Electronic Colloquium on Computational Complexity, Report No. 169 (2025)

On Proximity Gaps for Reed–Solomon Codes

Eli Ben-Sasson* Dan Carmon* Ulrich Haböck*

Swastik Kopparty† Shubhangi Saraf‡

November 7, 2025

Abstract

This paper is about the proximity gaps phenomenon for Reed-Solomon codes. Very roughly, the proximity gaps phenomenon for a code $\mathcal{C} \subseteq \mathbb{F}_q^n$ says that for two vectors $f, g \in \mathbb{F}_q^n$, if sufficiently many linear combinations $f + z \cdot g$ (with $z \in \mathbb{F}_q$) are close to $\mathcal{C}$ in Hamming distance, then so are both f and g, up to a proximity loss of ε^*. Determining the optimal quantitative form of proximity gaps for Reed–Solomon codes has recently become of great interest because of applications to interactive proofs and cryptography, and in particular, to scalable transparent arguments of knowledge (STARKs) and other modern hash based argument systems used on blockchains today.

Our main results show improved positive and negative results for proximity gaps for Reed–Solomon codes of constant relative distance $\delta \in (0, 1)$.

- For proximity gaps up to the unique decoding radius $\delta/2$, we show that arbitrarily small proximity loss $\varepsilon^* > 0$ can be achieved with only $O_{\varepsilon^*}(1)$ exceptional z's (improving the previous bound of $O(n)$ exceptions).
- For proximity gaps up to the Johnson radius $J(\delta)$, we show that proximity loss $\varepsilon^* = 0$ can be achieved with only $O(n)$ exceptional z's (improving the previous bound of $O(n^2)$ exceptions). This significantly reduces the soundness error in the aforementioned arguments systems.
- In the other direction, we show that for some Reed–Solomon codes and some δ, proximity gaps at or beyond the Johnson radius $J(\delta)$ with arbitrarily small proximity loss ε^* needs to have at least $\Omega(n^{1.99})$ exceptional z's.
- More generally, for all constants τ, we show that for some Reed–Solomon codes and some $\delta = \delta(\tau)$, proximity gaps at radius $\delta - \Omega_\tau(1)$ with arbitrarily small proximity loss ε^* needs to have n^τ exceptional z's.
- Finally, for all Reed–Solomon codes, we show that improved proximity gaps imply improved bounds for their list-decodability. This shows that improved bounds on the list-decoding radius of Reed–Solomon codes is a prerequisite for any new proximity gaps results beyond the Johnson radius.

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

†Department of Mathematics and Department of Computer Science, University of Toronto, Toronto, Canada. swastik.kopparty@utoronto.ca. Research supported by an NSERC grant.

‡Department of Mathematics and Department of Computer Science, University of Toronto, Toronto, Canada. Email: shubhangi.saraf@utoronto.ca. Research supported by an NSERC grant.

ISSN 1433-8092

Contents

1 Introduction

Reed–Solomon codes are classical error-correcting codes based on univariate polynomial evaluation, and are of central importance in both the theory and practice of coding theory. In the past few decades, they have also played a very important role in complexity theory and cryptography, utilizing their strong error-correcting properties combined with the universal expressive power of polynomials.

This paper is about the *proximity gaps* phenomenon for Reed–Solomon codes. Very roughly, the proximity gaps phenomenon for a code $\mathcal{C} \subseteq \mathbb{F}_q^n$ says that for two vectors $f, g \in \mathbb{F}_q^n$, if many linear combinations $f + z \cdot g$ (with $z \in \mathbb{F}_q$) are close to $\mathcal{C}$ in Hamming distance, then so are both f and g. The crucial quantitative aspects of this are: how close is "close", and how many is "many"? The question of understanding the optimal quantiative form of proximity gaps for Reed–Solomon codes is of great interest, has many potential applications, and is wide open.

Proximity gaps for linear codes in general, and Reed–Solomon codes in particular, have been very actively studied in recent years, especially in connection to STARKs and other modern hash based argument systems that are currently in use in blockchains. They capture interesting geometry of how the code $\mathcal{C}$ sits in the Hamming space $\mathbb{F}_q^n$: $\mathcal{C}$ has proximity gaps if for every line ℓ in the space $\mathbb{F}_q^n$, either all the points of ℓ are close to $\mathcal{C}$, or else most points of ℓ are far from $\mathcal{C}$. Consequently, a nominal one million dollar prize has been recently announce by the Ethereum Foundation to resolve these very questions.

