

DEEP-FRI: Sampling Outside the Box Improves Soundness

Eli Ben-Sasson* Lior Goldberg* Swastik Kopparty† Shubhangi Saraf‡

Abstract

Motivated by the quest for scalable and succinct zero knowledge arguments, we revisit worst-case-to-average-case reductions for linear spaces, raised by [Rothblum, Vadhan, Wigderson, STOC 2013]. The previous state of the art by [Ben-Sasson, Kopparty, Saraf, CCC 2018] showed that if some member of an affine space U is δ -far in relative Hamming distance from a linear code V — this is the worst-case assumption — then most elements of U are almost- δ -far from V — this is the average case. However, this result was known to hold only below the “double Johnson” function of the relative distance δ_V of the code V , i.e., only when $\delta < 1 - (1 - \delta_V)^{1/4}$.

First, we increase the soundness-bound to the “one-and-a-half Johnson” function of δ_V and show that the average distance of U from V is nearly δ for any worst-case distance δ smaller than $1 - (1 - \delta_V)^{1/3}$. This bound is tight, which is somewhat surprising because the one-and-a-half Johnson function is unfamiliar in the literature on error correcting codes.

To improve soundness further for Reed Solomon codes we sample outside the box. We suggest a new protocol in which the verifier samples a single point z outside the box D on which codewords are evaluated, and asks the prover for the value at z of the interpolating polynomial of a random element of U . Intuitively, the answer provided by the prover “forces” it to choose one codeword from a list of “pretenders” that are close to U . We call this technique Domain Extending for Eliminating Pretenders (DEEP).

The DEEP method improves the soundness of the worst-case-to-average-case reduction for RS codes up their list decoding radius. This radius is bounded from below by the Johnson bound, implying average distance is approximately δ for all $\delta < 1 - (1 - \delta_V)^{1/2}$. Under a plausible conjecture about the list decoding radius of Reed-Solomon codes, average distance from V is approximately δ for all δ . The DEEP technique can be generalized to all linear codes, giving improved reductions for capacity-achieving list-decodable codes.

Finally, we use the DEEP technique to devise two new protocols:

- An Interactive Oracle Proof of Proximity (IOPP) for RS codes, called DEEP-FRI. This soundness of the protocol improves upon that of the FRI protocol of [Ben-Sasson et al., ICALP 2018] while retaining linear arithmetic proving complexity and logarithmic verifier arithmetic complexity.
- An Interactive Oracle Proof (IOP) for the Algebraic Linking IOP (ALI) protocol used to construct zero knowledge scalable transparent arguments of knowledge (ZK-STARKs) in [Ben-Sasson et al., eprint 2018]. The new protocol, called DEEP-ALI, improves soundness of this crucial step from a small constant $< 1/8$ to a constant arbitrarily close to 1.

*StarkWare Industries Ltd. {eli,lior}@starkware.co

†Department of Mathematics and Department of Computer Science, Rutgers University. Research supported in part by NSF grants CCF-1253886, CCF-1540634, CCF-1814409 and CCF-1412958, and BSF grant 2014359. Some of this research was done while visiting the Institute for Advanced Study. swastik.kopparty@gmail.com

‡Department of Mathematics and Department of Computer Science, Rutgers University. Research supported in part by NSF grants CCF-1350572, CCF-1540634 and CCF-1412958, BSF grant 2014359, a Sloan research fellowship and the Simons Collaboration on Algorithms and Geometry. Some of this research was done while visiting the Institute for Advanced Study. shubhangi.saraf@gmail.com

1 Introduction

Arithmetization is a marvelous technique that can be used to reduce problems in computational complexity, like verifying membership in a nondeterministic language, to questions about membership of vectors in algebraic codes like Reed-Solomon (RS) and Reed-Muller (RM) codes [Raz87, LFKN92]. One of the end-points of such a reduction is the RS proximity testing (RPT) problem. It is a problem of inherent theoretical interest, but also of significant practical importance because it is used in recent constructions of succinct zero knowledge (ZK) arguments including Ligerio [AHIV17], Aurora [BCR⁺18], and Scalable Transparent ARguments of Knowledge (ZK-STARKs) [BBHR18a]. We discuss this connection after describing the problem and our results.

In the RPT problem a verifier is given oracle access to a function $f : D \rightarrow \mathbb{F}$, we call $D \subset \mathbb{F}$ the *evaluation domain*, and is tasked with distinguishing between the “good” case that f is a polynomial of degree at most d and the “bad” case in which f is δ -far in relative Hamming distance from all degree- d polynomials. To achieve succinct verification time, poly-logarithmic in d , we must allow the verifier some form of interaction with a prover — the party claiming that $\deg(f) \leq d$. Initially, this interaction took the form of oracle access to a probabilistically checkable proof of proximity (PCPP) [BGH⁺06] provided by the prover in addition to f . Indeed, in this model the RPT problem can be “solved” with PCPPs of quasilinear size $|D|\text{poly log } |D|$, constant query complexity and constant soundness [BS08, Din07]. However, the concrete complexity of prover time, verifier time and communication complexity are rather large, even when considering practical settings that involve moderately small instance sizes.

To improve prover, verifier, and communication complexity for concrete (non-asymptotic) size problems, the Interactive Oracle Proofs of Proximity (IOPP) model is more suitable [RRR16, BCS16, BCF⁺16]. This model can be viewed as a multi-round PCPP. Instead of having the prover write down a single proof π , in the IOPP setting the proof oracle is produced over a number of rounds of interaction, during which the verifier sends random bits and the prover responds with additional (long) messages to which the verifier is allowed oracle access. The additional rounds of interaction allow for a dramatic improvement in the asymptotic and concrete complexity of solving the RPT problem. In particular, the Fast RS IOPP (FRI) protocol of [BBHR18b] has linear prover arithmetic complexity, logarithmic verifier arithmetic complexity and constant soundness. Our goal here is to improve soundness of this protocol and to suggest better protocols in terms of soundness in the high-error regime (also known as the “list decoding” regime).

Soundness analysis of FRI reduces to the following natural “worst-case-to-average-case” question regarding linear spaces, which is also independently very interesting for the case of general (non-RS) codes. This question was originally raised in a different setting by [RVW13] and we start by discussing it for general linear codes before focusing on the special, RS code, case.

1.1 Maximum distance vs average distance to a linear code

Suppose that $U \subset \mathbb{F}^D$ is a “line”, a 1-dimensional¹ affine space over \mathbb{F} . Let $u^* \in \mathbb{F}^D$ denote the origin of this line and u be its slope, so that $U = \{u_x = u^* + xu \mid x \in \mathbb{F}\}$. For a fixed linear space $V \subset \mathbb{F}^D$, pick u^* to be the element in U that is farthest from V , denoting by δ_{\max} its relative Hamming distance (from V). This is our worst-case assumption. Letting $\delta_x = \Delta(u_x, V)$ where Δ

¹The generalization of our results to spaces U of dimension > 1 is straightforward by partitioning U into lines through u^* and applying these results to each line.

denotes relative Hamming distance, what can be said about the *expected* distance $\mathbf{E}_{x \in \mathbb{F}}[\delta_x]$ of u_x from V ?

Rothblum, Vadhan and Wigderson showed that $\mathbf{E}_x[\delta_x] \geq \frac{\delta_{\max}}{2} - o(1)$ for all spaces U and V , where, here and below, $o(1)$ denotes negligible terms that approach 0 as $|\mathbb{F}| \rightarrow \infty$ [RVW13]. A subset of the co-authors of this paper improved this to $\mathbf{E}[\delta_x] \geq 1 - \sqrt{1 - \delta_{\max}} - o(1)$, showing the average distance scales roughly like the Johnson list-decoding function of δ_{\max} , where $J(x) := 1 - \sqrt{1 - x}$ [BKS18a]. In both of these bounds the expected distance is strictly smaller than δ_{\max} . However, the latter paper also showed that when V is a (linear) error correcting code with large relative distance δ_V , if δ_{\max} is smaller than the “double Johnson” function of δ_V , given by $J^{(2)}(x) := J(J(x))$, then the average distance hardly deteriorates,

$$\mathbf{E}[\delta_x] \geq \min\left(\delta_{\max}, J^{(2)}(\delta_V)\right) - o(1) = \min\left(\delta_{\max}, 1 - \sqrt[4]{1 - \delta_V}\right) - o(1) \quad (1)$$

and the equation above summarizes the previous state of affairs on this matter.

Our first result is an improvement of Equation (1) to the “one-and-a-half-Johnson” function $J^{(1.5)}(x) = 1 - (1 - x)^{1/3}$. Lemma 3.1 says that for codes V of relative Hamming distance δ_V ,

$$\mathbf{E}[\delta_x] \geq \min\left(\delta_{\max}, J^{(1.5)}(\delta_V)\right) - o(1) = \min\left(\delta_{\max}, 1 - \sqrt[3]{1 - \delta_V}\right) - o(1). \quad (2)$$

Our second result shows that Equation (2) is tight, even for the special case of V being an RS code. We find this result somewhat surprising because the $J^{(1.5)}(x)$ function is not known to be related to any meaningful coding theoretic notion. The counter-example showing the tightness of Equation (2) arises for very special cases, in which (i) \mathbb{F} is a binary field (of characteristic 2), (ii) the rate ρ is precisely $1/8 = 2^{-3}$ and, most importantly, (iii) the evaluation domain D equals all of \mathbb{F} (see Section 3.1). Roughly speaking, the counter-example uses functions $u^*, u : \mathbb{F}_{2^n} \rightarrow \mathbb{F}_{2^n}$ that are $3/4 = 1 - \rho^{2/3}$ -far from polynomials of degree $\rho 2^n$ yet *pretend* to be low-degree because for all $x \in \mathbb{F}_{2^n} \setminus \{0\}$ the function $u^* + xu$ is $1/2 = \sqrt[3]{\rho}$ -close to a polynomial of degree $\rho 2^n$. See Lemma 3.3 for details.

Our next set of results, which we discuss below, show how to go beyond the above limitation through a new interactive proximity proving technique.

1.2 Domain Extension for Eliminating Pretenders (DEEP)

The case that interests us most is when V is an RS code (although we will return to the discussion of general linear codes later). Henceforth, the RS code of rate ρ evaluated over D is

$$\text{RS}[\mathbb{F}, D, \rho] := \{f : D \rightarrow \mathbb{F} \mid \deg(f) < \rho|D|\}.$$

RS codes are maximum distance separable (MDS), meaning that $\delta_V = 1 - \rho$ and so Equation (2) simplifies to

$$\mathbf{E}[\delta_x] \geq \min(\delta_{\max}, 1 - \sqrt[3]{\rho}) - o(1). \quad (3)$$

This improved bound can be translated, using some extra work, to FRI soundness analysis with similar guarantees. Specifically, Equation (3) implies that for $f : D \rightarrow \mathbb{F}$ that is δ -far from $\text{RS}[\mathbb{F}, D, \rho]$, the soundness error of a single invocation of the FRI QUERY test (which requires $\log |D|$ queries) is at most $\max\{1 - \delta, \sqrt[3]{\rho}\}$, and this can be plugged into ZK-STARKs like [BBHR18a] and ZK-SNARGs like Aurora [BCR⁺18]. Roughly speaking, if the rejection probability is of δ -far words

is $\max(\delta, \delta_0)$ then to reach soundness error less than $2^{-\lambda}$ for codes of blocklength n , communication complexity (and verifier complexity) scale roughly like $\frac{\lambda}{\log \delta_0} \cdot c \cdot \log n$ for some constant c . Thus, the improvement from Equation (1) to Equation (2) translates to a 25% reduction in verifier complexity (from $\frac{4\lambda}{\log \rho} \cdot c \cdot \log n$ to $\frac{3\lambda}{\log \rho} \cdot c \cdot \log n$).

To break the soundness bound of Equation (2) and thereby further reduce verifier complexity in the afore-mentioned systems, we suggest a new method. We discuss it first for RS codes, then generalize to arbitrary linear codes. If $u^*, u : D \rightarrow \mathbb{F}$ are indeed the evaluation of two degree d polynomials, say, P^* and P , our verifier will artificially *extend the domain* D to a larger one \bar{D} , sample uniformly $z \in \bar{D}$ and ask for the evaluation of $P^*(z)$ and $P(z)$. The answers provided by the prover can now be applied to modify each of u^* and u in a *local* manner to reflect the new knowledge, and along the way also prune down the large list of polynomials which u^* and u might pretend to be. If $\alpha_z^* = P^*(z), \alpha_z = P(z)$ are the honest prover's answers to the query z , then $(X - z)$ divides $P^*(X) - \alpha_z^*$ and likewise $(X - z) | P(X) - \alpha_z$. Letting $\alpha_x = \alpha^* + x\alpha$ and $P_x(X) = P^*(X) + xP(X)$ it follows that $(X - z) | P_x(X) - \alpha_x$. Consider now the soundness of this procedure. In the extreme case that u^* has a small list of polynomials that, each, somewhat agree with it, then with high probability over z , any answer provided by the prover will agree with at most one of the polynomials in this list. The proof of our main technical result, Theorem 4.1, formalizes this intuition. For radius δ , let L_δ^* be the maximal list size,

$$L_\delta^* = \max_{u^* \in \mathbb{F}^D} |\{v \in V \mid \Delta(u^*, v) < \delta\}|$$

where Δ denotes relative Hamming distance. Let $V|_{u_x(z)=\alpha_x}$ be the restriction of V to codewords that are evaluations of polynomials of degree at most d that, additionally, evaluate to α_x on z . Our main Theorem 4.1 shows that if $\Delta(u^*, V) = \delta_{\max}$ then for any pair of answers α_z^*, α_z given in response to query z ,

$$\mathbf{E}_{z,x} [\Delta(u_x, V|_{u_x(z)=\alpha_x})] \geq \delta_{\max} - L_\delta^* \cdot \left(\frac{\rho|D|}{|\bar{D}|} \right)^{1/3} - o(1). \quad (4)$$

The Johnson bound (Theorem 2.2) says that when $\delta < J(1 - \rho) = 1 - \sqrt{\rho}$ we have $L_\delta^* = O(1)$ and this improves the worst-case-to-average-case result from that of Equation (2) to a bound that matches the Johnson bound:

$$\mathbf{E}_{z,x} [\Delta(u_x, V|_{u_x(z)=\alpha_x})] \geq \min(\delta_{\max}, J(\delta_V)) - o(1) = \min(\delta_{\max}, \sqrt{\rho}) - o(1). \quad (5)$$

The exact behavior of the list size of Reed-Solomon codes beyond the Johnson bound is a famous open problem. It may be the case that the list size is small for radii far greater than the Johnson bound; in fact, for most domains D this is roughly known to hold [RW14]. If it holds that that list sizes are small all the way up to radius equal to the distance $\delta_V = 1 - \rho$ (i.e., if Reed-Solomon codes meet list-decoding capacity), then Equation (5) implies that the technique suggested here has optimal soundness for (nearly) all distance parameters.

Generalization to arbitrary linear codes The DEEP method can be used to improve worst-to-average-case reductions for general linear codes. Viewing codewords in V as evaluations of linear forms of a domain D , we ask for the evaluation of the linear forms that supposedly correspond to u^* and u on a random location $z \in \bar{D}$ where $|\bar{D}| \gg |D|$. Lemma 4.6 generalizes Theorem 4.1 and says that if V has near-capacity list-decoding radius (with small list size) and \bar{D} corresponds to

(columns of a generating matrix of) a good error correcting code, then we have $\mathbf{E}_x[\delta_x] \approx \delta_{\max}$. The main difference between the RS case and that of general linear codes is that in the former, the prover-answers $\alpha^*(x), \alpha(x)$ can be processed to modify locally the entries of u_x to reduce the degree of the resulting function; this is something we cannot carry out (to best of our understanding) for all linear codes.

