

Let $\Gamma_\psi : \mathcal{F} \to \mathcal{F}^d$ be the D-degree curve satisfying

$$\forall 1 \leq i \leq D, \;\; \Gamma_\psi(i) = \boldsymbol{x}_i$$

(actually it is a $(D-1)$-degree curve, but we write D for simplicity. We can force the curve to go through an additional arbitrary point, to make the discussion exact).

We add a set $\mathcal{D}_\psi$ of new variables, one for each linear combination of $\mathcal{V}[\Gamma_\psi(1)], \ldots, \mathcal{V}[\Gamma_\psi(s)]$, where $s \stackrel{def}{=} r_{yes}D$. For $z = (z_1, .., z_s) \in \mathcal{F}^s$, denote the variable that corresponds to $\sum_{i=1}^{s} z_i \mathcal{V}[\Gamma_\psi(i)]$ by $\mathcal{D}_\psi[z]$ (In particular, the variable $\mathcal{D}_\psi[(\alpha_1, .., \alpha_D, \bar{0})]$ corresponds to $\sum_{i=1}^{D} \alpha_i \mathcal{V}[x_i]$, which is the left-hand side of (ψ)).

Note that the variables of $\mathcal{D}_\psi$ are naturally idetified with the points of $\mathcal{F}^s$, and that $|\mathcal{F}|^s$ is polynomial in n, thanks to the condition on r_{yes} and D. We therefore introduce the LDF-assumptions $\forall\psi \quad \mathcal{L}_{\mathcal{D}_\psi}[1, R, s]$.
For every $a \in \mathcal{F}$, let ψ_a be the conjunction of the following two equations

$$\mathcal{D}_\psi[(\alpha_1, .., \alpha_D, \bar{0})] = \sum_{j=1}^{c} \beta_j v_j \qquad (\psi|_{\mathcal{D}_\psi})$$

and

$$\mathcal{D}_\psi[\bar{z}_a] = \mathcal{V}[\Gamma_\psi(a)] \qquad (a)$$

where $\bar{z}_a = (z_1, \ldots, z_s) \in \mathcal{F}^s$ is the interpolation vector, such that every $[r_{yes}D, 1]$-LDF, p, satisfies

$$p(a) = \sum_{i=1}^{r_{yes}D} \gamma_i \cdot p(i)$$

(we leave it to the reader to verify the existence of such coefficients).

The conjunction set for ψ is defined to be $\widetilde{\Psi}_\psi = \{\psi_a \mid a \in \mathcal{F}\}$, and define $\widetilde{\Psi} = \cup_{\psi\in\Psi} \widetilde{\Psi}_\psi$. Note that $\left|\widetilde{\Psi}_\psi\right|$ is the same for every $\psi \in \Psi$, as required. It is easy to see that $depend(\widetilde{\Psi}) = c + 3$. We have yet to verify the extenstion and restriction properties.

The Extension and Restriction Properties.

Denote $\mathcal{L}_\mathcal{V}[r_{yes}, r_{\mathtt{no}}, d] = (\mathsf{A}_{yes}, \mathsf{A}_{no})$, and denote $(\widetilde{\mathsf{A}}_{yes}, \widetilde{\mathsf{A}}_{no}) \stackrel{def}{=} \mathcal{L}_\mathcal{V}[r_{yes}, r_{\mathtt{no}}, d] \cap (\bigcap_{\psi\in\Psi}(\mathcal{L}_{\mathcal{D}_\psi}[1, r_{\mathtt{no}}, d]))$.

Extension: Let $A \in \mathsf{A}_{yes}$ be an assignment with $\omega(\Psi, A) = 1$. Let $\widetilde{\mathcal{A}}$ be the assignment that equals A on $\mathcal{V} \cup \mathcal{V}_{ext}$, and for every point $z = (z_1, .., z_s) \in \mathcal{F}^s$ assigns $\mathcal{D}_\psi[z]$ the value $\sum_{i=1}^{s} z_i \cdot A(\mathcal{V}[\Gamma_\psi(i)])$, obviously obeys $\widetilde{\mathcal{A}} \in \widetilde{\mathsf{A}}_{yes}$ and $\omega(\widetilde{\Psi}, \widetilde{\mathcal{A}}) = 1$.