Proximity gaps are also closely related to another fundamental geometric property of a code: list-decodability. A code $\mathcal{C} \subseteq \mathbb{F}_q^n$ has good list-decodability if for any point f in $\mathbb{F}_q^n$, there cannot be too many codewords of $\mathcal{C}$ close to it. Again the same quantitative aspects arise: how close is "close", and how many is "many"? The question of determining the optimal quantitative form of list decodability of Reed–Solomon codes is also of great interest, is much older, has been very well studied, and is wide open.

The study of proximity gaps originated in the works of Rothblum, Vadhan and Wigderson [RVW13] and Ames, Hazay, Ishai and Venkitasubramanian [AHIV17], who used it for delegating computation in sublinear time. The results of [RZ18, BKS18, BGKS20] further explored the quantitative aspects of proximity gaps for general linear codes. The work of Ben-Sasson, Carmon, Ishai, Kopparty and Saraf [BCI$^+$20] gave the state of the art proximity gaps for Reed–Solomon codes. These results were proved using *algorithms* for decoding Reed–Solomon codes: the proof of proximity gaps, which is a purely combinatorial result, used the Berlekamp–Welch unique decoding and Guruswami–Sudan list-decoding algorithms for Reed–Solomon codes. These results enabled new applications in interactive proofs, distributed storage, and cryptography. In particular, they led to the strongest known soundness analysis of the FRI protocol [BBHR18] for proving proximity to Reed–Solomon codes, a widely used protocol in modern argument systems, such as STARKs [BSBHR18], and validity proofs for blockchains.

Since [BCI$^+$20], several further applications of proximity gaps for Reed–Solomon codes have been discovered.

- In [ACY23], Arnon, Chiesa and Yogev used proximity gaps for Reed–Solomon codes to give interactive protocols for NP in the IOP model with linear proof size and inverse polynomial soundness.
- The STIR [ACFY24] and WHIR [ACFY25] protocols of Arnon, Chiesa, Fenzi and Yogev give improved theoretical and practical (1) protocols for proving proximity to (constrained) Reed–Solomon codes, and (2) polynomial commitment schemes.
- In [MZ25], Minzer and Zheng gave stronger round-by-round soundness for languages in NP, as well as a new protocol for further improved soundness for proving proximity to Reed–Solomon codes.
- In [Hab24], Haböck used proximity gaps for Reed–Solomon codes to give a strong soundness analysis for Basefold [ZCF23], a polynomial commitment scheme.
- Finally, the FRI protocol, first used in the Stone STARK prover of StarkWare as a scaling solution for Ethereum, has been widely adopted as the proximity proving protocol for a number of active

or upcoming blockchain "ZKVMs", such as S-two, SP1, Risc-0, Miden, Polygon Hermez, Plonky3, OpenVM, Matterlabs, Lita, Miden, Brevis, Linea and Zisk.

Given all these applications, and the strong connections to basic questions about list-decoding of Reed–Solomon codes, it is thus of great interest to prove quantitatively stronger proximity gaps for Reed–Solomon codes, as well as to understand the limits of what kind of proximity gaps are possible.

1.1 Error-Correcting Codes and Proximity Gaps

For an alphabet Σ and strings $f, g \in \Sigma^n$, we define the relative Hamming distance $\Delta(f, g)$ to be the fraction of $i \in [n]$ for which $f_i \neq g_i$. A linear code $\mathcal{C}$ over $\mathbb{F}_q$ of length n is a linear subspace of $\mathbb{F}_q^n$, and the minimum distance of $\mathcal{C}$ is the smallest value of $\Delta(x, y)$ as x, y vary over distinct elements of $\mathcal{C}$.

Let $\mathcal{C} \subseteq \mathbb{F}_q^n$ be a linear code with minimum distance δ. The minimum distance property immediately implies that for any $\gamma < \delta/2$ (the *unique decoding radius*), any Hamming ball of radius γ in $\mathbb{F}_q^n$ contains at most 1 codeword of $\mathcal{C}$.