1.3 DEEP-FRI

Applying the technique of domain extension for eliminating pretenders to the FRI protocol requires a modification that we discuss next. The FRI protocol can be described as a process of “randomly folding” an (inverse) Fast Fourier Transform (iFFT) computation. In the “classical” iFFT, one starts with a function $f^{(0)} : \langle \omega \rangle \rightarrow \mathbb{F}$ where ω generates a multiplicative group of order 2^k for integer k . The iFFT computes (in arithmetic complexity $O(k2^k)$) the interpolating polynomial $\tilde{f}(X)$ of the function f . This computation follows by computing (in linear time) a pair of functions $f_0, f_1 : \langle \omega^2 \rangle \rightarrow \mathbb{F}$, recalling $|\langle \omega^2 \rangle| = \frac{1}{2}|\langle \omega \rangle|$. Their interpolants \tilde{f}_0, \tilde{f}_1 are then used to compute in linear time the original interpolant \tilde{f} of f .

As explained in [BBHR18b], in the FRI protocol the prover first commits to f as above. Then the verifier samples a random $x^{(0)} \in \mathbb{F}$ and the protocol continues with the single function $f^{(1)} : \langle \omega^2 \rangle \rightarrow \mathbb{F}$ which is supposedly $f^{(1)} := f_0 + x^{(0)}f_1$. It turns out that if f is indeed of degree less than $\rho|\langle \omega \rangle|$ then for all x we have that $f^{(1)}$ is of degree less than $\rho|\langle \omega^2 \rangle|$ as well. The tricky part is showing that when f is δ -far from $\text{RS}[\mathbb{F}, \langle \omega \rangle, \rho]$ this also holds with high probability (over x) for $f^{(1)}$ and some δ' that is as close as possible to δ . (One can show that invariably we have $\delta' \leq \delta$, i.e., the green line of Figure 1 is an upper bound on soundness of both FRI and the new DEEP-FRI protocol described below.)

The worst-case-to-average-case results of Equation (2) and Lemma 3.1 can be converted to similar improvements for FRI, showing that for $\delta < 1 - \sqrt[3]{\rho}$ we have $\delta' \approx \delta$. This follows directly from the techniques of [BKS18a, Section 7] (see the red line in Figure 1). But to use the new DEEP technique of Equation (4) and Theorem 4.1 in order to improve soundness of an RS-IOPP, we need to modify the FRI protocol, leading to a new protocol that is aptly called DEEP-FRI. Instead of constructing $f^{(1)}$ directly, our verifier first samples $z^{(0)} \in \mathbb{F}$ and queries the prover for the evaluation of the interpolant of $f^{(0)}$ on $z^{(0)}$ and $-z^{(0)}$. After the answers $\alpha_{z^{(0)}}, \alpha_{-z^{(0)}}$ have been recorded, the verifier proceeds by sampling $x^{(0)}$ and expects the prover to submit $f^{(1)}$ which is the linear combination of f'_0, f'_1 derived from the modification f' of f that takes into account the answers $\alpha_{z^{(0)}}, \alpha_{-z^{(0)}}$. Assuming \tilde{f} is the interpolant of f , an honest prover would set $f'(X) := (\tilde{f}(X) - U(X))/Z(X)$ where $U(X)$ is the degree ≤ 1 polynomial that evaluates to $\alpha_{z^{(0)}}$ on $z^{(0)}$ and to $\alpha_{-z^{(0)}}$ on $-z^{(0)}$ and $Z(X)$ is the monic degree 2 polynomial whose roots are $z^{(0)}$ and $-z^{(0)}$. As shown in Section 5, the soundness bounds of Equation (4) and Theorem 4.1 now apply to DEEP-FRI. This shows that the soundness of DEEP-FRI, i.e., the rejection probability of words that are δ -far from $\text{RS}[\mathbb{F}, D, \rho]$ is roughly δ for any δ that is smaller than the maximal radius for which list-sizes are “small”. Figure 1 summarizes the results described here.

1.4 DEEP Algebraic Linking IOP (DEEP-ALI)

In Section 1.3 we only discussed results improving the soundness of Reed-Solomon Proximity Testing (RPT). We now discuss how to improve the soundness of IOP-based argument systems (such as [BBHR18a, BCR+18]) that use RPT solutions. In order to reap the benefits of the improved




Figure 1: FRI and DEEP-FRI soundness threshold δ_0 as a function of RS code rate ρ , for a single invocation of the QUERY phase, as field size $q \rightarrow \infty$. $\delta_0(\rho)$ is defined to be the largest distance parameter δ for which soundness (rejection probability) of a single invocation of the FRI/DEEP-FRI QUERY is $\delta - o(1)$. Higher lines are better. The top line is the trivial upper bound on soundness which applies to both FRI and DEEP-FRI; the bottom line is the soundness of the original analysis of [BBHR18b]. Dashed lines represent prior results. The red line is the (tight) soundness lower bound for FRI and the blue line is a lower bound on DEEP-FRI soundness. Under a plausible conjecture for Reed-Solomon list-decodability (Conjecture 2.3), the actual soundness is as high as the green line.

soundness of RPT, we need reductions that produce instances of the RPT problem *that are very far from the relevant RS code* when the input instance is unsatisfiable. One such protocol is the Algebraic Linking IOP (ALI) of [BBHR18a]. The instances of the RPT problem derived from an unsatisfiable instance in ALI are proven to be somewhat far from low-degree but the distance bound proved in that paper is less than $1/8$, even when used with RS codes of negligible rate ρ (nevertheless it is conjectured and assumed in both ZK-STARK [BBHR18a] and Aurora [BCR⁺18] that the distance is significantly greater). In Section 6 we use the DEEP technique to modify the ALI protocol in a manner similar to the DEEP-FRI modification. The result of this modification allows us to apply the soundness results of Equation (4) to the DEEP-ALI protocol and show that, when provided with unsatisfiable instances, the distance of the received words that result from that protocol is provably at least $1 - \sqrt{\rho} - o(1)$ (and may be greater, assuming more favourable bounds on the list decoding radius for RS codes, as in Conjecture 2.3).

Organization of the rest of the paper Section 2 presents general notation. Section 3 gives an improved worst-to-average case reductions for general spaces and shows that the bound in the reduction is tight (Lemma 3.3). Section 4 presents our main technical result, showing that the DEEP method improves worst-to-average case reductions for RS codes up to the Johnson bound (provably) and perhaps even beyond. Section 5 presents the DEEP-FRI protocol that obtains better soundness than the state of the art FRI protocol, and Section 6 discusses the DEEP-ALI protocol.

2 Preliminaries

Functions For a set D , we will be working with the space of functions $u : D \rightarrow \mathbb{F}$, denoted \mathbb{F}^D . For $u \in \mathbb{F}^D$ we use $u(z)$ to denote the z th entry of u , for $z \in D$. For $C \subset D$ we use $f|_C$ to denote the restriction of f to C . For two functions $f, g : D \rightarrow \mathbb{F}$ we write $f = g$ when the two functions are equal as elements in \mathbb{F}^D and similarly say $f|_C = g|_C$ when their restrictions are equal as elements in \mathbb{F}^C .

Distance We use $\Delta_D(u, v) = \Pr_{z \in D} [u(z) \neq v(z)]$ for relative Hamming distance, and omit D when it is clear from context. For a set $S \subset \mathbb{F}^D$ we use $\Delta_D(v, S) = \min_{s \in S} \Delta_D(v, s)$ and $\Delta_D(S) = \min_{s \neq s' \in S} \Delta_D(s, s')$ denotes the minimal relative distance of S . For $u \in \mathbb{F}^D$ let $B(u, \delta)$ denote the Hamming ball in \mathbb{F}^D of normalized radius δ centered at u ,

$$B(u, \delta) = \{u' \in \mathbb{F}^D \mid \Delta_D(u, u') < \delta\}.$$

Linear codes An $[n, k, d]_q$ -linear error correcting code is a linear space $V \subset \mathbb{F}_q^n$ of dimension k over \mathbb{F}_q with minimal Hamming distance d . A generating matrix for V is a matrix $G \in \mathbb{F}_q^{n \times k}$ of rank k such that $V = \{Gx \mid x \in \mathbb{F}_q^k\}$.

Polynomials and RS codes The *interpolant* of $f : D \rightarrow \mathbb{F}_q$ is the unique polynomial of degree $< |D|$ whose evaluation on D is f . The degree of f , denoted $\deg(f)$, is the degree of its interpolant. The RS code evaluated over domain $D \subset \mathbb{F}$ and rate ρ is denoted $\text{RS}[\mathbb{F}, D, \rho] = \{f : D \rightarrow \mathbb{F} \mid \deg(f) < \rho|D|\}$. Sometimes it will be more convenient to work with degree rather than rate, in which case we abuse notation and define $\text{RS}[\mathbb{F}, D, d] = \{f : D \rightarrow \mathbb{F} \mid \deg(f) < d\}$. We use capital letters like P, Q to denote polynomials and when we say $P \in \text{RS}[\mathbb{F}, D, \rho]$ we mean that $\deg(P) < \rho|D|$ and associate P with the RS codeword that is its evaluation on D . We also use \tilde{f} to denote the interpolant of a function f .

2.1 List Decoding

Definition 2.1 (List size for Reed-Solomon Codes). For $u \in \mathbb{F}^D$, a set $V \subset \mathbb{F}^D$, and distance parameter $\delta \in [0, 1]$, let $\text{List}(u, V, \delta)$ be the set of elements in V that are at most δ -far from u in relative Hamming distance. Formally, using $B(u, \delta)$ to denote the Hamming ball of relative radius δ centered around u , we have $\text{List}(u, V, \delta) = B(u, \delta) \cap V$.

The code V is said to be (δ, L) -list-decodable if $|\text{List}(u, V, \delta)| \leq L$ for all $u \in \mathbb{F}_q^D$.

For $D \subseteq \mathbb{F}_q$, let $\mathcal{L}(\mathbb{F}_q, D, d, \delta)$ be the maximum size of $\text{List}(u, V, \delta)$ taken over all $u \in \mathbb{F}_q^D$ for $V = \text{RS}[\mathbb{F}_q, D, \rho = d/|D|]$.

We recall the fundamental Johnson bound, which says that sets with large minimum distance have nontrivial list-decodability. The particular version below follows, e.g., from [Gur07, Theorem 3.3] by setting $d = (1 - \rho)|D|$ and $e = (1 - \sqrt{\rho} - \varepsilon)|D|$ there.

Theorem 2.2 (Johnson bound). Let $V \subset \mathbb{F}^D$ be a code with minimum relative distance $1 - \rho$, for $\rho \in (0, 1)$. Then V is $(1 - \sqrt{\rho} - \varepsilon, 1/(2\varepsilon\sqrt{\rho}))$ -list-decodable for every $\varepsilon \in (0, 1 - \sqrt{\rho})$.

In particular, for Reed-Solomon codes this implies the following list-decodability bound:

$$\mathcal{L}(\mathbb{F}_q, D, d = \rho|D|, 1 - \sqrt{\rho} - \varepsilon) \leq O\left(\frac{1}{\varepsilon\sqrt{\rho}}\right).$$

Extremely optimistically, we could hope that Reed-Solomon codes are list-decodable all the way up to their distance with moderate list sizes. Staying consistent with the known limitations [BSKR10], we have the following brave conjecture.

Conjecture 2.3 (List decodability of Reed-Solomon Codes up to Capacity). *For every $\rho > 0$, there is a constant C_ρ such that every Reed-Solomon code of length n and rate ρ is list-decodable from $1 - \rho - \varepsilon$ fraction errors with list size $\left(\frac{n}{\varepsilon}\right)^{C_\rho}$. That is:*

$$\mathcal{L}(\mathbb{F}_q, D, d = \rho|D|, 1 - \rho - \varepsilon) \leq \left(\frac{|D|}{\varepsilon}\right)^{C_\rho}.$$

3 Improved High-error Distance Preservation

Our first result gives better distance preservation results for linear codes V of relative distance λ . The previous state-of-the-art [BKS18a] said that when a 1-dimensional affine space U contains some element u^* that is $\delta_{\max} = \Delta(u^*, V)$ far from V , then

$$\mathbf{E}_{u \in U}[\Delta(u, V)] \geq \min(\delta_{\max}, 1 - J^{(2)}(\lambda)) - o(1).$$

The following lemma improves the average-case distance to

$$\mathbf{E}_{u \in U}[\Delta(u, V)] \geq \min(\delta_{\max}, 1 - J^{(1.5)}(\lambda)) - o(1).$$

Later on, in Section 3.1, we will show that this result is tight (for a sub-family of RS codes).

Lemma 3.1 (One-and-half Johnson distance preservation). *Let $V \subseteq \mathbb{F}_q^n$ be a linear code of distance $\lambda = \Delta(V)$. Let $\epsilon, \delta > 0$ with $\epsilon < 1/3$ and $\delta < 1 - (1 - \lambda + \epsilon)^{1/3}$.*

Suppose $u^ \in \mathbb{F}_q^n$ is such that $\Delta(u^*, V) > \delta + \epsilon$. Then for all $u \in \mathbb{F}_q^n$, there are at most $2/\epsilon^2$ values of $x \in \mathbb{F}_q$ such that $\Delta(u^* + xu, V) < \delta$.*

This result is the contra-positive statement of the following, more informative, version of it, that we prove below.

Lemma 3.2 (One-and-half Johnson distance preservation — positive form). *Let $V \subseteq \mathbb{F}_q^D$ be a linear code of distance $\lambda = \Delta(V)$. Let $\epsilon, \delta > 0$ with $\epsilon < 1/3$ and $\delta < 1 - (1 - \lambda + \epsilon)^{1/3}$. Let $u, u^* \in \mathbb{F}_q^D$ satisfy*

$$\Pr_{x \in \mathbb{F}_q} [\Delta(u^* + xu, V) < \delta] \geq \frac{2}{\epsilon^2 q}. \tag{6}$$

Then there exist $v, v^ \in V$ and $C \subseteq D$ such that the following three statements hold simultaneously:*

- $|C| \geq (1 - \delta - \epsilon)|D|$,
- $u|_C = v|_C$, and

- $u^*|_C = v^*|_C$.

Observe that if u, u^* satisfy Equation (6) then the v, v^*, C deduced by Lemma 3.2 have the property that for all $x \in \mathbb{F}_q$, we have $\Delta(u^* + xu, V) \leq \delta + \epsilon$. In other words, the existence of v, v^* and C almost completely explains Equation (6).

Proof. Let $u_x = u^* + xu$. Let

$$A = \{x \mid \Delta(u^* + xu, V) < \delta\}.$$

For each $x \in A$, let $v_x \in V$ be an element of V that is closest to u_x , and let $S_x \subseteq D$ be the agreement set of u_x and v_x , defined as $S_x = \{y \in D \mid u_x(y) = v_x(y)\}$.

For x, β, γ picked uniformly from A and y picked uniformly from D , we have:

$$\begin{aligned} \mathbf{E}_{x,\beta,\gamma}[|S_x \cap S_\beta \cap S_\gamma|/n] &= \mathbf{E}_{y,x,\beta,\gamma}[\mathbf{1}_{y \in S_x \cap S_\beta \cap S_\gamma}] \\ &= \mathbf{E}_y[\mathbf{E}_x[\mathbf{1}_{y \in S_x}]^3] \\ &\geq \mathbf{E}_{y,x}[\mathbf{1}_{y \in S_x}]^3 \\ &\geq (1 - \delta)^3 \\ &> 1 - \lambda + \epsilon. \end{aligned}$$

The second equality above follows from the independence of the events $y \in S_x, y \in S_\beta, y \in S_\gamma$ given $y \in D$. The first inequality is Jensen's and the last inequality is by assumption on δ, γ, ϵ .