Restriction: Let $\widetilde{\mathcal{A}} \in \widetilde{\mathsf{A}}_{no}$ be an arbitrary assignment, and take $\mathcal{A}$ to be its restriction to $\mathcal{V} \cup \mathcal{V}_{ext}$ (so in this case, the transformation $\mathcal{T}$ mentioned in the lemma-template is just the identity). Let $f(x) \stackrel{def}{=} \widetilde{\mathcal{A}}(\mathcal{V}[x])$ be an $[r_{\mathtt{no}}, d]$-LDF. Consider an arbitrary equation $\psi \in \Psi$. By definition of Γ_ψ it follows that $f \circ \Gamma_\psi : \mathcal{F} \to \mathcal{F}$ is an $[r_{\mathtt{no}}D, 1]$-LDF.

Define an $[r_{yes}D, 1]$-LDF $g : \mathcal{F} \to \mathcal{F}$, by

$$\forall 1 \le i \le r_{yes}D \qquad g(i) \stackrel{def}{=} \widetilde{\mathcal{A}}(\mathcal{D}_\psi[\Gamma_\psi(i)])$$

The second equation in the conjunctions of $\widetilde{\Psi}_\psi$ (labeled (a)), checks whether $f \circ \Gamma_\psi(a) = g(a)$ for a point $a \in \mathcal{F}$. If, for some equation $\psi \in \Psi$, these LDFs are different, they equal on at most $\frac{r_{\mathtt{no}}D}{|\mathcal{F}|}$ of their domain. This implies that a $1 - \frac{r_{\mathtt{no}}D}{|\mathcal{F}|}$ fraction of the conjunctions in $\widetilde{\Psi}_\psi$ will evaluate to `false`. However, whenever these LDFs are equal, the first equation in the conjunction (labeled $(\psi|_{\mathcal{D}_\psi})$) assures that unless ψ is satisfied by $\widetilde{\mathcal{A}}$, all of the conjunctions in $\widetilde{\Psi}_\psi$ will evaluate to `false`.

Altogether, we have that $\omega(\Psi, \mathcal{A}) > \omega(\widetilde{\Psi}, \widetilde{\mathcal{A}}) - \frac{r_{\mathtt{no}}D}{|\mathcal{F}|}$. It is easy to see the existence of a constant $c_2 > 0$ for which $\frac{r_{\mathtt{no}}D}{|\mathcal{F}|} \le |\mathcal{F}|^{-c_2}$.

To transform $\widetilde{\Psi}$ from a system of conjunctions to a system of equations, we substitute each conjunction for the set of linear combinations of its equations, as in the previous proofs.

■

5 Appendix

5.1 Cube vs. Point Test

To make the proof of lemma 3 complete, it is left to prove lemma 6. This is a cube-vs.point version of the low-degree-test of [RS97].

Lemma 6 (Cube-vs.-Point) *Let $\mathcal{D} = \mathcal{F}^d$ and $f : \mathcal{D} \to \mathcal{F}$. Let $\mathbf{x} = \{x_1, \ldots, x_D\} \subseteq \mathcal{D}$, and suppose we have an $[r, D+3]$-LDF, $F_{\mathcal{C}} : \mathcal{C} \to \mathcal{F}$, for each cube $\mathcal{C} \in \mathcal{S}_{\mathbf{x}}$.*

Choose a random point $x \in_R \mathcal{D}$ and a random cube $\mathcal{C} \in_R \mathcal{S}_{\mathbf{x} \cup \{x\}}$. The probability that $f(x) = F_{\mathcal{C}}(x)$, yet $F_{\mathcal{C}}$ is not $[r, \rho]$-permissible, is $O(\rho^{\frac{1}{6}})$. This is true for every choice of ρ which satisfies $\rho > 2\sqrt{\frac{rd}{|\mathcal{F}|}}$

We begin by citing a lemma by Raz and Safra, and then by a series of claims derived from it we get the required result.

For simplicity, we use the following notations and conventions throughout the proof.