One can also say something about Hamming balls of larger radii – this is related to list-decoding. Define *the Johnson radius* $J(\delta) = 1 - \sqrt{1-\delta}$ (and note that $\delta/2 < J(\delta) < \delta$). Then for any $\eta > 0$, the Johnson bound [Joh62] implies that any Hamming ball in $\mathbb{F}_q^n$ of radius $\gamma = J(\delta) - \eta$ contains at most $O_\eta(1)$ codewords of $\mathcal{C}$.

For Reed–Solomon codes of minimum distance δ, it is an extremely interesting question whether all Hamming balls in $\mathbb{F}_q^n$ of radius γ (for some γ larger than $J(\delta)$) are guaranteed to contain at most $\mathsf{poly}(n)$ codewords. This is the problem of determining the list-decoding radius of Reed–Solomon codes. For some Reed–Solomon codes, it is known that one can take γ arbitrarily close to δ; that this may be true for all Reed–Solomon codes is a tantalizing possibility.

For a general linear code $\mathcal{C}$, if $f, g \in \mathbb{F}_q^n$ are both in $\mathcal{C}$, then of course every linear combination $f + z \cdot g$ (with $z \in \mathbb{F}_q$) also lies in $\mathcal{C}$. The proximity gaps phenomenon [BCI+20] (see also [RVW13, AHIV17, BKS18, BGKS20]) provides a robust converse to this – it says that if for many z we have that $f + z \cdot g$ is close to $\mathcal{C}$, then f and g are themselves both close to $\mathcal{C}$ in a certain correlated manner.

Before the formal definition, we quickly introduce a little bit of notation. For two strings $f, g \in \Sigma^n$, we let $[f, g]$ denote the interleaved/correlated vector $h \in (\Sigma^2)^n$ with $h_i = (f_i, g_i)$. In the Hamming metric space $(\Sigma^2)^n$, the alphabet is Σ^2: thus $\Delta([f, g], [u, v]) \leq \lambda$ if and only if there exists some subset $S \subseteq [n]$, with $|S| \geq (1 - \lambda)n$, such that $f|_S = u|_S$ and $g|_S = v|_S$. Finally, for a code $\mathcal{C} \subseteq \Sigma^n$, we define the interleaved code $\mathcal{C}^2 \subseteq (\Sigma^2)^n$, given by:

$$\{[u, v] \in (\Sigma^2)^n \mid u, v \in \mathcal{C}\}.$$

It is easy to see that if $\Delta([f, g], \mathcal{C}^2) \leq \gamma$, then for all $z \in \mathbb{F}_q$, we have $\Delta(f + z \cdot g, \mathcal{C}) \leq \gamma$.

Definition 1.1 (Proximity gaps)**.** *We say the linear code $\mathcal{C} \subseteq \mathbb{F}_q^n$ has* proximity gaps *up to radius $\gamma \in [0, 1]$, with parameters $a \in \mathbb{N}$ and $\varepsilon^* \in [0, 1]$, if, for every $f, g \in \mathbb{F}_q^n$ and every $\gamma' \in [0, \gamma]$ we have the following:*

- *Whenever*
$$|\{z \in \mathbb{F}_q \mid \Delta(f + z \cdot g, \mathcal{C}) \leq \gamma'\}| \geq a,$$
then
$$\Delta([f, g], \mathcal{C}^2) \leq \gamma' + \varepsilon^*.$$

The quantity $\frac{a}{q}$ is called the soundness error*. The quantity ε^* is called the* proximity loss*.*

In applications, we will usually need γ to be a constant in $(0, 1)$, and proximity loss to tend to 0. Some applications need proximity loss equal to 0. There have been many results on proximity gaps, for both

general linear codes and Reed–Solomon codes. We give a quick summary of the prior work before we state our results.