Thus

$$\Pr_{x,\beta,\gamma}[|S_x \cap S_\beta \cap S_\gamma| \geq (1 - \lambda)|D|] \geq \epsilon.$$

Note that $\Pr_{x,\beta,\gamma}[x, \beta, \gamma \text{ are not all distinct}] < 3/|A|$. Since $|A| \geq 2/\epsilon^2 > \frac{6}{\epsilon}$, we have that $3/|A| \leq \epsilon/2$ and hence x, β, γ are all distinct with probability at least $1 - \epsilon/2$. Thus with probability at least $\epsilon/2$ over the choice of x, β, γ , we have that x, β, γ are all distinct and $|S_x \cap S_\beta \cap S_\gamma| > (1 - \lambda)|D|$.

This means that there are distinct x_0, β_0 such that

$$\Pr_\gamma[|S_{x_0} \cap S_{\beta_0} \cap S_\gamma| > (1 - \lambda)|D|] \geq \epsilon/2.$$

Fix a γ where this happens. Let $S = S_{x_0} \cap S_{\beta_0} \cap S_\gamma$. We have that

$$(x_0, u_{x_0}), (\beta_0, u_{\beta_0}), (\gamma, u_\gamma)$$

are collinear. Thus

$$(x_0, u_{x_0}|_S), (\beta_0, u_{\beta_0}|_S), (\gamma, u_\gamma|_S)$$

are all collinear. By definition of S , we get that:

$$(x_0, v_{x_0}|_S), (\beta_0, v_{\beta_0}|_S), (\gamma, v_\gamma|_S)$$

are all collinear. Since $|S| > (1 - \lambda)|D|$ (and recalling that λ is the distance of V), we get that v_γ is determined by $v_\gamma|_S$ via a linear map. This means that

$$(x_0, v_{x_0}), (\beta_0, v_{\beta_0}), (\gamma, v_\gamma)$$

are all collinear.

Thus $\epsilon/2$ -fraction of the $\gamma \in A$ have the “good” property that (γ, v_γ) is on the line passing through (x_0, v_{x_0}) and (β_0, v_{β_0}) . Write this line as $v^* + xv$ and notice that for all “good” γ we have $v_\gamma = v^* + \gamma v$. Let $A' \subseteq A$ denote the set of good elements for this line, recording that $|A'| \geq |A| \cdot \epsilon/2 \geq 1/\epsilon$.

Thus for $x \in A'$, $\Delta(u^* + xu, v^* + xv) < \delta$.

Consider the set $C \subset D$ defined by

$$C = \{y \in D \mid u^*(y) = v^*(y) \text{ AND } u(y) = v(y)\}.$$

For each $y \in D \setminus C$ there exists at most a single value of $x \in \mathbb{F}_q$ satisfying $u^*(y) + x \cdot u(y) = v^*(y) + x \cdot v(y)$ because

$$(u^*(y) - v^*(y)) + x \cdot (u(y) - v(y))$$

has at most one value x on which it vanishes.

This implies

$$\delta \geq \mathbf{E}_{x \in A'}[\Delta_D(u_x, v_x)] \geq \frac{|D \setminus C|}{|D|} \cdot \left(1 - \frac{1}{|A'|}\right) \geq \left(1 - \frac{|C|}{|D|}\right) \cdot (1 - \epsilon) \geq 1 - \frac{|C|}{|D|} - \epsilon.$$

Rearranging, we get $\frac{|C|}{|D|} \geq 1 - (\delta + \epsilon)$ and this completes the proof. \square

3.1 Tightness of the one-and-a-half Johnson bound

Lemma 3.1 says that when V is a linear code with minimum distance λ , and u^* is some element that is δ -far from V , then for any u we have with high probability

$$\Delta(u^* + xu, V) \geq \min(\delta, J^{(1.5)}(\lambda) = 1 - (1 - \lambda)^{1/3}).$$

The rightmost term seems quite strange, as the $J^{(1.5)}(\cdot)$ function is unfamiliar in other settings of coding theory. However, as we show next, in certain settings this function gives the correct bound!

Lemma 3.3 (Tightness of one-and-a-half Johnson bound). *For every member V_n of following family of RS codes $\{V_n = \text{RS}[\mathbb{F}_{2^n}, \mathbb{F}_{2^n}, \rho = 2^{-3}] \mid n \in \mathbb{N}\}$ there exist $u_n^*, u_n \in \mathbb{F}_{2^{2n}}^{\mathbb{F}_{2^n}}$ satisfying the following:*

- $\delta_{\max} \triangleq \Delta(u_n^*, V_n) = \Delta(u_n, V_n) = \frac{3}{4} = 1 - \rho^{2/3}$
- $\forall x \neq 0, \Delta(u_n^* + xu_n, V_n) \leq \frac{1}{2} = 1 - \rho^{1/3} = J^{(1.5)}(\Delta(V_n))$

Consequently, $\mathbf{E}[\delta_x] \leq J^{(1.5)}(V_n) + o(1) \leq \delta_{\max} - \frac{1}{4} + o(1)$.

We shall need to following claim in our proof of the lemma.

Claim 3.4. *For every $x \in \mathbb{F}_{2^n} \setminus \{0\}$ there exists a polynomial $P_x(Y) \in \mathbb{F}_{2^n}[Y]$ of the form*

$$P_x(Y) = Y^{2^{n-1}} + xY^{2^{n-2}} + \tilde{P}_x, \quad \deg(\tilde{P}_x) < 2^{n-3}.$$

that has 2^{n-1} distinct roots in \mathbb{F}_{2^n} .

Proof. For $x \neq 0$ let $\beta_x = 1/x^2$, noticing β_x is unique because the map $\beta \mapsto \beta^2$ is bijective on \mathbb{F}_{2^n} . Let $\text{Tr}(Z) \triangleq \sum_{i=0}^{n-1} Z^{2^i}$ be the trace function from \mathbb{F}_{2^n} to \mathbb{F}_2 . Define

$$S_x = \{y \in \mathbb{F}_{2^n} \mid \text{Tr}(\beta_x y) = 0\}.$$

It is well known that $|S_x| = 2^{n-1}$ because the trace function has 2^{n-1} roots in \mathbb{F}_{2^n} . So we define

$$P_x(Y) = \frac{1}{\beta_x^{2^{n-1}}} \cdot \text{Tr}(\beta_x Y) = Y^{2^{n-1}} + \frac{1}{\beta_x^{2^{n-2}}} Y^{2^{n-2}} + \tilde{P}_x = Y^{2^{n-1}} + xY^{2^{n-2}} + \tilde{P}_x, \quad \deg(\tilde{P}_x(Y)) < 2^{n-3}.$$

The last equality follows because $\beta_x^{2^{n-2}} = x$. □

Proof of Lemma 3.3. Consider V_n in this family and let $\mathbb{F} = \mathbb{F}_{2^n}$. Define $u^* : \mathbb{F} \rightarrow \mathbb{F}$ to be the function $u^*(y) = y^{2^{n-1}}$ and let $u : \mathbb{F} \rightarrow \mathbb{F}$ be the function $u(y) = y^{2^{n-2}}$.

By Claim 3.4, for every $x \in \mathbb{F} \setminus \{0\}$ there is some $v_x \in V_n$ and P_x with 2^{n-1} roots in \mathbb{F} such that

$$P_x - (u^* + xu) + v_x = 0.$$

Then

$$\Delta(u^* + xu, v_x) = \Pr_{y \in \mathbb{F}}[u^*(y) + xu(y) \neq v_x(y)] = \Pr_y[P_x(y) \neq 0] = 1/2.$$

Thus we get that for all $x \in \mathbb{F} \setminus \{0\}$

$$\Delta(u^* + xu, V) \leq 1/2.$$

On the other hand,

$$\Delta(u, V) \geq 3/4,$$

because for all $v \in V_n$, $u - v$ is a polynomial of degree at most $2^{n-2} = |\mathbb{F}|/4$. This completes the proof. □

Remark 3.5. *Since this example is based on Reed-Solomon codes, it also easily translates into a limitation on the soundness of FRI. In particular, it means that the improvement to the soundness of FRI given in Remark 5.2 is optimal.*

Discussion Lemma 3.3 raises the question of whether the one-and-a-half Johnson bound of Lemma 3.1 is tight for all RS codes, including non-binary fields and evaluation domains that are strict subsets of the ambient field. We point out that the technique used to prove Lemma 3.3 deteriorates rapidly even for binary fields, and even when the evaluation domain is an \mathbb{F}_2 -linear space which resembles the case above.

Indeed, consider an evaluation domain $D \subset \mathbb{F}_{2^n}$ that is a $d + 1$ -dimensional linear space over \mathbb{F}_2 , where $n > d + 1$. There are 2^{d+1} subspaces of dimension d in V . For such $U \subset V$, $\dim(U) = d$, the polynomial $P_U(X) = \prod_{\alpha \in U} (X - \alpha)$ is of the form

$$P_U(Y) = Y^{2^d} + x_U Y^{2^d-1} + \hat{P}_U(Y)$$

which resembles the structure of Claim 3.4. Moreover, as was the case there, for $U' \neq U, U' \subset V, \dim(U') = d$ we have $x_U \neq x_{U'}$. This is because $P_U - P_{U'}$ is a non-zero polynomial with

2^{d-1} roots, because $\dim(U \cap U') = d - 1$. Thus, we cannot have $x_U = x_{U'}$ as this would imply $\deg(P_U - P_{U'}) \leq 2^{d-2} < 2^{d-1}$, contradiction.

As in Lemma 3.3, taking u^* to be the evaluation of Y^{2^d} on D and u be the evaluation of $Y^{2^{d-1}}$ on the same space, we conclude there exists a set $A \subset \mathbb{F}$, $|A| = 2^d$, such that for $x \in A$ we have that $u^* + xu$ agrees with some RS codeword of rate 2^{-3} on half of the evaluation domain.

However, notice that $|A|/2^n = 2^{-(n-d)}$, meaning that the probability of sampling $x \in A$ deteriorates exponentially with the difference $n - d$. Thus the above counterexample fails to rule out an improvement to Lemma 3.1 when the length of the code n is much smaller than the size of the field q .

Conceivably, both Lemma 3.1 and the analysis of FRI can be improved significantly under the assumption $n \ll q$. This is the case of most importance to practical implementations of STARKs.

4 The DEEP Theorem — Using Domain Extension for Eliminating Pretenders (DEEP) and Improving Soundness

We now come to the statement of our improved-soundness distance preservation result. We describe it first for the special case of RS codes. A weighted variant of the theorem is shown in Section 4.2 because it is used later in the DEEP-FRI protocol (Section 5). We end with Section 4.3 in which we present a general version of the following result, that applies to all linear codes.

4.1 DEEP Theorem for RS codes

The vectors u^*, u discussed in the previous section are now viewed as functions $u^*, u : D \rightarrow \mathbb{F}_q$ and we are interested in the distance of a random linear combination $u_x = u^* + x \cdot u$ from the code $V = \text{RS}[\mathbb{F}_q, D, \rho]$, where $x \in \mathbb{F}_q$ is sampled uniformly. Lemma 3.1 established that if $\max(\Delta(u^*, V), \Delta(u, V)) = \delta_{\max}$, then with high probability (over x), the function u_x will have distance at least $\approx \min(\delta_{\max}, 1 - \rho^{1/3})$ from V .

Lemma 3.2 roughly gets used in the following way in the FRI protocol. There are two functions $u^*, u : D \rightarrow \mathbb{F}_q$ and there is a prover who claims that both are evaluations of low degree polynomials. In order to verify this, the verifier uniformly samples $x \in \mathbb{F}_q$ and considers the function $u_x = u^* + x \cdot u$. Lemma 3.2 shows that if any of u^*, u is far from being evaluations of a low degree polynomial, then so is $u^* + x \cdot u$. This then gets exploited in the FRI protocol using FFT type ideas.

We now precede the random process of sampling $x \in \mathbb{F}_q$ with a step of *domain extension*, explained next. Assume a prover claims that both u and u^* are evaluations of low degree polynomials (say $P(Y)$ and $P^*(Y)$). So these polynomials can be evaluated also outside of D . Based on this, a verifier first samples $z \in \mathbb{F}_q$ uniformly and asks the prover to reply with two field elements $a^*, a \in \mathbb{F}_q$ which are supposedly equal to $P^*(z), P(z)$, respectively. After receiving these answers, the verifier proceeds as before, sampling uniformly $x \in \mathbb{F}_q$. Then, setting $b = a^* + x \cdot a$, we examine the distance of u_x from the sub-code $V_{z,b} \subset V$ comprised of all members of V whose interpolating polynomial evaluates to b on input z . The code $V_{z,b}$ is the additive coset (shifted by b) of a low-degree ideal, the ideal generated by $(X - z)$ (cf. Lemma 5.3).

Using the Johnson Bound (Theorem 2.2) we prove that with high probability u_x is at least $\approx \min(\delta_{\max}, 1 - \rho^{1/2})$ far from $V_{z,b}$. Assuming RS codes have a larger list-decoding radius (Conjecture 2.3), we show that with high probability u_x is $\approx \delta_{\max}$ -far from $V_{z,b}$ for nearly all values of δ_{\max} .

Later, in Section 5, we shall use the improved distance preservation to construct the DEEP-FRI protocol for testing proximity to the RS code with improved soundness.

The statement we give below is given more generally in terms of the list size bound $\mathcal{L}(\mathbb{F}_q, D, d = \rho|D|, \delta)$; we instantiate it later with the Johnson bound and with Conjecture 2.3. It is useful to keep in mind that this will be used in a setting where q is much larger than $|D|$ (and hence d), and where L_δ^* is small.

Theorem 4.1 (DEEP method for RS codes). *Let $\rho > 0$ and let $V = \text{RS}[\mathbb{F}_q, D, \rho]$. For $z, b \in \mathbb{F}_q$, we let*

$$V_{z,b} = \{Q(Y)|_D \in V \mid Q(z) = b\}.$$

For $\delta > 0$ let $L_\delta^* = \mathcal{L}(\mathbb{F}_q, D, d = \rho|D|, \delta)$.

Let $u, u^* \in \mathbb{F}_q^D$. For each $z \in \mathbb{F}_q$, let $B_z(X) \in \mathbb{F}_q[X]$ be an arbitrary linear function. Suppose that for some $1/3 > \epsilon > 0$ the following holds,

$$\Pr_{x,z \in \mathbb{F}_q} [\Delta(u^* + xu, V_{z,B_z(x)}) < \delta] \geq \max \left(2L_\delta^* \left(\frac{d}{q} + \epsilon \right)^{1/3}, \frac{4}{\epsilon^2 q} \right), \quad (7)$$

Then there exist $v, v^* \in V$ and $C \subset D$ such that:

- $|C| \geq (1 - \delta - \epsilon)|D|$,
- $u|_C = v|_C$,
- $u^*|_C = v^*|_C$.

Consequently, we have $\Delta(u, V), \Delta(u^*, V) \leq \delta + \epsilon$.

Proof. To simplify notation set $\eta = \max \left(2L_\delta^* \left(\frac{d}{q} + \epsilon \right)^{1/3}, \frac{4}{\epsilon^2 q} \right)$, and let $u_x = u^* + xu$.

Let $\mathcal{E}[x, z]$ denote the event “ $\exists P(Y) \in \text{List}(u_x, V, \delta), P(z) = B_z(x)$ ”.

The assumption of Equation (7) now reads as

$$\Pr_{x,z \in \mathbb{F}_q} [\mathcal{E}[x, z]] \geq \eta.$$

Thus we get,

$$\Pr_x [\Pr_z [\mathcal{E}[x, z]] \geq \eta/2] \geq \eta/2 \quad (8)$$

Let

$$A = \left\{ x \in \mathbb{F}_q \mid \Pr_z [\mathcal{E}[x, z]] \geq \eta/2 \right\}$$

and notice $|A| \geq \eta q/2$.

For $x \in \mathbb{F}_q$, pick $P_x \in V$ to be a member $P \in \text{List}(u_x, V, \delta)$ that maximizes $\Pr_{z \in \mathbb{F}_q} [P(z) = B_z(x)]$. Let $S_x = \{z \in \mathbb{F}_q \mid P_x(z) = B_z(x)\}$ and set $\mu_x = |S_x|/q$. By definition, $|\text{List}(u_x, V, \delta)| \leq L_\delta^*$, and so by the pigeonhole principle, for each $x \in A$ we have $\mu_x \geq \frac{\eta}{2L_\delta^*}$.