1. An LDF, in this section, refers to an $[r, t]$-LDF.
2. By a ρ-permissible LDF over a subspace of $\mathcal{F}^d$ we mean a restriction of a ρ-permissible LDF over $\mathcal{F}^d$ to that subspace.
3. Whenever we refer to a random choice of a member from a set (e.g. $y \in_R \mathcal{F}^d$) without specifying the probability distribution, the uniform distribution is used.
4. The claims we make about ρ-permissible polynomials over $\mathcal{F}^d$, are only made for $\rho > 2\sqrt{\frac{rd}{|\mathcal{F}|}}$.

Suppose that each of the planes (i.e. affine subspaces of dimension 2) in a domain $\mathcal{D} = \mathcal{F}^d$ is provided with an LDF, i.e. a polynomial function of degree r. The following lemma states that, in a sense, these LDFs may agree with each other (on plane intersections) only as far as they originate from restrictions of global LDFs over $\mathcal{D}$ to the planes.

Lemma 18 ([RS97]) *Let $|\mathcal{F}|$, r, d and ρ satisfy the conditions above. Let $\mathcal{D} = \mathcal{F}^d$. Let $\mathcal{S}$ be the set of planes (2-cubes) on $\mathcal{D}$, with each $\wp \in \mathcal{S}$ being provided with an LDF $f_\wp : \wp \to \mathcal{F}$. A global LDF over $\mathcal{D}$ is called* **ρ-planewise-permissible** *if it agrees with ρ fraction of the LDFs over the planes. Let $\wp_1, \wp_2 \in_R \mathcal{S}$ be two random planes with a (random) line intersection. The probability that $f_{\wp_1}$ and $f_{\wp_2}$ agree on the intersection and yet neither of them is (a restriction of) a ρ-planewise-permissible function is $O(\rho)$.*

We now begin a series of claims culminating in the desired lemma. First of all, since comparing two planes on the whole intersection is costly in terms of the depend parameter, we show that it is enough to compare them on one random point.

Claim 19 *Let $\mathcal{D}$ and $\mathcal{S}$ be as above, with an LDF $f_\wp : \wp \to \mathcal{F}$ provided for each $\wp \in \mathcal{S}$. Let $\wp_1, \wp_2 \in_R \mathcal{S}$ be two random planes with a (random) line intersection $\mathcal{L}$, and $y \in_R \mathcal{L}$ a random point on that line. The probability that $f_{\wp_1}$ and $f_{\wp_2}$ agree on y and yet $f_{\wp_1}$ is not a ρ-planewise-permissible function is $O(\rho)$.*

Proof: Since the restrictions of $f_{\wp_1}$ and $f_{\wp_2}$ to $\mathcal{L}$ are in particular LDFs, if these two differ then they differ on all but at most r of the points $y \in \mathcal{L}$. Since $\frac{r}{|\mathcal{F}|} \ll \rho$ the claim follows directly from Lemma 18. ■

In the previous claim, two LDFs provided for two planes are read, and then their values are compared on one point. Since in the main body of the proof this type of comparison needs to be iterated several times in a recursive manner, this causes an (exponential) inflation of the dependency. In order to avoid this, we may not access more than one LDF at a time. Thus, instead of comparing two plane functions with each other, we shall compare one such LDF with a new function over $\mathcal{D}$, which is not assumed to be an LDF. So only one instance of reading a value of an LDF is performed.

Claim 20 *Let $\mathcal{D}$, $\mathcal{S}$ and $\{f_\wp | \wp \in \mathcal{S}\}$ be as above, and let $f : \mathcal{D} \to \mathcal{F}$ be an additional function, not necessarily an LDF. Let $\wp \in_R \mathcal{S}$ be a random plane and $y \in_R \wp$ a random point from that plane. The probability that $f_\wp$ agrees with f on y and yet $f_\wp$ is not ρ-planewise-permissible is $O(\rho^{1/2})$.*