Arbitrary linear codes The first result of a proximity gap flavor was a very influential lemma of Rothblum, Vadhan, Wigderson [RVW13]. They showed that that for any linear code $\mathcal{C}$ (with no assumption about the distance), proximity gaps holds with arbitrary γ, $a = 2$ and proximity loss $\varepsilon^* = \gamma$. This was used in their results on delegating computation in sublinear time. A generalization of the [RVW13] lemma was given in [BKS18], using ideas from list-decoding and the Johnson bound. This again holds for every linear code $\mathcal{C}$ and arbitrary γ: the result said that proximity gaps holds with any a and proximity loss $\varepsilon^* = \frac{1}{a-1}(\gamma - \gamma^2)$. Qualitatively, this gives a nontrivial proximity gap result even for very large radii γ close to 1. Notice that both these results do not have vanishing proximity loss when γ is a fixed constant in $(0, 1)$.

Linear codes with good distance A different line of work extending [RVW13] showed proximity gap results for linear codes with good distance δ, provided the radius γ is small enough in terms of δ. The main new feature was that these results could achieve vanishing proximity loss, and this feature enables many more applications.

Ames, Hazay, Ishai and Venkitasubramaniam [AHIV17], in their work on sublinear time delegation (Ligero), showed that for any linear code $\mathcal{C}$ of minimum distance δ, $\mathcal{C}$ has proximity gaps at any radius $\gamma < \delta/4$ with proximity loss $\varepsilon^* = 0$ and any $a = \gamma n + 2$. The same argument shows shows that $\mathcal{C}$ has proximity gaps at radius $\gamma < \delta/4$ with arbitrarily small proximity loss ε^* and $a = 1 + \frac{\gamma}{\varepsilon^*}$. The $\delta/4$ threshold was improved to $\delta/3$ by Roth–Zemor [RZ18] and Ben-Sasson–Kopparty–Saraf [BKS18].

This latter work [BKS18] also showed a proximity gap result at a radius γ that could get close to 1: for any linear code of minimum distance δ, for $\gamma = J_2(\delta) - \eta$ (where $J_2(\delta) = J(J(\delta)) = 1 - (1 - \delta)^{1/4}$ is the "double Johnson" radius, and $\eta > 0$ is arbitrary), there are proximity gaps with proximity loss ε^* and[1] $a = O_\eta(1/\varepsilon^*) \cdot \gamma$. This result is incomparable to the above mentioned result on proximity gaps up to radius $\delta/3$: for small δ we have $\delta/3 > J_2(\delta)$, while for large δ we have $\delta/3 < J_2(\delta)$.

Both these bounds got a common improvement in the paper of Ben-Sasson, Goldberg, Kopparty and Saraf [BGKS20]. That paper showed that proximity gaps hold at radius $\gamma = J_{1.5}(\delta) - \eta$ (where $J_{1.5}(\delta) = 1 - (1 - \delta)^{1/3}$ is the "one-and-a-half Johnson radius"), with proximity loss ε^* and $a = O_\eta(1/\varepsilon^*) \cdot \gamma$.

[BGKS20] also gave a matching lower bound. They gave instances of infinitely many Reed Solomon codes with distance $\delta = 7/8$, such that when γ is equal to the one-and-a-half Johnson radius $J_{1.5}(\delta)$, a $o(1)$ proximity loss cannot be guaranteed even if a is taken as large as $\Omega(n)$.

Reed–Solomon codes with good distance Many of the cryptographic applications where the above mentioned results were used took $\mathcal{C}$ to be a Reed–Solomon code. It was thus of interest to obtain proximity gaps results for just Reed–Solomon codes.

In [BCI+20], Ben-Sasson, Carmon, Ishai, Kopparty and Saraf showed proximity gaps results for Reed–Solomon codes going beyond what was known for general linear codes. The crucial ingredient that enabled these results was the well developed theory of *decoding algorithms* for Reed–Solomon codes. Specifically, they used the Berlekamp–Welch algorithm [WB86] (for unique decoding) and the Guruswami–Sudan algorithm [GS99] (for list decoding up to the Johnson radius) as a tool (within the proof) to get a criterion for a given string to be close to the code, and used this to analyze proximity gaps. There were two kinds of proximity gaps that they showed:

- up to the unique decoding radius, $\gamma < \delta/2$, proximity gaps holds with $a = n$ and proximity loss $\varepsilon^* = 0$,
- up to the Johnson radius $\gamma < J(\delta)$, proximity gaps holds with $a = O_{\gamma,\delta}(n^2)$ and proximity loss $\varepsilon^* = 0$.