For x, β, γ picked uniformly from A , and z picked uniformly from \mathbb{F}_q , we have:

$$\mathbf{E}_{x,\beta,\gamma} [|S_x \cap S_\beta \cap S_\gamma|/q] = \mathbf{E}_{z,x,\beta,\gamma} [\mathbf{1}_{z \in S_x \cap S_\beta \cap S_\gamma}]$$

$$\begin{aligned}
&= \mathbf{E}_z[\mathbf{E}_x[1_{z \in S_x}]^3] \\
&\geq \mathbf{E}_{z,x}[1_{z \in S_x}]^3 \\
&\geq \left(\frac{\eta}{2L_\delta^*}\right)^3 \\
&> \frac{d}{q} + \epsilon.
\end{aligned}$$

The second equality above follows from the independence of x, β, γ . The first inequality is an application of Jensen's inequality and the last inequality is by assumption on η .

Thus

$$\Pr_{x,\beta,\gamma}[|S_x \cap S_\beta \cap S_\gamma| > d] \geq \epsilon.$$

Note that $\Pr_{x,\beta,\gamma}[x, \beta, \gamma \text{ are not all distinct}] < 3/|A|$. Since $|A| \geq \eta q/2 \geq 2/\epsilon^2 \geq 6/\epsilon$ we have $3/|A| \leq \epsilon/2$. Thus $\Pr_{x,\beta,\gamma}[x, \beta, \gamma \text{ are all distinct and } |S_x \cap S_\beta \cap S_\gamma| > d] \geq \epsilon/2$.

This means that there are distinct x_0, β_0 such that

$$\Pr_\gamma[|S_{x_0} \cap S_{\beta_0} \cap S_\gamma| > d] \geq \epsilon/2.$$

Consider some γ where this happens. Let $S = S_{x_0} \cap S_{\beta_0} \cap S_\gamma$. By construction we know that for all $z \in \mathbb{F}_q$,

$$(x_0, B_z(x_0)), (\beta_0, B_z(\beta_0)), (\gamma, B_z(\gamma))$$

are collinear. So, in particular, for $z \in S$ this holds.

By definition of S , we get that for each $z \in S$,

$$(x_0, P_{x_0}(z)), (\beta_0, P_{\beta_0}(z)), (\gamma, P_\gamma(z)) \in \mathbb{F}_q \times \mathbb{F}_q$$

are collinear. Since $|S| > d$, we have that P_γ is uniquely determined by $P_\gamma|_S$ by a linear map. This allows us to conclude that

$$(x_0, P_{x_0}), (\beta_0, P_{\beta_0}), (\gamma, P_\gamma) \in \mathbb{F}_q \times \mathbb{F}_q[Y]$$

are collinear in the \mathbb{F}_q -vector space $\mathbb{F}_q \times \mathbb{F}_q[Y]$.

Thus, an $\epsilon/2$ -fraction of the $\gamma \in A$ have the “good” property that (γ, P_γ) is on the line passing through (x_0, P_{x_0}) and (β_0, P_{β_0}) . Write this line as $P^* + xP$ and notice that for all “good” γ we have $P_\gamma = P^* + \gamma P$. Let $A' \subseteq A$ denote the set of good elements for this line, recording that $|A'| \geq |A| \cdot \epsilon/2 \geq 1/\epsilon$. By definition of $\text{List}(u_x, V, \delta)$ and the assumption $P_x \in \text{List}(u_x, V, \delta)$, we have that $\Delta(u_x, P_x) < \delta$ for $x \in A'$.

Consider the set $C \subset D$ defined by

$$C = \{y \in D \mid u^*(y) = P^*(y) \text{ AND } u(y) = P(y)\}.$$

For each $y \in D \setminus C$ there exists at most a single value of $x \in \mathbb{F}_q$ satisfying $u_x(y) = P_x(y)$ because

$$u_x(y) - P_x(y) = (u^*(y) - P^*(y)) + x \cdot (u(y) - P(y))$$

has at most one value x on which it vanishes. This implies

$$\delta \geq \mathbf{E}_{x \in A'}[\Delta_D(u_x, v_x)] \geq \frac{|D \setminus C|}{|D|} \cdot \left(1 - \frac{1}{|A'|}\right) \geq \left(1 - \frac{|C|}{|D|}\right) \cdot (1 - \epsilon) \geq 1 - \frac{|C|}{|D|} - \epsilon.$$

Rearranging, we get $\frac{|C|}{|D|} \geq 1 - (\delta + \epsilon)$. Taking $v = P$ and $v^* = P^*$ completes the proof. \square

Remark 4.2. One could extend the domain even further, and sample z from an extension field \mathbb{F}_{q^a} . This gives even better soundness; the expression $2L_\delta^* \cdot \left(\frac{d}{q} + \epsilon\right)^{1/3}$ by $2L_\delta^* \cdot \left(\frac{d}{q^a} + \epsilon\right)^{1/3}$. This can give interesting results even if $L_\delta^* = q^{O(1)}$ by taking $a = O(1)$.

4.2 Weighted version

For application to Reed-Solomon Proximity Testing, it is more convenient to have a weighted version of the previous result. We briefly introduce some notation for dealing with weights, and then state the new version.

Let $u, v \in \mathbb{F}_q^D$. Let $\eta \in [0, 1]^D$ be a vector of weights. We define the η -agreement between u and v by:

$$\text{agree}_\eta(u, v) = \frac{1}{|D|} \sum_{i \in D | u_i = v_i} \eta(i).$$

For a subspace $V \subseteq \mathbb{F}_q^n$, we define

$$\text{agree}_\eta(u, V) = \max_{v \in V} \text{agree}_\eta(u, v).$$

Theorem 4.3. Let $\rho > 0$ and let $V = \text{RS}[\mathbb{F}_q, D, d = \rho \cdot |D|]$. For $z, b \in \mathbb{F}_q$, we let

$$V_{z,b} = \{Q(Y) \in V \mid Q(z) = b\}.$$

For $\alpha < 1$, let $L^* = \mathcal{L}(\mathbb{F}_q, D, d = \rho|D|, 1 - \alpha)$ be the list-size for list-decoding V from $(1 - \alpha)$ -fraction errors (without weights).

Let $u, u^* \in \mathbb{F}_q^D$. For each $z \in \mathbb{F}_q$, let $B_z(X) \in \mathbb{F}_q[X]$ be an arbitrary linear function. Suppose that

$$\Pr_{x, z \in \mathbb{F}_q} [\text{agree}_\eta(u^* + xu, V_{z, B_z(x)}) > \alpha] \geq \max \left(2L^* \left(\frac{d}{q} + \epsilon \right)^{1/3}, \frac{4}{\epsilon^2 q} \right), \quad (9)$$

Then there exist $v, v^* \in V$ and $C \subset D$ such that:

- $\sum_{y \in C} \eta(y) > (\alpha - \epsilon)|D|$,
- $u|_C = v|_C$,
- $u^*|_C = v^*|_C$.

Consequently, we have $\text{agree}_\eta(u, V), \text{agree}_\eta(u^*, V) \geq \alpha - \epsilon$.

The proof is nearly identical to the proof of Theorem 4.1 so we only highlight the changes. First, we observe that if $\eta_1 : D \rightarrow [0, 1]$ is the the constant function with value 1, then $\text{agree}_{\eta_1}(u, v) \leq \text{agree}_{\eta_1}(u, v) = 1 - \Delta(u, v)$. Thus the set

$$\{Q(Y) \in \mathbb{F}_q[Y] \mid \deg(Q) \leq d, \text{agree}_\eta(u^* + xu, Q) > \alpha\}$$

is contained in

$$\{Q(Y) \in \mathbb{F}_q[Y] \mid \deg(Q) \leq d, \Delta(u^* + xu, Q) < 1 - \alpha\}.$$

The size of this latter set is bounded by L^* , and thus the size of the former set is too. The proof then proceeds as before, until the very end, where we have a set $A' \subseteq \mathbb{F}_q$, with $|A'| \geq \frac{2}{\epsilon}$, and

polynomials $P, P^* \in V$ such that for each $x \in A'$, $\text{agree}_\eta(u^* + xu, P^* + xP) > \alpha$. Then we take $C = \{y \in C \mid u^*(y) = P^*(y), u(y) = P(y)\}$, and our goal is to show that $\sum_{y \in C} \eta(y) > (\alpha - \epsilon)|D|$. To this end, consider:

$$\begin{aligned}
\alpha &< \frac{1}{|A'|} \sum_{x \in A'} \text{agree}_\eta(u^* + xu, P^* + xP) \\
&= \frac{1}{|D| \cdot |A'|} \sum_{x \in A'} \sum_{y \in D} (\eta(y) \cdot 1_{u^*(y) + xu(y) = P^*(y) + xP(y)}) \\
&= \frac{1}{|D|} \sum_{y \in D} \eta(y) \left(\frac{1}{|A'|} \sum_{x \in A'} 1_{u^*(y) + xu(y) = P^*(y) + xP(y)} \right) \\
&\leq \frac{1}{|D|} \sum_{y \in C} \eta(y) + \frac{1}{|D|} \sum_{y \in D \setminus C} \eta(y) \cdot \frac{1}{|A'|} \\
&\leq \frac{1}{|D|} \sum_{y \in C} \eta(y) + \epsilon/2.
\end{aligned}$$

This implies that $\sum_{y \in C} \eta(y) > (\alpha - \epsilon)|D|$, and the rest of the proof is the same as before.

4.3 DEEP Lemma for general linear codes

Theorem 4.1 can be generalized to apply to arbitrary linear codes, and this is the focus of this section. We explain the basic principles for an $[n, k, d]_q$ -linear code V with generating matrix $G \in \mathbb{F}_q^{k \times n}$, viewing codewords as evaluations of linear forms on the columns of G .

Let $D \subset \mathbb{F}_q^k$ be the set of columns of G . A linear form $\ell \in \mathbb{F}_q^k$ can be “evaluated” at any any element x of D . Similarly, if we fix a set of points $S \subseteq \mathbb{F}_q^k$ (thinking $|S| \gg |D|$), we may evaluate the linear form ℓ at any point of S – this corresponds to evaluation outside the original domain D .

If we are given a function $u : D \rightarrow \mathbb{F}_q$ which is supposed to be the evaluations of a linear form ℓ on D , we can ask about what the evaluation of this linear form at a point $z \in S$ is. This is the viewpoint from which the DEEP lemma generalizes to general codes.

We start with two functions $u, u^* : D \rightarrow \mathbb{F}_q$ (which are supposed to correspond to linear forms, say $\ell \in \mathbb{F}_q^k$ and $\ell^* \in \mathbb{F}_q^k$). We have a verifier who samples $z \in S$ and asking for $a = \ell(z)$ and $a^* = \ell^*(z)$. Given these answers, the verifier now samples $x \in \mathbb{F}_q$ and computes $b = a^* + xa$ which is supposedly equal to $\ell^*(z) + \ell(z)$ (if u^* and u are indeed codewords of V). The result below says that if S is the set of columns of an error correcting code with good distance, and V has small list size for list-decoding up to radius δ , then with high probability, the function $u_x = u^* + xu$ has distance at least $\approx \min\{\Delta(u^*, V), \delta\}$ from the sub-code of V corresponding to the linear forms that evaluate to b on z .

Definition 4.4 (Robust). *A set $S \subseteq \mathbb{F}_q^k$ is called σ -robust if every subset of S of size σ contains a basis for \mathbb{F}_q^k .*

The following claim is well-known in coding theory (cf. [Rot06, Problem 2.8]).

Claim 4.5. *Fix a full-rank matrix $G \in \mathbb{F}_q^{k \times N}$, $N \geq k$, and let $C = \{x \cdot M \mid x \in \mathbb{F}_q^k\}$ be the linear code generated by it. Then the set of columns of G is σ -robust if and only if the minimum distance of C is at least $N - \sigma + 1$.*

Lemma 4.6 (DEEP method for general linear codes). *Let V be an $[n, k, d]_q$ -code that is (δ, L_δ^*) -list decodable for some $\delta > 0$, and fix $G \in \mathbb{F}_q^{k \times n}$ to be its generating matrix. Let $S \subset \mathbb{F}_q^k$ be a σ -robust set of size N . For $z \in S, b \in \mathbb{F}_q$, let*

$$V_{z,b} = \{v \in V \mid v = G \cdot \ell_v \text{ AND } \langle \ell_v, z \rangle = b\}$$

where $\langle v, z \rangle = \sum_{i=1}^k v_i z_i$.

Let $u, u^* \in \mathbb{F}_q^n$. For each $z \in S$, let $B_z(X) \in \mathbb{F}_q[X]$ be an arbitrary linear function. Suppose that for some $\epsilon > 0$ the following holds,

$$\Pr_{x \in \mathbb{F}_q, z \in S} [\Delta(u^* + xu, V_{z, B_z(x)}) < \delta] \geq \max \left(2L_\delta^* \left(\frac{\sigma}{N} + \epsilon \right)^{1/3}, \frac{4}{\epsilon^2 q} \right), \quad (10)$$

Then there exist $v, v^* \in V$ and $C \subset [n]$ such that:

- $|C| \geq (1 - \delta - \epsilon)n$,
- $u|_C = v|_C$,
- $u^*|_C = v^*|_C$.

Consequently, we have $\Delta(u, V), \Delta(u^*, V) \leq \delta + \epsilon$.

The proof is analogous to the proof in the Reed-Solomon case, and appears in Appendix A.

Discussion For the special case of RS codes, the DEEP method can be used to locally modify the problem and reduce degree. Indeed, the subcode $V_{z,b}$ in the case of RS codes corresponds to evaluations of polynomials of degree d whose interpolating polynomial P_f satisfies $P_f(z) = b$. From such a codeword, one can construct a new codeword $f_{z,b} : D \rightarrow \mathbb{F}$ defined by $f_{z,b}(x) = \frac{f(x) - b}{z - x}$, which is well-defined for all $x \neq z$. Notice that the transformation from f to $f_{z,b}$ is *1-local*, meaning that each entry of $f_{z,b}$ is constructed by making a single query to f . Furthermore, this transformation maps a subset of the code $RS[\mathbb{F}, D, d]$ to the code $RS[\mathbb{F}, D, d - 1]$, so we may use this transformation in RS IOPPs (as will be done in the following section).

In contrast, for a general k -dimensional linear code V , the subcode $V_{z,b}$, while being an affine subspace of V , has less structure. In particular, it is not clear how to locally convert this subcode to a “nice” code of dimension $k - 1$. An interesting middle ground, left to future work, is the case of algebraic codes like Reed Muller codes and Algebraic Geometry codes which resemble RS codes.

5 DEEP-FRI

In this section we describe the new fast RS IOPP, called DEEP-FRI. We start by recalling the FRI protocol from [BBHR18b], describing it nearly verbatim as in [BKS18b, Section 7].

5.1 FRI

Our starting point is a function $f^{(0)} : L^{(0)} \rightarrow \mathbb{F}$ where \mathbb{F} is a finite field, the evaluation domain $L^{(0)} \subset \mathbb{F}$ is a coset of a group² contained in \mathbb{F} , and $|L^{(0)}| = 2^{k^{(0)}}$. We assume the target rate is $\rho = 2^{-\mathcal{R}}$ for some positive integer \mathcal{R} . The FRI protocol is a two-phase protocol (the two phases are called COMMIT and QUERY) that convinces a verifier that $f^{(0)}$ is close to the Reed-Solomon code $\text{RS}[\mathbb{F}, L^{(0)}, \rho]$.

The COMMIT phase of the FRI protocol involves $r = k^{(0)} - \mathcal{R}$ rounds. Before any communication, the prover and verifier agree on a sequence of (cosets of) sub-groups $L^{(i)}$, where $|L^{(i)}| = 2^{k^{(0)} - i}$. Let $\text{RS}^{(i)}$ denote the Reed-Solomon code $\text{RS}[\mathbb{F}, L^{(i)}, \rho|L^{(i)}|]$.