Proof: Let us denote by E_1 the event that, in the procedure described here, $f_\wp$ agrees with f on y and yet is an impermissible function. Considering now the process described in Claim 19, of choosing two random planes $\wp_1$, $\wp_2$ intersecting on a line $\mathcal{L}$ and a point $y \in_R \mathcal{L}$, let us denote by E_2 the event that $f_{\wp_1}$ and $f_{\wp_2}$ agree on y and yet $f_{\wp_1}$ is impermissible. Remember that $\Pr(E_2) = O(\rho)$ by that claim. We prove for every fixed $y_0 \in \mathcal{D}$ that

$$\Pr(E_1 | y = y_0) \leq (\Pr(E_2 | y = y_0))^{1/2} .$$

The claim then follows since $f(\beta) = \beta^{1/2}$ is a concave function, so

$$\begin{aligned} \Pr(E_1) &= |\mathcal{D}|^{-1} \sum_{y_0 \in \mathcal{D}} \Pr(E_1 | y = y_0) \\ &\leq |\mathcal{D}|^{-1} \sum_{y_0 \in \mathcal{D}} \Pr(E_2 | y = y_0)^{1/2} \\ &\leq \left(|\mathcal{D}|^{-1} \sum_{y_0 \in \mathcal{D}} \Pr(E_2 | y = y_0) \right)^{1/2} = O(\rho^{1/2}) \end{aligned}$$

To prove the claim for a fixed y_0, denote by $\mathcal{S}'_{y_0}$ the set of planes which contain y_0 for which impermissible functions were provided. Fixing $\{f_\wp | \wp \in \mathcal{S}\}$, one can assume that $f(y_0)$ is the most common value $f_\wp(y_0)$ corresponding to planes $\wp \in \mathcal{S}'_{y_0}$; this is the worst case for the provided functions for $\mathcal{S}$ since the definition of planewise permissibility is not dependent on $f : \mathcal{D} \to \mathcal{F}$. The fraction of pairs of members of $\mathcal{S}'_{y_0}$ for which the provided functions agree on y_0 is bounded below by the square of the fraction of planes for which the provided functions receive on y_0 the most common value, so the claim follows. ■

The following two claims show that distinct planewise permissible functions differ greatly on their restrictions to planes, and so there are (for a given assignment to $\mathcal{S}$) only a few of them. The fewness of planewise permissible functions allows us later to make claims for a slightly different family of functions, with only a slight increase of the probabilities involved.

Claim 21 *Two distinct LDFs over $\mathcal{D}$ will disagree on all but at most $\frac{rd}{|\mathcal{F}|}$ of possible restrictions to planes.*

Proof: Since two distinct LDFs will disagree on all but at most $\frac{rd}{|\mathcal{F}|}$ of the points, they will, for a choice of a random $\wp \in_R \mathcal{S}$ and then a random $y \in_R \wp$, produce a disagreement with probability at least $1 - \frac{rd}{|\mathcal{F}|}$. Thus, a random plane will with probability at least $1 - \frac{rd}{|\mathcal{F}|}$ produce a disagreement on at least one of its points. ■

Claim 22 *For a given set of functions provided for $\mathcal{S}$, there are less than $2\rho^{-1}$ of ρ-planewise-permissible LDFs over $\mathcal{D}$ in all.*

Proof: Supposing there are $2\rho^{-1}$ such functions, by claim 21 and the assumption $\rho > 2\sqrt{\frac{rd}{|\mathcal{F}|}}$ which implies in particular $2\rho^{-1}\frac{rd}{|\mathcal{F}|} < \frac{1}{2}\rho$, one sees that for each function there are more than $\frac{1}{2}\rho|\mathcal{S}|$ planes for which the provided functions agree with it alone. This is a contradiction. ■

In order to have more freedom when dealing with the planes (they will be replaced shortly by affine subspaces of a greater dimension), and in order to make proofs easier, we define a new notion of permissibility which depends on f rather than $\{f_{\wp}|\wp \in \mathcal{S}\}$. The concept of permissibility used outside this section is called here pointwise permissibility, to differentiate it from the planewise permissibility defined in the previous claims.