[1] Here $O_K(\cdot)$ means that the constant in the $O(\cdot)$ may depend on K.

Improving the radius up to which these proximity gap results holds is of is of great importance for the above mentioned practical hash-based arguments of knowledge.

Conjectured proximity gaps for Reed–Solomon codes A central guiding conjecture for the area has been that proximity gaps should hold all the way to radius δ with $o(1)$ proximity loss and a being polynomial in n. The conjecture is parametrized by a constant $\tau > 0$, the exponent of the polynomial bound.

Conjecture 1.2 (n^τ-bounded proximity gaps)**.** *Let $\delta \in (0,1)$ be a constant.*

For every Reed Solomon code $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ with length $|\mathcal{D}| = n$ and distance δ, and for every $\eta > 0$,

$$\mathcal{C} \text{ has proximity gaps up to radius } \gamma = \delta - \eta, \text{ with proximity loss } \varepsilon^* = o_\eta(1) \text{ and } a = O_\eta(n^\tau).$$

A stronger form of the conjecture, which is needed for some applications, asks for $\varepsilon^* = 0$.

If the above conjecture holds with a certain value of τ, then in applications of proximity gaps one would choose the field size q to be $\Omega_\eta(\frac{n^\tau}{\beta})$, where β is the target soundness error.

The negative result of [BGKS20] showed that the conjecture is false for $\tau < 1$. Diamond and Gruen [DG25b] showed that a strengthened form of the conjecture is false for $\tau < 1$ and certain $\delta = 1 - o(1)$. We also just learned of an elegant new result of Diamond and Gruen [DG25a] who disproved a stronger version of the above conjecture (made explicitly in [BCI+20]) for *all* τ: the stronger version asked for a to also be polynomially bounded in η.

To compare the conjecture with the positive results that are known, we list below how the strongest results in this direction fall short of this.

- For proximity gaps with $a = \mathsf{poly}(n)$, the previously best known results required γ to be bounded below $J(\delta)$, and achieved $a = O(n^2)$ and 0-proximity loss.
- For proximity gaps with $a = O(n)$, the previously best known results required γ to be bounded below $\max\{\delta/2, J_{1.5}(\delta)\}$ and achieved 0-proximity loss.

With this state of affairs, we now describe our results.

1.2 Overview of Results

Our paper is motivated by Conjecture 1.2. On the one hand, we prove positive results establishing proximity gaps for general Reed–Solomon codes with much stronger quantitative behavior than previously known. In particular, these results reduce the a parameter (and thus the needed field size for the same target soundness error) in known results by a factor n. On the other hand, we show negative results establishing the limits of proximity gaps: there are instances where any proximity gap result must have a proximity loss of $\Omega(1)$ unless the a parameter is chosen very large. In particular, we refute Conjecture 1.2 for all $\tau = O(1)$.

We give a brief summary of our main results:

- For $\gamma < \delta/2$, we show that proximity gaps up to radius γ holds with proximity loss ε^* and $a = \max\{\Omega(\frac{1}{\delta/2-\gamma}), 1 + \frac{\gamma}{\varepsilon^*}\}$. Notably, this bound on a is $O(1)$ if ε^* and $\delta/2 - \gamma$ are positive constants. The previous result in this range of γ, of [BCI+20], could not give any proximity gaps with this small an a: it needed $a = n$ to conclude anything in this range of γ (and then it gave proximity gaps with $\varepsilon^* = 0$).

  Our method also gives the optimal bound on a for the $\varepsilon^* = 0$ case. If $\delta/2 - \gamma$ is a positive constant, then proximity gaps up to radius γ holds with proximity loss $\varepsilon^* = 0$ with $a = \gamma n + 2$, matching the exact (and tight) bound on a that was proven by [AHIV17, RZ18] for $\gamma < \delta/3$.