The main ingredient of the FRI protocol is a special algebraic hash function H_x , which takes a seed $x \in \mathbb{F}$, and given as input a function $f : L^{(i)} \rightarrow \mathbb{F}$, it produces as output a hash whose length is 1/2 as long as f . More concretely, $H_x[f]$ is a function

$$H_x[f] : L^{(i+1)} \rightarrow \mathbb{F}$$

with the following properties:

1. **locality:** For any $s \in L^{(i+1)}$, $H_x[f](s)$ can be computed by querying f at just two points in its domain (these two points are $(q^{(i)})^{-1}(s)$).
2. **completeness:** If $f \in \text{RS}^{(i)}$, then for all $x \in \mathbb{F}$, we have that $H_x[f] \in \text{RS}^{(i+1)}$.
3. **soundness:** If f is far from $\text{RS}^{(i)}$, then with high probability over the choice of seed x , $H_x[f]$ is quite far from $\text{RS}^{(i+1)}$.

These last two properties roughly show that for random x , H_x preserves distance to Reed-Solomon codes. For the precise description of H_x see Appendix B and [BKS18a].

The high-level idea of the FRI protocol can then be described as follows. First we are in the COMMIT phase of the protocol. The verifier picks a random $x^{(0)} \in \mathbb{F}$ and asks the prover to write down the hash $H_{x^{(0)}}[f^{(0)}] : L^{(1)} \rightarrow \mathbb{F}$. By Properties 2 and 3 above, our original problem of estimating the distance of $f^{(0)}$ to $\text{RS}^{(0)}$ reduces to estimating the distance of $H_{x^{(0)}}[f^{(0)}]$ to $\text{RS}^{(1)}$ (which is a problem of 1/2 the size). This process is then repeated: the verifier picks a random $x^{(1)} \in \mathbb{F}$ and asks the prover to write down $H_{x^{(1)}}[H_{x^{(0)}}[f^{(0)}]]$, and so on. After r rounds of this, we are reduced to a constant sized problem which can be solved in a trivial manner. However, the verifier cannot blindly trust that the functions $f^{(1)}, \dots$ that were written down by the prover truly are obtained by repeatedly hashing $f^{(0)}$. This has to be checked, and the verifier does this in the QUERY phase of the protocol, using Property 1 above.

We describe the phases of the protocol below.

COMMIT Phase:

1. For $i = 0$ to $r - 1$:
 - (a) The verifier picks uniformly random $x^{(i)} \in \mathbb{F}$ and sends it to the prover.
 - (b) The prover writes down a function $f^{(i+1)} : L^{(i+1)} \rightarrow \mathbb{F}$. (In the case of an honest prover, $f^{(i+1)} = H_{x^{(i)}}[f^{(i)}]$.)

²The group can be additive, in which case \mathbb{F} is a binary field, or multiplicative, in which case it is not.

2. The prover writes down a value $C \in \mathbb{F}_q$. (In the case of an honest prover, $f^{(r)}$ is the constant function with value $= C$).

QUERY Phase: (executed by the Verifier)

1. Repeat ℓ times:
 - (a) Pick $s^{(0)} \in L^{(0)}$ uniformly at random.
 - (b) For $i = 0$ to $r - 1$:
 - i. Define $s^{(i+1)} \in L^{(i+1)}$ by $s^{(i+1)} = q^{(i)}(s^{(i)})$.
 - ii. Compute $H_{x^{(i)}}[f^{(i)}](s^{(i+1)})$ by making 2 queries to $f^{(i)}$.
 - iii. If $f^{(i+1)}(s^{(i+1)}) \neq H_{x^{(i)}}[f^{(i)}](s^{(i+1)})$, then REJECT.
 - (c) If $f^{(r)}(s^{(r)}) \neq C$, then REJECT.
2. ACCEPT

The previous state of the art regarding the soundness of FRI is given by the following statement from [BKS18a]. In what follows let $J_\epsilon(x) = 1 - \sqrt{1 - x(1 - \epsilon)}$.

Theorem 5.1 (FRI soundness (informal)). *Suppose $\delta^{(0)} \triangleq \Delta(f^{(0)}, \text{RS}^{(0)}) > 0$. Let $n = |L^{(0)}|$. Then for any $\epsilon > 0$ there exists $\epsilon' > 0$ so that with probability at least*

$$1 - \frac{2 \log n}{\epsilon^3 |\mathbb{F}|} \tag{11}$$

over the randomness of the verifier during the COMMIT phase, and for any (adaptively chosen) prover oracles $f^{(1)}, \dots, f^{(r)}$, the QUERY protocol with repetition parameter ℓ outputs accept with probability at most

$$\left(1 - \min \left\{ \delta^{(0)}, 1 - (\rho^{1/4} + \epsilon') \right\} + \epsilon \log n \right)^\ell. \tag{12}$$

Remark 5.2. *Using the improved distance preservation of Lemma 3.2 in the analysis of FRI from [BKS18a], one immediately improves the factor 1/4 in the exponent in Equation (14) to an exponent of 1/3 (details omitted).*

5.2 DEEP-FRI

We now describe our variation of FRI, that we call DEEP-FRI, for which we can give improved soundness guarantees, at the cost of a small increase in the query complexity (but no increase in the proof length or the number of queries to committed proofs – which is important in applications).

Before we can describe our protocol we introduce the operation of “quotienting”, which allows us to focus our attention on polynomials taking certain values at certain points.

5.2.1 Quotienting

Suppose we a set $L \subseteq \mathbb{F}_q$ and a function $f : L \rightarrow \mathbb{F}_q$. Suppose further that we are given a point $z \in \mathbb{F}_q$ and a value $b \in \mathbb{F}_q$.

We define the function $\text{QUOTIENT}(f, z, b) : L \rightarrow \mathbb{F}_q$ as follows. Let $Z(Y) \in \mathbb{F}_q[Y]$ be the polynomial $Z(Y) = Y - z$. Then we define $\text{QUOTIENT}(f, z, b)$ to be the function $g : L \rightarrow \mathbb{F}_q$ given by:

$$g(y) = \frac{f(y) - b}{Z(y)}$$

(or more succinctly, $g = \frac{f-b}{Z}$).

Lemma 5.3. *Let $L \subseteq \mathbb{F}_q$. Let $z \in \mathbb{F}_q$ with $z \notin L$. Let $d \geq 1$ be an integer.*

Let $f : L \rightarrow \mathbb{F}_q$, and $b \in \mathbb{F}_q$. Let $g = \text{QUOTIENT}(f, z, b)$. Then the following are equivalent:

- *There exists a polynomial $Q(X) \in \mathbb{F}_q[X]$ of degree at most $d - 1$ such that $\Delta(g, Q) < \delta$.*
- *There exists a polynomial $R(X) \in \mathbb{F}_q[X]$ of degree at most d such that $\Delta(f, R) < \delta$ and $R(z) = b$.*

Proof. If there is such a polynomial Q , $\deg(Q) \leq d - 1$ that agrees with g on all but a δ -fraction of entries, we can take $R = QZ + b$. Notice $\deg(R) \leq d$ because $\deg(Z) = 1$.

Conversely, if there is such a polynomial R that agrees with f on all but a δ -fraction of entries, we can take $Q = (R - b)/Z$. This is indeed a polynomial because $R - b$ vanishes on z , so $Z|(R - b)$ in the ring of polynomials.

Finally, by construction R agrees with f whenever g agrees with R and this completes the proof. \square

5.3 DEEP-FRI

Recall: We have linear spaces $L^{(0)}, L^{(1)}, \dots, L^{(r)}$, with dimensions $k, k - 1, \dots, k - r$. We further have 1 dimensional subspaces $L_0^{(0)}, L_0^{(1)}, \dots, L_0^{(r)}$ with $L_0^{(i)} \subseteq L^{(i)}$.

For this, it will be helpful to keep in mind the case that the domain $L^{(0)}$ is much smaller than the field \mathbb{F}_q (maybe $q = |L^{(0)}|^{\Theta(1)}$).

Protocol 5.4 (DEEP-FRI).

Input: a function $f^{(0)} : L^{(0)} \rightarrow \mathbb{F}_q$ which is supposed to be of degree $< d^{(0)}$.

COMMIT Phase:

1. For each $i \in [0, r - 1]$:
 - (a) The verifier picks a uniformly random $z^{(i)} \in \mathbb{F}_q$.
 - (b) The prover writes down a degree one polynomial $B_{z^{(i)}}^{(i)}(X) \in \mathbb{F}_q[X]$ (which is supposed to be such that $B_{z^{(i)}}^{(i)}(x)$ equals the evaluation of the low degree polynomial $H_x[f^{(i)}]$ at $z^{(i)}$).
 - (c) The verifier picks uniformly random $x^{(i)} \in \mathbb{F}_q$.

(d) The prover writes down a function

$$f^{(i+1)} : L^{(i+1)} \rightarrow \mathbb{F}_q.$$

(which on input y is supposed to equal $\text{QUOTIENT}(H_{x^{(i)}}[f^{(i)}], z^{(i)}, B_{z^{(i)}}^{(i)}(x)).$)

2. The prover writes down a value $C \in \mathbb{F}_q$.

QUERY Phase:

1. Repeat ℓ times:

(a) The verifier picks a uniformly random $s^{(0)} \in D$.

(b) For each $i \in [0, r-1]$:

i. Define $s^{(i+1)} \in L^{(i+1)}$ by $s^{(i+1)} = q^{(i)}(s^{(i)})$.

ii. Compute $H_{x^{(i)}}[f^{(i)}](s^{(i+1)})$ by making 2 queries to $f^{(i)}$.

iii. If $H_{x^{(i)}}[f^{(i)}](s^{(i+1)}) \neq f^{(i+1)}(s^{(i+1)}) \cdot (s^{(i+1)} - z^{(i)}) + B_{z^{(i)}}^{(i)}(x^{(i)})$, then REJECT.

(c) If $f^{(r)}(s^{(r)}) \neq C$, then REJECT.

2. ACCEPT.

5.4 Analysis

The following theorem proves the soundness of the DEEP-FRI protocol.

Theorem 5.5 (DEEP-FRI). *Fix degree bound $d^{(0)} = 3 \cdot 2^r - 2$ and $\text{RS}^{(0)} = \text{RS}[\mathbb{F}_q, L^{(0)}, d^{(0)}]$. Let $n = |L^{(0)}|$.*

For some $\epsilon, \delta > 0$, let

$$\begin{aligned} \delta^* &= \delta - 2r\epsilon, \\ \mathcal{L}^* &= \mathcal{L}(\mathbb{F}_q, L^{(0)}, d^{(0)}, \delta^*), \\ \nu^* &= 2\mathcal{L}^* \left(\frac{d^{(0)}}{q} + \epsilon \right)^{1/3} + \frac{4}{\epsilon^2 q}. \end{aligned}$$

Then the following properties hold when the DEEP-FRI protocol is invoked on oracle $f^{(0)} : L^{(0)} \rightarrow \mathbb{F}_q$,

1. **Prover complexity** is $O(n)$ arithmetic operations over \mathbb{F}
2. **Verifier complexity** is $O(\log n)$ arithmetic operations over \mathbb{F} for a single invocation of the QUERY phase; this also bounds communication and query complexity (measured in field elements).
3. **Completeness** If $f^{(0)} \in \text{RS}^{(0)}$ and $f^{(1)}, \dots, f^{(r)}$ are computed by the prover specified in the COMMIT phase, then the DEEP-FRI verifier outputs accept with probability 1.

4. **Soundness** Suppose $\Delta(f^{(0)}, \text{RS}^{(0)}) > \delta$. Then with all but probability

$$\text{err}_{\text{COMMIT}} \leq r \cdot \nu^* \leq (\log n) \cdot \nu^*. \quad (13)$$

and for any (adaptively chosen) prover oracles $f^{(1)}, \dots, f^{(r)}$, the QUERY protocol with repetition parameter ℓ outputs accept with probability at most

$$\text{err}_{\text{QUERY}} \leq (1 - \delta^* + (\log n) \cdot \epsilon)^\ell \quad (14)$$

Consequently, the soundness error of FRI is at most

$$\text{err}(\delta) \leq (\log n) \cdot \nu^* + (1 - \delta^* + (\log n) \cdot \epsilon)^\ell \quad (15)$$

We give a consequence below with a specific setting of parameters based on the Johnson bound.

Example 5.6. Continuing with the notation of Theorem 5.5, fix degree bound $d^{(0)} = 3 \cdot 2^r - 2$ and assume $n = |L^{(0)}| < \sqrt{q}$. Let $\text{RS}^{(0)} = \text{RS}[\mathbb{F}_q, L^{(0)}, d^{(0)}]$ and let $\rho = d^{(0)}/n$ be its rate.

Let $f^{(0)} : L^{(0)} \rightarrow \mathbb{F}_q$ be a function, and let $\delta^{(0)} = \Delta(f^{(0)}, \text{RS}^{(0)})$. Then with all but probability $\text{err}_{\text{COMMIT}} \leq O(q^{-\Omega(1)})$, the query phase will accept with probability at most: $\text{err}_{\text{QUERY}} \leq (\max(1 - \delta^{(0)}, \sqrt{\rho}) + o(1))^\ell$ as $n \rightarrow \infty$.

Proof. Note that $d^{(0)} \leq n \leq \sqrt{q}$.

Set $\delta = \min(\delta^{(0)}, 1 - \sqrt{\rho} - q^{-1/13})$, and apply the previous theorem. Theorem 2.2 implies that $\mathcal{L}^* < q^{1/13}/(2\sqrt{\rho}) = O(q^{1/13})$. Set $\epsilon = q^{-6/13}$. Hence

$$\nu^* = 2\mathcal{L}^* \left(d^{(0)}q^{-1} + q^{-6/13} \right)^{1/3} + 4q^{-6/13} = O(q^{-1/13}),$$

which implies $\text{err}_{\text{COMMIT}} \leq \tilde{O}(q^{-1/13})$.

If $\delta = \delta^{(0)}$, then $1 - \delta^* + (\log n)\epsilon = 1 - \delta + o(1)$. Otherwise $\delta = 1 - \sqrt{\rho} - q^{-1/13}$, and so

$$1 - \delta^* + (\log n)\epsilon = \sqrt{\rho} + q^{-1/13} + (\log n)\epsilon = \sqrt{\rho} + q^{-1/13} + (\log n)q^{-6/13}.$$

Thus $\text{err}_{\text{QUERY}} \leq (\max(1 - \delta, \sqrt{\rho}) + o(1))^\ell$. \square

We now give an example setting of DEEP-FRI under the optimistic Conjecture 2.3.

Example 5.7. Assume Conjecture 2.3. Continuing with the notation of Theorem 5.5, fix degree bound $d^{(0)} = 3 \cdot 2^r - 2$ and $n = |L^{(0)}|$. Let $\text{RS}^{(0)} = \text{RS}[\mathbb{F}_q, L^{(0)}, d^{(0)}]$ and let $\rho = d^{(0)}/n$ be its rate.

Let $C = C_\rho$ be the constant given by Conjecture 2.3. Suppose $q > n^{24C}$.

Let $f^{(0)} : L^{(0)} \rightarrow \mathbb{F}_q$ be a function, and let $\delta^{(0)} = \Delta(f^{(0)}, \text{RS}^{(0)})$. Then with all but probability $\text{err}_{\text{COMMIT}} \leq O(q^{-\Omega(1)})$, the query phase will accept with probability at most: $\text{err}_{\text{QUERY}} \leq (1 - \delta^{(0)} + o(1))^\ell$ as $n \rightarrow \infty$.

Proof. Set $\epsilon = q^{-1/(6C)}$.