Claim 23 *Let $\mathcal{D} = \mathcal{F}^d$ and $\mathcal{S}$ be as above, and provided with the functions f and $\{f_{\wp}|\wp \in \mathcal{S}\}$ as above. Define an LDF over $\mathcal{D}$ to be ρ-***pointwise-permissible** *if it agrees with f on at least ρ of the points $y \in \mathcal{D}$. Let $\wp \in_R \mathcal{S}$ be a random plane and $y \in_R \wp$ a random point from that plane. The probability that $f_{\wp}$ agrees with f on y and yet is not a ρ-pointwise-permissible function is $O(\rho^{1/3})$.*

Proof: Any LDF which agrees with some ρ fraction of the pairs in the set $\{(\wp, y)|\wp \in \mathcal{S}, y \in \wp\}$ (i.e. agrees with $f_{\wp}$ and f on y) agrees in particular with f for a ρ fraction of the points $y \in \mathcal{D}$, which means that it is ρ-pointwise-permissible.

By claim 20 the probability that $f_{\wp}$ agrees with f on y and yet is not $\rho^{2/3}$-planewise-permissible is $O(\rho^{1/3})$. By claim 22, there are no more than $2\rho^{-2/3}$ of $\rho^{2/3}$-planewise-permissible functions in all. By the above discussion, for each of these functions which is not ρ-pointwise-permissible, only with probability ρ or less $f_{\wp}$ will agree with it and also with f on y.

Let us now consider all $\rho^{2/3}$-planewise-permissible functions (of $\mathcal{D}$) which are also ρ-pointwise-permissible. The probability that $f_{\wp}$ will agree with f on y and yet will not be from this list is bounded according to the above discussion by

$$O(\rho^{1/3}) + 2\rho^{-2/3} \cdot \rho = O(\rho^{1/3})$$

so the claim follows. ■

In the final lemma the planes are replaced with cubes of some other fixed dimension. In order to bridge the gap, we first need to allow some randomeness in the functions $\{f_{\wp}|\wp \in \mathcal{S}\}$ provided for the planes.

Claim 24 *Let $\mathcal{D}$ and $\mathcal{S}$ be as above, but now for each $\wp \in \mathcal{S}$ let $f_{\wp} : \wp \to \mathcal{F}$ be chosen from a set according to some random distribution specified for $\wp$. Let now $\wp \in_R \mathcal{S}$ be a random plane and $y \in_R \wp$ a random point from that plane. The probability that (the chosen) $f_{\wp}$ agrees with f on y and yet is not ρ-pointwise-permissible is $O(\rho^{1/3})$.*

Proof: The process of first choosing a random plane and then a random provided function is isomorphic to the following: First choose a random function to each plane $\wp_0 \in \mathcal{S}$ (according to the distribution specified for each $\wp_0$), thus getting a random $\{f_{\wp}|\wp \in \mathcal{S}\}$. Then choose a random plane $\wp \in_R \mathcal{S}$, with the appropriate function from that set. With the definition of pointwise permissibility being independent of $\{f_{\wp}|\wp \in \mathcal{S}\}$, allowing this set to be random does not change the bound. ■

We next show a bound on the number of ρ-pointwise-permissible LDFs that is used in other parts of the proof. Here it is used for ensuring that an LDF over a cube which agrees with permissible LDFs on many of the cube's points (and hence must agree on many points with one such LDF), is a restriction of one permissible LDF by the interpolation property of LDFs.

Claim 25 *For the parameters involved, there are less than $2\rho^{-1}$ of ρ-pointwise-permissible functions in all.*

Proof: Similar to claim 22. ■

We now show that certain distributions, which appear naturally in our setting, are very close to the uniform ones. Thus previous claims remain applicable.

Claim 26 *Let $\mathbf{x} = \{x_1, \ldots, x_D\} \subseteq \mathcal{F}^d$ be a fixed arbitrary tuple. Let us choose a random $(D+3)$-cube $\mathcal{C}$ that contains $\mathbf{x}$, a random plane $\wp$ contained in $\mathcal{C}$ and a random point y contained in $\wp$. In this process,*

$$\sum_{\wp_0 \in \mathcal{S}} |\Pr(\wp = \wp_0) - |\mathcal{S}|^{-1}| \leq O(|\mathcal{F}|^{-1})$$

and

$$\sum_{y_0 \in \mathcal{F}^d} |\Pr(y = y_0) - |\mathcal{F}|^{-d}| \leq O(|\mathcal{F}|^{-1})$$

Proof: We shall prove here the second inequality. The proof of the first one is similar though more tedious.