- For $\gamma < J(\delta) - \eta$, we show that proximity gaps up to radius γ holds with 0 proximity loss and $a = O_{\eta,\delta}(n)$, The previous result in this range of γ, of [BCI+20], needed $a = \Omega_{\eta,\delta}(n^2)$. This significantly reduces the soundness error in the practical argument systems where these results are used.
- On the negative side, we show that for any constant τ, the n^τ-bounded proximity gaps conjecture (Conjecture 1.2) is false. Specifically for each integer τ, letting $\lambda_\tau = 2^{-(\tau+2)}$, for an infinite family of Reed–Solomon codes with distance $\delta = 1 - \lambda_\tau$, we show that proximity gaps at radius $\gamma = 1 - 4\lambda_\tau$ must have a proximity loss of at least $2\lambda_\tau$ if a is smaller than $n^{\tau - o(1)}$. The main qualitative feature is that for every constant τ, there are Reed-Solomon codes with some constant distance $\delta = \delta_\tau$, such that proximity gaps at radius $\delta - \Omega_\tau(1)$ and $a = O(n^\tau)$ must have a proximity loss of $\Omega_\tau(1)$.
- An instantiation of the above negative result for $\tau = 2$ gives us a tight negative result for proximity gaps at the Johnson radius: for an infinite family of Reed–Solomon codes with distance $\delta = 15/16$, we get that proximity gaps at radius $\gamma = J(\delta)$ must have a proximity loss of at least $\Omega(1)$ if a is smaller than $n^{2-o(1)}$. Thus a must jump from $O(n)$ to $\Omega(n^{2-o(1)})$ as γ increases past $J(\delta)$.
- Finally, we show that proximity gaps for a Reed–Solomon code at the list decoding radius (for list size q) and $o(1)$ proximity loss needs a at least $\frac{q}{2n}$ (and so the soundness error is at least $\frac{1}{2n}$, independent of q). This means that improving the proximity gap radius beyond the Johnson radius with $a < \frac{q}{2n}$ for any Reed–Solomon code implies that the list-decoding radius of that code is larger than the Johnson radius.

Additional results include: some weaker limitations to proximity gaps over prime fields, a sharp threshold behavior for proximity gaps at radius $\delta/3$ (in the $\delta = o(1)$ regime), and some attacks on STARKs for some simple constraint satisfaction problems, showing that the bound from [BGKS20] on their soundness error cannot be improved much.

The following tables show our new results on proximity gaps alongside what was known before.

Table 1: Proximity gaps for Reed–Solomon codes of distance δ. The shaded rows hold for arbitrary linear codes of distance δ.

Reference	requirement on		ε^*	**Notes**
	γ	a		
[RVW13]	arbitrary	$a \geq 2$	γ	
[BKS18]	arbitrary	$a \geq 2$	$\frac{a}{a-1} \cdot (\gamma - \gamma^2)$	$\varepsilon^* \to J^{-1}(\gamma) - \gamma$ as $a \to \infty$
Trivial	$< \delta$	$a \geq 2^{\Omega_\delta(n)}$	0	
[AHIV17]	$< \delta/4$	$a \geq 2$	$\frac{1}{a-1} \cdot \gamma$	$\varepsilon^* = 0$ for $a > \gamma n + 1$
[RZ18, BKS18]	$< \delta/3$	$a \geq 2$	$\frac{1}{a-1} \cdot \gamma$	$\varepsilon^* = 0$ for $a > \gamma n + 1$
[BCI+20]	$< \delta/2$	$a > n$	0	
This work	$< \delta/2$	$a \geq \frac{1}{\delta/2 - \gamma}$	$\frac{1}{a-1} \cdot \gamma$	$\varepsilon^* = 0$ for $a > \gamma n + 1$, $\frac{\delta}{2} - \gamma \geq \Omega(\frac{1}{\sqrt{n}})$
[BKS18]	$< 1 - (1-\delta)^{1/4} - \eta$	$a \geq \Theta_\delta\left(\frac{1}{\eta^2}\right)$	$O_\delta\left(\frac{1}{\eta^2 a}\right) \cdot \gamma$	$\varepsilon^* = 0$ for $a = \Omega_\delta\left(\frac{n}{\eta^2}\right)$
[BGKS20]	$< 1 - (1-\delta)^{1/3} - \eta$	$a \geq \Theta_\delta\left(\frac{1}{\eta}\right)$	$O_\delta\left(\frac{1}{\eta a}\right) \cdot \gamma$	$\varepsilon^* = 0$ for $a = \Omega_\delta\left(\frac{n}{\eta}\right)$
[BCI+20]	$< 1 - (1-\delta)^{1/2} - \eta$	$a \geq \Theta_\delta\left(\frac{n^2}{\eta^7}\right)$	0	
This work	$< 1 - (1-\delta)^{1/2} - \eta$	$a \geq \Theta_\delta\left(\frac{n}{\eta^5}\right)$	0	

Table 2: Parameters where large proximity losses are unavoidable, for Reed–Solomon codes of distance δ. LDR is the list-decoding radius, see Definition 1.8.