Set $\delta = \min(\delta^{(0)}, 1 - \rho - q^{-1/(6C)})$. Conjecture 2.3 gives us:

$$\mathcal{L}^* < n^C q^{1/6}.$$

We then apply the previous theorem. We get $\nu^* \ll O(\mathcal{L}^* \cdot (d/q + \epsilon)^{1/3} + \frac{1}{\epsilon^2 q}) \ll q^{-1/12}$, and this gives us the claimed bound on $\text{err}_{\text{COMMIT}}$.

For the bound on $\text{err}_{\text{QUERY}}$, we note that $\delta = \delta^{(0)} + o(1)$. This is because *every* function is within distance $1 - \rho$ of $\text{RS}^{(0)}$ (this follows easily from polynomial interpolation). Thus

$$1 - \delta^* + (\log n)\epsilon = \delta + o(1),$$

and we get the desired bound on $\text{err}_{\text{COMMIT}}$. \square

The prover and verifier complexity as well as completeness follow by construction (see, e.g., [BBHR18b] for detailed analysis of these aspects). In the rest of the section we prove the soundness bound of Theorem 5.5.

5.5 Preparations

We do the analysis below for the case $\ell = 1$. The generalization to arbitrary ℓ easily follows.

Define $d^{(0)} = 3 \cdot 2^r - 2$, and $d^{(i+1)} = d^{(i)}/2 - 1$. It is easy to check that $d^{(r)} = 1$. Define $\text{RS}^{(i)} = \text{RS}[\mathbb{F}_q, L^{(i)}, d^{(i)}]$. In the case of the honest prover (when $f^{(0)} \in \text{RS}^{(0)}$), we will have $f^{(i)} \in \text{RS}^{(i)}$ for all i .

Our analysis of the above protocol will track the agreement of $f^{(i)}$ with $\text{RS}^{(i)}$. This agreement will be measured in a certain weighted way, which we define next.

5.5.1 The success probability at $s \in L^{(i)}$

There is a natural directed forest that one can draw on the vertex set

$$L^{(0)} \cup L^{(1)} \cup \dots \cup L^{(r)},$$

namely, where $s \in L^{(i)}$ is joined to $q^{(i)}(s) \in L^{(i+1)}$ (and we say that s is a child of $q^{(i)}(s)$). Note that every vertex not in $L^{(0)}$ has two children.

Let $i \leq r - 1$ and $s_0 \in L^{(i)}$. Let $s \in L^{(i+1)}$ be the parent of s_0 , and let $s_1 \in L^{(i)}$ the sibling of s_0 . We color s_0 GREEN if $f^{(i+1)}(s)$ is consistent with $f^{(i)}|_{\{s_0, s_1\}}$ according to the test

$$H_{x^{(i)}}[f^{(i)}](s) = f^{(i+1)}(s) \cdot (s - z^{(i)}) + B_{z^{(i)}}^{(i)}(x^{(i)})$$

and we color s_0 RED otherwise. Notice that a vertex and its sibling get the same color.

For $s \in L^{(r)}$, we color s GREEN if $f^{(r)}(s) = C$ and RED otherwise.

The QUERY phase of the protocol can be summarized as follows: we pick a uniformly random $s^{(0)} \in L^{(0)}$ and consider the path $s^{(0)}, s^{(1)}, s^{(2)}, \dots, s^{(r)}$ going through all its ancestors. If all these vertices are GREEN, then we ACCEPT, otherwise we REJECT.

To capture this, we define functions $\eta^{(i)} : L^{(i)} \rightarrow \mathbb{R}$ as follows. For $s \in L^{(i)}$, let $\eta^{(i)}(s)$ be the fraction of leaf-descendants $s^{(0)}$ of s for which the path from $s^{(0)}$ to s (including $s^{(0)}$ but not including s) consists exclusively of GREEN vertices. Observe that $p_{\text{ACCEPT}} = \mathbb{E}_{s \in L^{(r)}}[\eta^{(r)}(s) \cdot \mathbf{1}_{f^{(r)}(s)=C}]$ equals the probability that the QUERY phase accepts.

The exact quantity that we will track is as i increases is the weighted agreement:

$$\alpha^{(i)} = \text{agree}_{\eta^{(i)}}[f^{(i)}, \text{RS}^{(i)}].$$

Notice that

$$\alpha^{(0)} = 1 - \Delta(f^{(0)}, \text{RS}^{(0)}),$$

and the acceptance probability, p_{ACCEPT} satisfies:

$$p_{ACCEPT} \leq \alpha^{(r)}.$$

Our main intermediate claim is that with high probability over the choice of $x^{(i)}, z^{(i)}, B_{z^{(i)}}^{(i)}$, we have that $\alpha^{(i+1)}$ is not much more than $\alpha^{(i)}$. This gives us that p_{ACCEPT} is not much more than $1 - \Delta(f^{(0)}, \text{RS}^{(0)})$, as desired.

5.5.2 Operations AVG and ZERO

We define two important operations.

1. **AVG.** For a function $w : L^{(i-1)} \rightarrow \mathbb{R}$, we define the function $\text{AVG}[w] : L^{(i)} \rightarrow \mathbb{R}$ as follows. Let $s \in L^{(i)}$, and let $\{s_0, s_1\} = (q^{(i-1)})^{-1}(s)$. Then define:

$$\text{AVG}[w](s) = \frac{w(s_0) + w(s_1)}{2}.$$

2. **ZERO.** For a function $w : L^{(i)} \rightarrow \mathbb{R}$, and a set $S \subseteq L^{(i)}$, we define the function $\text{ZERO}[w, S] : L^{(i)} \rightarrow \mathbb{R}$ as follows. For $s \in L^{(i)}$, we set:

$$\text{ZERO}[w, S](s) = \begin{cases} 0 & s \in S \\ w(s) & s \notin S \end{cases}.$$

We can use these two operations to express $\eta^{(i+1)}$ in terms of $\eta^{(i)}$. Let $E^{(i+1)}$ denote the set of all $s \in S^{(i+1)}$ both of whose children are RED (i.e., the test

$$H_{x^{(i)}}[f^{(i)}](s) = f^{(i+1)}(s) \cdot (s - z^{(i)}) + B_{z^{(i)}}^{(i)}(x^{(i)})$$

fails at s).

Define $\theta^{(i+1)} : L^{(i+1)} \rightarrow \mathbb{R}$ by

$$\theta^{(i+1)} = \text{AVG}[\eta^{(i)}].$$

Then we have:

$$\eta^{(i+1)} = \text{ZERO}(\theta^{(i+1)}, E^{(i+1)}).$$

Analogous to our definition of

$$\alpha^{(i)} = \text{agree}_{\eta^{(i)}}(f^{(i)}, \text{RS}^{(i)}),$$

we define

$$\beta^{(i+1)} = \text{agree}_{\theta^{(i+1)}}(H_{x^{(i)}}[f^{(i)}], \{P(Y) \in \mathbb{F}_q[Y] \mid \deg(P) \leq d^{(i+1)} \text{ and } P(z^{(i)}) = B_{z^{(i)}}^{(i)}(x^{(i)})\}).$$

The following two lemmas control the growth of $\alpha^{(i)}$ and $\beta^{(i)}$.

Lemma 5.8. *For all i , with probability at least $1 - \nu^*$ over the choice of $x^{(i)}, z^{(i)}$, we have:*

$$\beta^{(i+1)} \leq \max(\alpha^{(i)}, 1 - \delta^*) + \epsilon.$$

We prove this using Theorem 4.3 in Appendix C.

Lemma 5.9. *For all i ,*

$$\alpha^{(i)} \leq \beta^{(i)}.$$

We prove this using Lemma 5.3 in Appendix C.

We can now complete the proof of Theorem 5.5.

Proof. As observed earlier, $\alpha^{(0)} = 1 - \Delta(f^{(0)}, \text{RS}^{(0)}) < 1 - \delta$.

The two lemmas above imply that with probability at least $1 - r\nu^*$,

$$\alpha^{(r)} \leq \max(\alpha^{(0)}, 1 - \delta^*) + r \cdot \epsilon < (1 - \min(\delta, \delta^*) + r \cdot \epsilon).$$

Finally, we use the observation that $p_{\text{ACCEPT}} \leq \alpha^{(r)}$ to complete the proof. \square

6 The DEEP Algebraic Linking IOP (DEEP-ALI) protocol

The techniques used earlier in Theorem 4.1 and Section 5 can also be used to improve soundness in other parts of an interactive oracle proof (IOP) protocol. We apply them here to obtain a Scalable Transparent IOP of Knowledge (STIK) [BBHR18a, Definition 3.3] with better soundness than the prior state of the art, given in [BBHR18a, Theorem 3.4].

Proof systems typically use a few steps of reduction to convert problems of membership in a nondeterministic language L to algebraic problems regarding proximity of a function (or a sequence of functions) to an algebraic code like Reed-Solomon (or, in earlier works, Reed-Muller). The goal of such a reduction is to maintain a large *proximity gap* γ , meaning that for instances in L , an honest prover will provide information that leads to codewords, whereas for instances not in L , any oracles submitted by the prover will be converted by the reduction, with high probability, to functions that are γ -far from the code. Considerable effort is devoted to increasing γ because it is the input to the proximity protocols (like FRI and DEEP-FRI) and the soundness of those protocols is correlated to γ (as discussed earlier, e.g., in Theorem 5.5).

The STIK protocol is a special case of this paradigm. It requires the prover to provide oracle access to a function $f : D \rightarrow \mathbb{F}$ that is supposedly an RS encoding of a witness for membership of the input instance in L . A set of t -local constraints is applied to f to construct a function $g : D \rightarrow \mathbb{F}$, along with a gap-guarantee: If f is indeed an encoding of a valid witness for the instance, then the resulting function $g : D \rightarrow \mathbb{F}$ is also be a member of an RS code. One of the tests that the verifier performs is a *consistency test* between f and g , and, prior to this work, this consistency test was applied to the functions f and g *directly*. This leads to a rather small gap $\gamma \leq \frac{1}{8}$ which results in a small soundness guarantee from the RPT protocol applied to f, g later on.

In this section we apply the DEEP technique to this setting. After f and g have been provided, the verifier samples a random $z \in \mathbb{F}_q$ and asks for the values of the interpolating polynomials of f, g on all t entries needed to check the consistency test. Our verifier now applies the QUOTIENT operation to f, g , using the information obtained from the prover. Crucially, we prove that a *single* consistency test, conducted over a large domain $D' \supset D$, suffices to improve the proximity gap to roughly $1 - \sqrt{\rho}$, a value that approaches 1 as $\rho \rightarrow 0$. Assuming Conjecture 2.3 the proximity gap is nearly-optimal, at $\gamma \approx 1 - \rho$ (compare with with the value $\gamma \leq 1/8$ obtained by prior works). Details follow.

We focus on the the Algebraic linking IOP protocol (ALI) of [BBHR18a, Theorem B.15], and present a new protocol that we call DEEP-ALI (Protocol 6.4) that obtains the aforementioned improved proximity gap(s).

In what follows, we will first recall (a variant of) the language (or, more accurately, binary relation) which was the input to the ALI protocol of [BBHR18a] and is likewise the input to our DEEP-ALI protocol. The description of the protocol follows in Section 6.2. Its basic properties are specified in Section 6.3 and we analyze its soundness in Theorem 6.2 and Section 6.4.

6.1 The Algebraic Placement and Routing (APR) Relation

In what follows we use the notation \tilde{f} to refer to a polynomial in $\mathbb{F}[x]$. Note that the operator $|_D$ for $D \subseteq \mathbb{F}$ takes a polynomial to a function: $\tilde{f}|_D: D \rightarrow \mathbb{F}$.

We start by defining a simplified version of the Algebraic placement and routing relation (APR). See [BBHR18a, Definition B.10]. In particular, we only use one witness polynomial. This relation will be the input to the reduction used in Protocol 6.4.

Definition 6.1. *The relation R_{APR} is the set of pairs $(\mathfrak{x}, \mathfrak{w})$ satisfying:*

1. **Instance format:** *The instance \mathfrak{x} is a tuple $(\mathbb{F}_q, d, \mathcal{C})$ where:*

- \mathbb{F}_q is a finite field of size q .
- d is an integer representing a bound on the degree of the witness.
- \mathcal{C} is a set of $|\mathcal{C}|$ tuples (M^i, P^i, Q^i) representing constraints. M^i is the mask which is a sequence of field elements $M^i = \{M_j^i \in \mathbb{F}_q\}_{j=1}^{|M^i|}$. P^i is the condition of the constraint which is a polynomial with $|M^i|$ variables. $Q^i \in \mathbb{F}_q[x]$ is the domain polynomial of the constraint which should vanish on the locations where the constraint should hold.

We further introduce the following notation:

- Let $\mathcal{M} = \{M_j^i \mid 1 \leq i \leq |\mathcal{C}| \text{ and } 1 \leq j \leq |M^i|\} \subseteq \mathbb{F}_q$ be the full mask.
 - Let $d_{\mathcal{C}} = \max_i \deg(P^i)$ be the maximal total degree of the P^i s.
 - Let $Q_{\text{lcm}} \in \mathbb{F}_q[x]$ be the least common multiple of the Q^i s.
2. **Witness format:** *The witness \mathfrak{w} is a polynomial $\tilde{f} \in \mathbb{F}_q[x]$. A constraint (M, P, Q) is said to hold at a location $x \in \mathbb{F}_q$ if $P(\tilde{f}(x \cdot M_1), \dots, \tilde{f}(x \cdot M_{|M|})) = 0$. We say that \tilde{f} satisfies the constraint if the constraint holds at every $x \in \mathbb{F}_q$ for which $Q(x) = 0$.*

We say that \mathfrak{w} satisfies the instance if and only if $\deg(\tilde{f}) < d$ and \tilde{f} satisfies all of the constraints.

To see that the notion of the R_{APR} relation defined above is strong enough, we follow the ideas from [BBHR18a] and show a reduction from an Algebraic Intermediate Representation (AIR, see [BBHR18a, Definition B.3]) to an APR. The following uses the notation from [BBHR18a, Definition B.3]. Let $\mathfrak{x} = (\mathbb{F}_q, T, \mathfrak{w}, \mathcal{P}, \mathcal{C}, \mathcal{B})$ be an instance of R_{AIR} . Pick a multiplicative subgroup $\langle \gamma \rangle \subseteq \mathbb{F}_q^\times$ of size $T \cdot \mathfrak{w}$ and pick \tilde{f} such that $\tilde{f}(\gamma^{t\mathfrak{w}+j}) = w_j(t)$ for $t \in [T]$ and $i \in [\mathfrak{w}]$ (here $[n] = \{0, \dots, n-1\}$). For all the constraints in \mathcal{P} , choose the mask $M = \{1, \gamma, \dots, \gamma^{2\mathfrak{w}-1}\}$ and choose the domain polynomial whose zeros are $\{\gamma^{t\mathfrak{w}}\}_{t \in [T-1]}$ ($Q(x) = (x^T - 1)/(x - \gamma^{-\mathfrak{w}})$). Replace each boundary constraint $(i, j, \alpha) \in \mathcal{B}$ with a regular constraint with mask $M = \{1\}$, $P(x) = x - \alpha$ and $Q(x) = x - \gamma^{i\mathfrak{w}+j}$.

6.2 The DEEP-ALI protocol

We now describe our new protocol, that achieves improved soundness, as stated in the following theorem.

Theorem 6.2 (DEEP-ALI soundness). *Fix a code rate $0 < \rho < 1$ and a distance parameter $0 < \delta \leq 1 - \rho$. Let $D, D' \subseteq \mathbb{F}_q$ be two evaluation domains such that $|D| = d\rho^{-1}$ and $|D'| = d \cdot d_C \rho^{-1}$. Let $\text{RPT}_D, \text{RPT}_{D'}$ be two IOPPs with perfect completeness for the codes $\text{RS}[\mathbb{F}_q, D, (d - |\mathcal{M}|)/|D|]$ and $\text{RS}[\mathbb{F}_q, D', (dd_C - 1)/|D'|]$ respectively. Let ϵ, ϵ' be the bounds on the soundness error (acceptance probability) for words that are at least δ -far from the corresponding code. Denote*

$$L = \max\{\mathcal{L}(\mathbb{F}_q, D, d, \delta), \mathcal{L}(\mathbb{F}_q, D', d \cdot d_C, \delta)\}.$$

Then, there exists an IOP for R_{APR} with perfect completeness and soundness error $\epsilon + \epsilon' + \frac{2L^2(d \cdot d_C + \deg(Q_{\text{lcm}}))}{q}$.