Let us look at the process of choosing a random $(D+3)$-cube $\mathcal{C}$ that contains $\mathbf{x}$ and then a random $y \in_R \mathcal{C}$. Denote by X the affine subspace of $\mathcal{F}^d$ spanned by $\mathbf{x}$. This will be subspace of dimension $D-1$ or less, which is contained in all cubes containing $\mathbf{x}$. Hence,

$$\Pr(y_0 \in X) = |\mathcal{F}|^{-D-3}|X| \leq O(|\mathcal{F}|^{-1})$$

Note also that $\Pr(y = y_0) > |\mathcal{F}|^{-d}$ for each $y_0 \in X$.

For each $y_0 \in \mathcal{F}^d - X$, one notes that

$$|\mathcal{F}|^{-d} > \Pr(y = y_0) = \frac{1 - \Pr(y_0 \in X)}{|\mathcal{F}|^d - |X|} \geq (1 - O(|\mathcal{F}|^{-1}))|\mathcal{F}|^{-d}$$

and so

$$\begin{aligned} & \sum_{y_0 \in \mathcal{F}^d} |\Pr(y = y_0) - |\mathcal{F}|^{-d}| \\ = & \sum_{y_0 \in X} (\Pr(y = y_0) - |\mathcal{F}|^{-d}) + \sum_{y_0 \in \mathcal{F}^d - X} (|\mathcal{F}|^{-d} - \Pr(y = y_0)) \\ < & \Pr(y \in X) \cdot 1 + |\mathcal{F}|^d \cdot O(|\mathcal{F}|^{-d-1}) \;=\; O(|\mathcal{F}|^{-1}) \end{aligned}$$

■

The following proposition is *almost* the same as the desired lemma 6, other than a small difference in probability distributions.

Proposition 27 *Suppose now we are provided with an LDF $f_{\mathcal{C}}$ for each $(D+3)$-cube $\mathcal{C}$ of $\mathcal{D}$, and any function $f : \mathcal{D} \to \mathcal{F}$. Let $\mathbf{x} = \{x_1, \ldots, x_D\} \subseteq \mathcal{F}^d$ be any fixed arbitrary tuple. Choose a random cube $\mathcal{C}$ uniformly from the set of cubes that contain $\mathbf{x}$ and then (uniformly) a random point $y \in_R \mathcal{C}$. The probability that f and $f_{\mathcal{C}}$ agree on y and yet $f_{\mathcal{C}}$ is not (a restriction of) a ρ-pointwise-permissible function is $O(\rho^{1/6})$.*

Proof: First, let us look at the following process: A random cube $\mathcal{C}'$ containing $\mathbf{x}$ is chosen, and then a random plane $\wp$ which is contained in $\mathcal{C}'$ is chosen, and finally a random point y' from $\wp$ is chosen. One notes that the probability distribution of $\mathcal{C}'$ and y' in this process is the same as that of $\mathcal{C}$ and y in the process described in the lemma to be proven. One also notes that the probability distributions of y' and $\wp$ are very close to the uniform

ones by claim 26. Define $f_{\wp}$ as restriction of $f_{\mathcal{C}'}$ to $\wp$. For a fixed $\wp$, $f_{\wp}$ is considered a random function since it depends on the chosen $\mathcal{C}'$.

Denote by E_1 the event that $f_{\wp}$ agrees with f on y' and yet is an impermissible function. Applying Claim 24, $\Pr(E_1) = O(\rho^{1/3})$. Since this probability is an average over all cubes $\mathcal{C}_0$ containing $\mathbf{x}$ of the probabilities $\Pr(E_1|\mathcal{C} = \mathcal{C}_0)$, a $1 - O(\rho^{1/6})$ fraction of the cubes $\mathcal{C}_0$ are such that $\Pr(E_1|\mathcal{C} = \mathcal{C}_0) \leq \rho^{1/6}$.

Denote now by E_2 the event that $f_{\wp}$ agrees with f on y'. For each cube $\mathcal{C}_0$ as above, if $\Pr(E_2|\mathcal{C} = \mathcal{C}_0) \geq 2\rho^{1/6}$, then $\rho^{1/6}$ of the restrictions of $f_{\mathcal{C}}$ to planes $\wp \subset \mathcal{C}$ are permissible, and so for $\rho^{1/6}$ of the points $y' \in \mathcal{C}_0$ the value $f_{\mathcal{C}_0}(y')$ is permissible. There are less than $2\rho^{-1}$ permissible functions in all by Lemma 25, and the assumption $\rho > 2\sqrt{\frac{rd}{|\mathcal{F}|}}$ implies $2\rho^{-1}\frac{rd}{|\mathcal{F}|} < \rho^{1/6}$. Thus, by interpolation of LDFs, for a cube $\mathcal{C}_0$ for which $\Pr(E_1|\mathcal{C} = \mathcal{C}_0) \leq \rho^{1/6}$ and $\Pr(E_2|\mathcal{C} = \mathcal{C}_0) \geq 2\rho^{1/6}$, $f_{\mathcal{C}_0}$ is (a restriction to $\mathcal{C}_0$ of) one of the permissible functions, since in particular on more than $\frac{rd}{|\mathcal{F}|}$ of its points it agrees with one permissible function.

Finally, view $p_1(\mathcal{C}) = \Pr(E_1|\mathcal{C}' = \mathcal{C})$ and $p_2(\mathcal{C}) = \Pr(E_2|\mathcal{C}' = \mathcal{C})$ as functions of $\mathcal{C}$. We now assume that $\mathcal{C}$ is chosen randomly from all cubes containing $\mathbf{x}$, and define (new) events using these functions. Define E_3 as the event that $p_1(\mathcal{C}) > \rho^{1/6}$, and E_4 as the event that $p_1(\mathcal{C}) \leq \rho^{1/6}$ and $p_2(\mathcal{C}) < 2\rho^{1/6}$. Let $y \in_R \mathcal{C}$ be chosen (randomly, uniformly) too. By the above discussion, if neither E_3 nor E_4 occures, the function $f_{\mathcal{C}}$ is permissible. Whenever E_4 occures, the probability that $f_{\mathcal{C}}$ and f agree on y is bounded above by $2\rho^{1/6}$. Thus the probability that $f_{\mathcal{C}}$ agrees with f on y and yet $f_{\mathcal{C}}$ is an impermissible function is bounded by

$$\Pr(E_3) + 2\rho^{1/6}\Pr(E_4) \leq O(\rho^{1/6}) + 2\rho^{1/6} \cdot 1 = O(\rho^{1/6})$$

so the lemma follows. ■

Finally, we can restate and prove the lemma that is used for the main body of the proof, to assure that the values read there indeed correspond to values from the short list of pointwise permissible LDFs. This lemma differs from the previous proposition only in regards to the probability distribution by which the cube and point are chosen.

Lemma 6 (Cube-vs.-Point) *Let $\mathcal{D} = \mathcal{F}^d$ and $f : \mathcal{D} \to \mathcal{F}$. Let $\mathbf{x} = \{x_1, \ldots, x_D\} \subseteq \mathcal{D}$, and suppose we have an $[r, D+3]$-LDF, $F_{\mathcal{C}} : \mathcal{C} \to \mathcal{F}$, for each cube $\mathcal{C} \in \mathcal{S}_{\mathbf{x}}$.*

Choose a random point $x \in_R \mathcal{D}$ and a random cube $\mathcal{C} \in_R \mathcal{S}_{\mathbf{x}\cup\{x\}}$. The probability that $f(x) = F_{\mathcal{C}}(x)$, yet $F_{\mathcal{C}}$ is not $[r, \rho]$-permissible, is $O(\rho^{\frac{1}{6}})$. This is true for every choice of ρ which satisfies $\rho > 2\sqrt{\frac{rd}{|\mathcal{F}|}}$

Proof: The lemma follows from proposition 27 with a slight modification of distributions. Instead of first choosing a cube $\mathcal{C} \in_R \mathcal{S}_{\mathbf{x}}$ and then a point in it $x \in_R \mathcal{C}$, we first choose a point $x \in_R \mathcal{D}$, and then a cube $\mathcal{C} \in_R \mathcal{S}_{\mathbf{x}\cup\{x\}}$.