Example 6.3. *Fix a code rate $0 < \rho < 1$. Choosing DEEP-FRI as the RPT protocol and setting $\delta = 1 - \sqrt{\rho} - q^{-1/13}$ as in Example 5.6, using ℓ repetitions, we obtain an IOP for R_{APR} with perfect completeness and soundness error that approaches $2\rho^{\ell/2}$ as $q \rightarrow \infty$ assuming the parameters $d, d_C, \deg(Q_{\text{lcm}})$ of the APR are constant with respect to q .*

Proof. Theorem 2.2 implies that $L \leq q^{1/13}/(2\sqrt{\rho}) = O(q^{1/13})$. Hence the expression $2L^2(d \cdot d_C + \deg(Q_{\text{lcm}}))/q$ approaches 0 as $q \rightarrow \infty$. Moreover, Example 5.6 implies that ϵ, ϵ' approach $\rho^{\ell/2}$. \square

We now describe the protocol that achieves the soundness of Theorem 6.2.

Protocol 6.4 (DEEP-ALI).

1. The prover sends an oracle $f : D \rightarrow \mathbb{F}$ (which should be $\tilde{f}|_D$).
2. The verifier sends random coefficients $\alpha = (\alpha_1, \dots, \alpha_{|\mathcal{C}|}) \in \mathbb{F}_q^{|\mathcal{C}|}$.
3. The prover sends an oracle $g_\alpha : D' \rightarrow \mathbb{F}$ (which should be $\tilde{g}_\alpha|_{D'}$, where

$$\tilde{g}_\alpha(x) = \sum_{i=1}^{|\mathcal{C}|} \alpha_i \cdot \frac{P^i(\tilde{f}(x \cdot M_1^i), \dots, \tilde{f}(x \cdot M_{|M^i|}^i))}{Q^i(x)}. \quad (16)$$

Note that $\deg(\tilde{g}_\alpha) < d \cdot d_C$.

4. The verifier sends a random value $z \in \mathbb{F}_q$.
5. Denote $\mathcal{M}_z = \{z \cdot M_j^i \mid 1 \leq i \leq |\mathcal{C}| \text{ and } 1 \leq j \leq |M^i|\}$. The prover sends $a_{\alpha, z} : \mathcal{M}_z \rightarrow \mathbb{F}$ (which should be $\tilde{f}|_{\mathcal{M}_z}$). The verifier deduces $b_{\alpha, z}$, the alleged value of $\tilde{g}_\alpha(z)$, using Equation (16).
6. Let $U(x), Z(x)$ as defined in Section 5.2.1 for $\text{QUOTIENT}(f, a_{\alpha, z})$ and let

$$h^1(x) = h_{\alpha, z}^1(x) = \text{QUOTIENT}(f, a_{\alpha, z}) = \frac{f(x) - U(x)}{Z(x)},$$

$$h^2(x) = h_{\alpha, z}^2(x) = \text{QUOTIENT}(g_\alpha, \{z \mapsto b_{\alpha, z}\}) = \frac{g_\alpha(x) - b_{\alpha, z}}{x - z},$$

and note that the verifier has oracle access to h^1 and h^2 using the oracles f and g_α .

7. They use RPT_D and $\text{RPT}_{D'}$ to prove that h^1 is at most δ -far from $\text{RS}[\mathbb{F}_q, D, (d - |\mathcal{M}|)/|D|]$ (in other words, it is close to a polynomial of degree $< d - |\mathcal{M}|$) and that h^2 is at most δ -far from $\text{RS}[\mathbb{F}_q, D', (dd_C - 1)/|D'|]$.

6.3 Properties of DEEP-ALI

Note that in the original ALI protocol, the equivalent to the expression $P^i(\tilde{f}(x \cdot M_1^i), \dots, \tilde{f}(x \cdot M_{|\mathcal{M}|}^i))/Q^i(x)$ is sampled at Q random locations from the evaluation domain, where Q is the number of queries.

The main idea in DEEP-ALI is to use Quotienting to allow the verifier to choose *one* random element z from the *entire* field, and check the consistency between \tilde{f} and \tilde{g} only at $x = z$.

The fact that DEEP-ALI allows to sample from the entire field introduces several advantages over the ALI protocol from [BBHR18a]:

Soundness As described above, the reduction in ALI has lower bound $1/8$ on the distance from the code for inputs that are not in the language, even for $\rho \rightarrow 0$. In DEEP-ALI the lower bound on the distance is $1 - \sqrt{\rho}$.

Query complexity In ALI the verifier queries $|\mathcal{M}| \cdot Q$ field elements as we need $|\mathcal{M}|$ elements to evaluate $P^i(\tilde{f}(x \cdot M_1^i), \dots, \tilde{f}(x \cdot M_{|\mathcal{M}|}^i))$. In DEEP-ALI the verifier queries $O(|\mathcal{M}| + Q)$ field elements as the evaluation of P^i is done once.

Verifier complexity Previously, the verifier complexity was $\Omega(Q \cdot T_{\text{arith}})$ (where T_{arith} is the arithmetic complexity of evaluating all the constraints). The verifier complexity in DEEP-ALI depends on $Q + T_{\text{arith}}$ as we evaluate the constraints only once.

Prover complexity It is possible to alter Definition 6.1 and DEEP-ALI to work with several witness polynomials f_1, \dots, f_w (as was done in ALI). The prover complexity in this case will depend on $(w\rho^{-1} + d_C\rho^{-1} + T_{\text{arith}}d_C) \cdot d$ instead of $(wd_C\rho^{-1} + T_{\text{arith}}d_C) \cdot d$ (in ALI).

6.4 Soundness analysis

The proof of Theorem 6.2 will follow from the following lemma:

Lemma 6.5. *Let \mathcal{E} be the event that the DEEP-ALI verifier accepts. If*

$$\Pr[\mathcal{E}] \geq \epsilon + \epsilon' + \frac{2L^2(d \cdot d_C + \deg(Q_{\text{lcm}}))}{q},$$

then there exists a polynomial of degree $< d$ satisfying the constraints.

Proof. Let $L(f) \subseteq \text{RS}[\mathbb{F}_q, D, \rho]$ be the set of codewords that are at most δ -far from f . Similarly define $L(g_\alpha)$. We have $|L(f)|, |L(g_\alpha)| \leq L$.

Let \mathcal{E}_1 be the event where the verifier accepts and h^1 and h^2 are at most δ -far from the corresponding codes. Denote $\eta = 2L^2(d \cdot d_C + \deg(Q_{\text{lcm}}))/q$. Then, $\Pr[\mathcal{E}_1] \geq \eta$. \mathcal{E}_1 implies that there exists a polynomial $\tilde{h}^1 = \tilde{h}_{\alpha, z}^1$ of degree $< d - |\mathcal{M}|$ such that $|\{x \in D : \tilde{h}^1(x) \neq \frac{f(x) - U(x)}{Z(x)}\}| < \delta|D|$. Hence $Z(x) \cdot \tilde{h}^1(x) + U(x) \in L(f)$. Similarly there exists a polynomial $\tilde{h}^2 = \tilde{h}_{\alpha, z}^2$ of degree $< d \cdot d_C - 1$ such that $(x - z)\tilde{h}^2(x) + b \in L(g_\alpha)$.

Fix $\tilde{r}^1(x)$ (independent of α and z) to be the element in $L(f)$ maximizing the probability that $\tilde{r}^1(x) = Z(x) \cdot \tilde{h}^1(x) + U(x)$ given \mathcal{E}_1 . Let $\mathcal{E}_2 \subseteq \mathcal{E}_1$ be the event that $\tilde{r}^1(x) = Z(x) \cdot \tilde{h}^1(x) + U(x)$. It follows that $\Pr[\mathcal{E}_2] \geq \eta/L$.

Fix $\tilde{r}_\alpha^2(x) \in L(g_\alpha)$ maximizing the probability that $\tilde{r}_\alpha^2(x) = (x - z)\tilde{h}^2(x) + b$ given \mathcal{E}_2 (note that \tilde{r}_α^2 depends on α as the oracle g_α was sent only after the verifier sent α), and let $\mathcal{E}_3 \subseteq \mathcal{E}_2$ be the event where $\tilde{r}_\alpha^2(x) = (x - z)\tilde{h}^2(x) + b$. We have $\Pr[\mathcal{E}_3] \geq \eta/L^2$. This implies, $\Pr_\alpha[\Pr_z[\mathcal{E}_3] \geq \eta/(2L^2)] \geq \eta/(2L^2)$.

The event \mathcal{E}_3 implies

$$\begin{aligned}\tilde{r}^1|_{\mathcal{M}_z} &= U|_{\mathcal{M}_z} = a_{\alpha,z}, \\ \tilde{r}_\alpha^2(z) &= b_{\alpha,z}.\end{aligned}$$

Recall that $b_{\alpha,z}$ was defined according to (16), so

$$b_{\alpha,z} = \sum_{i=1}^{|\mathcal{C}|} \alpha_i \cdot \frac{P^i(a_{\alpha,z}(z \cdot M_1^i), \dots, a_{\alpha,z}(z \cdot M_{|M^i|}^i))}{Q^i(z)}.$$

Substituting values for $a_{\alpha,z}$ and $b_{\alpha,z}$ and multiplying by $Q_{\text{lcm}}(z)$ we obtain:

$$Q_{\text{lcm}}(z) \cdot \tilde{r}_\alpha^2(z) = \sum_{i=1}^{|\mathcal{C}|} \alpha_i \cdot P^i(\tilde{r}^1(z \cdot M_1^i), \dots, \tilde{r}^1(z \cdot M_{|M^i|}^i)) \cdot \frac{Q_{\text{lcm}}(z)}{Q^i(z)}. \quad (17)$$

Both sides of the equation are polynomials of degree $< d \cdot d_{\mathcal{C}} + \deg(Q_{\text{lcm}})$ in z . For every α for which $\Pr_z[\mathcal{E}_3] \geq \eta/(2L^2) = (d \cdot d_{\mathcal{C}} + \deg(Q_{\text{lcm}}))/q$, we have at least $d \cdot d_{\mathcal{C}} + \deg(Q_{\text{lcm}})$ many z 's satisfying (17) and thus the two polynomials in (17) are identical. Let $G_\alpha(x)$ denote the the right-hand side of (17) (replacing z with x).

So far we have:

$$\Pr[G_\alpha(x) \text{ is divisible by } Q_{\text{lcm}}(x)] \geq \eta/(2L^2) > 1/q.$$

Note that the set of α 's satisfying this event forms a vector space. If its dimension was less than $|\mathcal{C}|$ then the probability would have been $\leq 1/q$. Hence this event holds for every α . Substituting the elements of the standard basis, we get that for every $1 \leq i \leq |\mathcal{C}|$,

$$P^i(\tilde{r}^1(x \cdot M_1^i), \dots, \tilde{r}^1(x \cdot M_{|M^i|}^i)) \cdot \frac{Q_{\text{lcm}}(x)}{Q^i(x)} \text{ is divisible by } Q_{\text{lcm}}(x).$$

Substituting any x for which $Q^i(x) = 0$ gives $P^i(\tilde{r}^1(x \cdot M_1^i), \dots, \tilde{r}^1(x \cdot M_{|M^i|}^i)) = 0$ which implies that \tilde{r}^1 satisfies all the constraints, as required. \square

6.5 Further optimizations for practical implementation

As we saw, it makes sense to work with several witness polynomials rather than one, as it improves the prover complexity. Another optimization is to apply the RPT only once for both h^1 and h^2 by taking a random linear combination of the two (and using Theorem 4.1). To make this work, the prover writes the degree $< d \cdot d_{\mathcal{C}}$ polynomial $\tilde{g}(x)$ as:

$$\tilde{g}(x) = \sum_{i=0}^{d_{\mathcal{C}}-1} x^i \tilde{g}_i(x^{d_{\mathcal{C}}}),$$

where the \tilde{g}_i s are of degree $< d$, and it sends oracles to $\tilde{g}_i|_D$ instead of $\tilde{g}|_{D'}$. In total, we will have to run RPT on $w + d_C$ polynomials of degree $< d$, so we choose only one evaluation domain $D \subseteq \mathbb{F}_q$ satisfying $|D| = d\rho^{-1}$.

References

- [AHIV17] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In *Proceedings of the 24th ACM Conference on Computer and Communications Security*, October 2017.
- [BBHR18a] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity. Cryptology ePrint Archive, Report 2018/046, 2018. Available at <https://eprint.iacr.org/2018/046>.
- [BBHR18b] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast Reed-Solomon Interactive Oracle Proofs of Proximity. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP)*, 2018.
- [BCF⁺16] Eli Ben-Sasson, Alessandro Chiesa, Michael A. Forbes, Ariel Gabizon, Michael Riabzev, and Nicholas Spooner. On probabilistic checking in perfect zero knowledge. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:156, 2016.
- [BCR⁺18] Eli Ben-Sasson, Alessandro Chiesa, Michael Riabzev, Nicholas Spooner, Madars Virza, and Nicholas P. Ward. Aurora: Transparent succinct arguments for R1CS. *IACR Cryptology ePrint Archive*, 2018:828, 2018.
- [BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part II*, pages 31–60, 2016.
- [BGH⁺06] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil P. Vadhan. Robust PCPs of proximity, shorter PCPs, and applications to coding. *SIAM Journal on Computing*, 36(4):889–974, 2006.
- [BKS18a] Eli Ben-Sasson, Swastik Kopparty, and Shubhangi Saraf. Worst-case to average case reductions for the distance to a code. In *33rd Computational Complexity Conference, CCC 2018, June 22-24, 2018, San Diego, CA, USA*, pages 24:1–24:23, 2018.
- [BKS18b] Eli Ben-Sasson, Swastik Kopparty, and Shubhangi Saraf. Worst-case to average case reductions for the distance to a code. *Electronic Colloquium on Computational Complexity (ECCC)*, 25:90, 2018.
- [BS08] Eli Ben-Sasson and Madhu Sudan. Short PCPs with polylog query complexity. *SIAM Journal on Computing*, 38(2):551–607, 2008. Preliminary version appeared in STOC '05.
- [BSKR10] Eli Ben-Sasson, Swastik Kopparty, and Jaikumar Radhakrishnan. Subspace polynomials and limits to list decoding of reed-solomon codes. *IEEE Trans. Information Theory*, 56(1):113–120, 2010.

- [Din07] Irit Dinur. The PCP theorem by gap amplification. *Journal of the ACM*, 54(3):12, 2007.
- [Gur07] Venkatesan Guruswami. Algorithmic results in list decoding. *Foundations and Trends® in Theoretical Computer Science*, 2(2):107–195, 2007.
- [LFKN92] Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992.
- [Raz87] Alexander A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987.
- [Rot06] Ron M. Roth. *Introduction to coding theory*. Cambridge University Press, 2006.
- [RRR16] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 49–62, 2016.
- [RVW13] Guy N. Rothblum, Salil Vadhan, and Avi Wigderson. Interactive proofs of proximity: delegating computation in sublinear time. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 793–802. ACM, 2013.
- [RW14] Atri Rudra and Mary Wootters. Every list-decodable code for high noise has abundant near-optimal rate puncturings. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 764–773, 2014.

A Proof of the DEEP lemma for general codes

Proof of Lemma 4.6. To simplify notation set $\eta = \max\left(2L_\delta^* \left(\frac{\sigma}{N} + \epsilon\right)^{1/3}, \frac{4}{\epsilon^2 q}\right)$, and let $u_x = u^* + xu$. Let $\mathcal{E}[x, z]$ denote the event “ $\exists v \in \text{List}(u_x, V, \delta), \langle v, z \rangle = B_z(x)$ ”.