It is obvious that in both distributions, $\mathcal{C}$ is chosen uniformly from $\mathcal{S}_{\mathbf{x}}$. It is easy to see that for any point $x \notin span(\mathbf{x})$ the pair $(\mathcal{C}, x)$ is equally likely. Since $\Pr_{x\in\mathcal{C}}(x \in span(\mathbf{x})) = \frac{|span(\mathbf{x})|}{|\mathcal{F}|^{D+3}} \leq \frac{1}{|\mathcal{F}|^3} \ll \rho^{\frac{1}{6}}$, this only adds a negligible error. ■

5.2 A Proposition

Proposition 28 (Linear Combinations) *Let $k \leq O(\log_{|\mathcal{F}|} n)$. There exists a polynomial algorithm that inputs a system of conjunctions, Ψ, and outputs a system of* equations, *rather of conjunctions, $\widetilde{\Psi}$, with the same depend and the same degree, such that*

$$\omega(\Psi) \leq \omega(\widetilde{\Psi}) \leq \omega(\Psi) + \frac{1}{|\mathcal{F}|}$$

Proof: We take 'random' linear-combinations of the equations in each conjunction: We replace every conjunction

$(\psi_{i_1}) \wedge \cdots \wedge (\psi_{i_k})$ by the following $|\mathcal{F}|^k$ equations

$$\{\sum_{j=1}^{k} \alpha_j \psi_{i_j} \ : \ \alpha_1, .., \alpha_k \in \mathcal{F}\}$$

where a linear combination of equations is defined in the obvious way.

If the conjunction evaluates to false, then no more than $\frac{1}{|\mathcal{F}|}$ of the equations will sum up to zero. Hence the error probability increases by no more than $\frac{1}{|\mathcal{F}|}$.

Note that using an efficient linear code (rather than going through all the linear combinations) gives a more size-efficient result. This is not necessary for our purposes. ■

References

[ABMP98] M. Alekhnovich, S. Buss, S. Moran, and T. Pitassi. Minimum propositional proof length is NP-hard to linearly approximate. Manuscript, 1998.

[ALM+92] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and intractability of approximation problems. In *Proc. 33rd IEEE Symp. on Foundations of Computer Science*, pages 13–22, 1992.

[AS92] S. Arora and S. Safra. Probabilistic checking of proofs: A new characterization of NP. In *Proc. 33rd IEEE Symp. on Foundations of Computer Science*, pages 2–13, 1992.

[AS97] Sanjeev Arora and Madhu Sudan. Improved low degree testing and its applications. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 485–495, El Paso, Texas, 4–6 May 1997.

[BFL91] L. Babai, L. Fortnow, and C. Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.

[BGLR93] M. Bellare, S. Goldwasser, C. Lund, and A. Russell. Efficient multi-prover interactive proofs with applications to approximation problems. In *Proc. 25th ACM Symp. on Theory of Computing*, pages 113–131, 1993.

[BGS98] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, PCPs, and nonapproximability—towards tight results. *SIAM Journal on Computing*, 27(3):804–915, June 1998.

[DS98] I. Dinur and S. Safra. Monotone-minimum-satisfying assignment is NP-hard for almost polynomial factors. Manuscript, 1998.

[FGL+91] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Approximating clique is almost NP-complete. In *Proc. 32nd IEEE Symp. on Foundations of Computer Science*, pages 2–12, 1991.

[FL92] U. Feige and L. Lovász. Two-prover one-round proof systems: Their power and their problems. In *Proc. 24th ACM Symp. on Theory of Computing*, pages 733–741, 1992.

[Hås97] Johan Håstad. Some optimal inapproximability results. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 1–10, El Paso, Texas, 4–6 May 1997.

[HPS93] J. Hastad, R. Phillips, and S. Safra. A well-characterized approximation problem. *Information Processing Letters*, 47:301–305, 1993.

[LS91] D. Lapidot and A. Shamir. Fully parallelized multi prover protocols for NEXPTIME. In *Proc. 32nd IEEE Symp. on Foundations of Computer Science*, pages 13–18, 1991.

[LY94] Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41(5):960–981, 1994.

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, June 1998.

[RS97] 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.

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'