The assumption of Equation (10) now reads as

$$\Pr_{x \in \mathbb{F}_q, z \in S} [\mathcal{E}[x, z]] \geq \eta.$$

Thus we get,

$$\Pr_{x \in \mathbb{F}_q} [\Pr_{z \in S} [\mathcal{E}[x, z]] \geq \eta/2] \geq \eta/2 \tag{18}$$

Let

$$A = \left\{ x \in \mathbb{F}_q \mid \Pr_{z \in S} [\mathcal{E}[x, z]] \geq \eta/2 \right\}$$

and notice $|A| \geq \eta q/2$.

For $x \in \mathbb{F}_q$, pick $v_x \in V$ to be a member of $\text{List}(u_x, V, \delta)$ that maximizes $\Pr_{z \in S} [P(z) = B_z(x)]$. Let $S_x = \{z \in S \mid \langle v_x, z \rangle = B_z(x)\}$ and set $\mu_x = |S_x|/s$. By definition, $|\text{List}(u_x, V, \delta)| \leq L_\delta^*$, and so by the pigeonhole principle, for each $x \in A$ we have $\mu_x \geq \frac{\eta}{2L_\delta^*}$.

For x, β, γ picked uniformly from A we have

$$\begin{aligned}
\mathbf{E}_{x, \beta, \gamma \in A} \left[\frac{|S_x \cap S_\beta \cap S_\gamma|}{s} \right] &= \mathbf{E}_{z \in S, x, \beta, \gamma \in \mathbb{F}_q} [\mathbf{1}_{z \in S_x \cap S_\beta \cap S_\gamma}] \\
&= \mathbf{E}_{z \in S} [\mathbf{E}_{x \in \mathbb{F}_q} [\mathbf{1}_{z \in S_x}]^3] \\
&\geq \mathbf{E}_{z \in S, x \in \mathbb{F}_q} [\mathbf{1}_{z \in S_x}]^3 \\
&\geq \left(\frac{\eta}{2L_\delta^*} \right)^3 \\
&> \frac{\sigma}{N} + \epsilon.
\end{aligned}$$

The second equality above follows from the independence of x, β, γ . The first inequality is an application of Jensen's inequality and the last inequality is by assumption on η .

Thus

$$\Pr_{x, \beta, \gamma} [|S_x \cap S_\beta \cap S_\gamma| > \sigma] \geq \epsilon.$$

Note that $\Pr_{x, \beta, \gamma} [x, \beta, \gamma \text{ are not all distinct}] < 3/|A|$. Since $|A| \geq \eta q/2 \geq 2/\epsilon^2 \geq 6/\epsilon$ we have $3/|A| \leq \epsilon/2$. Thus $\Pr_{x, \beta, \gamma} [x, \beta, \gamma \text{ are all distinct and } |S_x \cap S_\beta \cap S_\gamma| > \sigma] \geq \epsilon/2$.

This means that there are distinct x_0, β_0 such that

$$\Pr_\gamma [|S_{x_0} \cap S_{\beta_0} \cap S_\gamma| > d] \geq \epsilon/2.$$

Consider some γ where this happens. Let $\tilde{S} = S_{x_0} \cap S_{\beta_0} \cap S_\gamma$. Extend each of u^*, u to functions over domain S by defining for all $z \in S \setminus [n]$ $u^*(z) = B_z(0)$ and $u(z) = B_z(1)$, and for $x \in \mathbb{F}_q$ let $u_x(z) = u^*(z) + xu(z)$. Since V is systematic, we define $v_x(z) = \langle v_x|_{[k]}, z \rangle$ and thus extend v_x to domain \tilde{S} . By construction we know

$$(x_0, u_{x_0}), (\beta_0, u_{\beta_0}), (\gamma, u_\gamma)$$

are collinear. So, in particular,

$$(x_0, u_{x_0}|_{\tilde{S}}), (\beta_0, u_{\beta_0}|_{\tilde{S}}), (\gamma, u_\gamma|_{\tilde{S}}) \in \mathbb{F}_q \times \mathbb{F}_q^{\tilde{S}}$$

are likewise collinear, as a special case. By definition of \tilde{S} , we get that:

$$(x_0, v_{x_0}|_{\tilde{S}}), (\beta_0, v_{\beta_0}|_{\tilde{S}}), (\gamma, v_\gamma|_{\tilde{S}}) \in \mathbb{F}_q \times \mathbb{F}_q^{\tilde{S}}$$

are also collinear. Since $|\tilde{S}| > \sigma$ and S is σ -robust we conclude that v_γ is uniquely determined by $v_\gamma|_{\tilde{S}}$. This allows us to conclude that

$$(x_0, v_{x_0}), (\beta_0, v_{\beta_0}), (\gamma, v_\gamma) \in \mathbb{F}_q \times \mathbb{F}_q^n$$

are all collinear, recalling that $v_{x_0} \in \text{List}(u_{x_0}, V, \delta)$.

Thus, an $\epsilon/2$ -fraction of the $\gamma \in A$ have the ‘‘good’’ property that (γ, v_γ) is on the line passing through (x_0, v_{x_0}) and (β_0, v_{β_0}) . Write this line as $v^* + xv$ and notice that for all ‘‘good’’ γ we have $v_\gamma = v^* + \gamma v$. Let $A' \subseteq A$ denote the set of good elements for this line, recording that

$|A'| \geq |A| \cdot \epsilon/2 \geq 1/\epsilon$. By definition of $\text{List}(u_x, V, \delta)$ and the assumption $v_x \in \text{List}(u_x, V, \delta)$, we have that $\Delta(u_x, v_x) < \delta$ for $x \in A'$.

Consider the set $C \subset [n]$ defined by

$$C = \{y \in [n] \mid u^*(y) = v^*(y) \text{ AND } u(y) = v(y)\}.$$

For each $y \in [n] \setminus C$ there exists at most a single value of $x \in \mathbb{F}_q$ satisfying $u_x(y) = v_x(y)$ because

$$u_x(y) - v_x(y) = (u^*(y) - v^*(y)) + x \cdot (u(y) - v(y))$$

has at most one value x on which it vanishes. This implies

$$\delta \geq \mathbf{E}_{x \in A'}[\Delta_{[n]}(u_x, v_x)] \geq \frac{|[n] \setminus C|}{n} \cdot \left(1 - \frac{1}{|A'|}\right) \geq \left(1 - \frac{|C|}{n}\right) \cdot (1 - \epsilon) \geq 1 - \frac{|C|}{n} - \epsilon.$$

Rearranging, we get $\frac{|C|}{n} \geq 1 - (\delta + \epsilon)$ and this completes the proof. \square

B The algebraic hash function

We now describe the algebraic hash function H_x .

The description of the hash function requires fixing some choices of certain subspaces. For each $i \in [0, r]$ we choose \mathbb{F}_2 -subspaces $L_0^{(i)}$ and $L^{(i)}$, satisfying the following properties.

1. $L_0^{(i)} \subseteq L^{(i)}$ with $\dim(L_0^{(i)}) = 1$,
2. $L^{(i+1)} = q^{(i)}(L^{(i)})$, where $q^{(i)}(X)$ is the *subspace polynomial* of $L_0^{(i)}$,

$$q^{(i)}(X) = \prod_{\alpha \in L_0^{(i)}} (X - \alpha),$$

thus this is an \mathbb{F}_2 -linear map with kernel $L_0^{(i)}$. In particular, $\dim(L^{(i+1)}) = \dim(L^{(i)}) - 1$.

Let $\mathcal{S}^{(i)}$ denote the set of cosets of $L_0^{(i)}$ contained in $L^{(i)}$.

Given $x \in \mathbb{F}$ and $f : L^{(i)} \rightarrow \mathbb{F}$, the hash of f with seed x is defined to be the function $H_x[f] : L^{(i+1)} \rightarrow \mathbb{F}$ as follows. For $s \in L^{(i+1)}$, let $s_0, s_1 \in L^{(i)}$ be the two roots of $q^{(i)}(X) - s$. Let $P_{f,s}(X) \in \mathbb{F}[X]$ be the unique degree ≤ 1 polynomial satisfying

$$P_{f,s}(s_0) = f(s_0),$$

$$P_{f,s}(s_1) = f(s_1).$$

Then we define

$$H_x[f](s) = P_{f,s}(x). \tag{19}$$

Observe that $H_x[f](s)$ can be computed by querying f on the set $\{s_0, s_1\}$ (this set is a coset of $L_0^{(i)}$, and we denote it by $S_s^{(i)}$).

To understand H_x better, it is instructive to see what it does to $\text{RS}^{(i)}$. Let $f \in \text{RS}^{(i)}$. The underlying polynomial $f(X)$ thus has degree at most $\rho|L^{(i)}|$. We may write $f(X)$ in base $q^{(i)}(X)$ as:

$$f(X) = a_0(X) + a_1(X)q^{(i)}(X) + \dots + a_t(X)(q^{(i)}(X))^t, \tag{20}$$

where each $a_i(X)$ has degree at most 1, and $t \leq \rho|L^{(i)}|/2$. Since the polynomials $f(X)$ and $P_{f,s}(X)$ agree on the roots of $q(X) - s$, we get that $f(X) \equiv P_{f,s}(X) \pmod{(q^{(i)}(X) - s)}$. From Equation (20), we get that

$$P_{f,s}(X) = a_0(X) + a_1(X)s + \dots + a_t(X)s^t.$$

In particular, for all $x \in \mathbb{F}$,

$$H_x[f](s) = P_{f,s}(x) = a_0(x) + a_1(x)s + \dots + a_t(x)s^t,$$

and thus

$$H_x[f] \in \text{RS}^{(i+1)}.$$

C Proof of Lemma 5.8 and Lemma 5.9

We first prove Lemma 5.8.

Proof. Set $\gamma = \max(\alpha^{(i)}, 1 - \delta^*)$.

For simplicity, denote $f^{(i)}$ by f .

Recall the notation $P_{f,s}$ from the definition of the algebraic hash function H_x in Section B. We have that for each $s \in L^{(i+1)}$, $H_x[f](s) = P_{f,s}(x)$ is a linear function of x . Thus we can write $H_x[f] = u^* + xu$ for $u^*, u \in \mathbb{F}_q^{L^{(i+1)}}$, and for any fixed s , we have the formal polynomial equality $P_{f,s}(X) = u^*(s) + Xu(s)$.

We are interested in bounding the probability of the event $\beta^{(i+1)} > \gamma + \epsilon$. In other words, we want to bound the probability that there exists a polynomial $Q(Y) \in \mathbb{F}_q[Y]$ with $\deg(Q) < d^{(i+1)} + 1$ such that:

- $\text{agree}_{\theta^{(i+1)}}(u^* + xu, Q) > \gamma + \epsilon$,
- $Q(z^{(i)}) = B_{z^{(i)}}^{(i)}(x)$.

This is exactly the scenario of Theorem 4.3. That Lemma tells us that if the probability in question is larger than ν^* , then there exist polynomials $P(Y), P^*(Y)$ of degree $\leq d^{(i+1)}$ and a set $T \subseteq L^{(i+1)}$ such that:

- $$\frac{1}{|L^{(i+1)}|} \sum_{s \in L^{(i+1)}} \theta^{(i+1)} > \gamma,$$
- $u|_T = P|_T$,
- $u^*|_T = P^*|_T$.

Let

$$\hat{P}(X, Y) \triangleq P^*(Y) + X \cdot P(Y)$$

and notice that $\deg_X(\hat{P}) \leq 1$, $\deg_Y(\hat{P}) \leq d^{(i+1)}$.

Consider the polynomial $R(X) \triangleq \hat{P}(X, q^{(i)}(X))$. We have

$$\deg(R) \leq 2d^{(i+1)} + 1 = d^{(i)} - 1 < d^{(i)}.$$

We claim that R agrees with f on $\tilde{T} = \bigcup_{s \in T} S_s^{(i)}$.

Take any $s \in T$ and let $S_s^{(i)} = \{s_0, s_1\} \in \mathcal{S}^{(i)}$ be the pair of roots of the polynomial $q^{(i)}(X) - s$.

First we show that the polynomials $P_{f,s}(X)$ and $\hat{P}(X, s)$ are identical. Indeed, $\hat{P}(X, s) = P^*(s) + XP(s) = u^*(s) + Xu(s) = P_{f,s}(X)$. It follows that

$$f(s_0) = \hat{P}(s_0, s) = \hat{P}(s_0, q^{(i)}(s_0)) = R(s_0)$$

and similarly $f(s_1) = R(s_1)$. Therefore, R and f agree on \tilde{T} , as claimed.

We now use the above information to show that $\alpha^{(i)} = \text{agree}_{\eta^{(i)}}(f, R) > \gamma$, which contradicts the definition of γ . Indeed,

$$\begin{aligned} \text{agree}_{\eta^{(i)}}(f, R) &= \frac{1}{|L^{(i)}|} \sum_{r \in L^{(i)} | f(r)=R(r)} \eta^{(i)}(r) \\ &\geq \frac{1}{|L^{(i)}|} \sum_{r \in \tilde{T}} \eta^{(i)}(r) \\ &= \frac{1}{|L^{(i)}|} \sum_{s \in T} \sum_{r \in S_s^{(i)}} \eta^{(i)}(r) \\ &= \frac{1}{|L^{(i)}|} \sum_{s \in T} |S_s^{(i)}| \cdot \theta^{(i)}(s) \quad \text{Since } \theta(s) \text{ equals the average of } \eta(r) \mid r \in S_s^{(i)} \\ &= \frac{1}{|L^{(i+1)}|} \sum_{s \in T} \theta^{(i)}(s) \\ &> \gamma. \end{aligned}$$

This is the desired contradiction. □

Next we prove Lemma 5.9.

Proof. By definition,

$$\beta^{(i)} = \text{agree}_{\theta^{(i)}}(H_{x^{(i-1)}}[f^{(i-1)}], \{P(Y) \in \mathbb{F}_q[Y] \mid \deg(P) \leq d^{(i)} \text{ and } P(z^{(i-1)}) = B_{z^{(i-1)}}^{(i-1)}(x^{(i-1)})\})$$

Next, by the properties of quotienting, Lemma 5.3,

$$\begin{aligned} \beta^{(i)} &= \text{agree}_{\theta^{(i)}}(H_{x^{(i-1)}}[f^{(i-1)}], \{P(Y) \in \mathbb{F}_q[Y] \mid \deg(P) \leq d^{(i)} \text{ and } P(z^{(i-1)}) = B_{z^{(i-1)}}^{(i-1)}(x^{(i-1)})\}) \\ &= \text{agree}_{\theta^{(i)}}(\text{QUOTIENT}(H_{x^{(i-1)}}[f^{(i-1)}], z^{(i-1)}, B_{z^{(i-1)}}^{(i-1)}(x^{(i-1)})), \{P(Y) \in \mathbb{F}_q[Y] \mid \deg(P) \leq d^{(i)} - 1\}). \end{aligned}$$

Now observe that $\eta^{(i)}$ is obtained from $\theta^{(i)}$ by zeroing out coordinates in $E^{(i)}$, and the only coordinates where $f^{(i)}$ can differ from $\text{QUOTIENT}(H_{x^{(i-1)}}[f^{(i-1)}], z^{(i-1)}, B_{z^{(i-1)}}^{(i-1)}(x^{(i-1)}))$ are in $E^{(i)}$. Thus:

$$\begin{aligned} \beta^{(i)} &\geq \text{agree}_{\theta^{(i)}}(f^{(i)}, \{P(Y) \in \mathbb{F}_q[Y] \mid \deg(P) \leq d^{(i)} - 1\}) \\ &= \text{agree}_{\theta^{(i)}}(f^{(i)}, \text{RS}^{(i)}) \\ &= \alpha^{(i)}. \end{aligned}$$

This completes the proof. □