Reference	δ	γ	a	ε^*	**Notes**
[BGKS20]	$7/8$	$1-(1-\delta)^{1/3}=1/2$	$n-1=q-1$	$1/4$	q even
This work	$15/16$	$1-(1-\delta)^{1/2}=3/4$	$n^{2-o(1)}=(1-o(1))q$	$1/8$	q even
This work	$1-\lambda_\tau$	$1-4\lambda_\tau$	$n^{\tau-o(1)}=(1-o(1))q$	$2\lambda_\tau$	any τ, $\lambda_\tau=2^{-(\tau+2)}$, q even
This work	any δ	$\gamma=\mathsf{LDR}_{\mathbb{F}_q,n,q}(\delta)+\frac{2}{n}$	$\frac{q}{2n}$	$\delta-\gamma-\frac{1}{n}$	
This work	$\approx 1/2$	$\gamma=\delta-\frac{1}{\log_2 n}$	$n=q$	$\frac{1}{2\log_2 n}$	Mersenne prime q
[DG25b]	$1-1/n$	$\delta/2-\frac{1}{2n}$	n	$\frac{1}{n}$	
[DG25a]	$1-n^{-\epsilon}$	$\delta-n^{-\epsilon}$	$n^\tau=\frac{q}{n}$	$n^{-\epsilon}$	any τ, some $\epsilon>1/2$
This work	c/n	$\delta/3$	$n/c=(q-1)/c$	$\delta/3$	prime q, any $c=O(1)$
This work	n^ϵ/n	$\delta/2$	$n^{2/(1-4\epsilon)}=q-1$	$\delta/4$	prime q, any $\epsilon\in(0,1/4)$

1.3 Statements of Positive Results

We summarize our concrete improvements over [BCI+20].

1.3.1 Up to half minimum distance

Up to the unique decoding radius of the Reed–Solomon code, we obtain the following main result:

Theorem 1.3. *Let $\mathcal{C}$ be the code $\mathsf{RS}[\mathbb{F}_q,\mathcal{D},k]$ of block-length $n=|\mathcal{D}|$ and minimum distance $\delta=1-\frac{k}{n}$. Let $\gamma\in\left[\frac{\delta}{3},\frac{\delta}{2}-\frac{1}{n}\right]$. Suppose $u_0,u_1:\mathcal{D}\to\mathbb{F}_q$ are functions such that $S=\{z\in\mathbb{F}_q\mid\Delta(u_0+z\cdot u_1,\mathcal{C})\le\gamma\}$ is of size*

$$a\ge\left(\frac{\delta}{\gamma}-1\right)\cdot\frac{1}{\delta-2\gamma}$$

Then

$$\Delta([u_0,u_1],\mathcal{C}^2)\le\left(1+\frac{1}{a-1}\right)\cdot\gamma.$$

In other words, for distance loss ε^*, it suffices to take

$$a\ge\max\left(\left(\frac{\delta}{\gamma}-1\right)\cdot\frac{1}{\delta-2\gamma},\,1+\frac{\gamma}{\varepsilon^*}\right).$$

Using that $\frac{\delta}{\gamma}\le 3$, one obtains the coarser bound $\frac{1}{\delta/2-\gamma}$ on a, as cited in Table 1.

The most notable feature of this theorem is that for ε^* and $\delta/2-\gamma$ being positive constants, the a needed for the proximity gap result is just a constant. This saves a factor of $\Theta(n)$ in soundness error compared to the previous best result for this range, in [BCI+20], at the cost of introducing some arbitrarily small constant proximity loss.

Choosing $\varepsilon^*<\frac{1}{n}$ in Theorem 1.3 yields the following lossless result, for not too small distances.

Corollary 1.4. *For $\delta \geq \frac{3\cdot\sqrt{2}}{\sqrt{n}}$ and $\gamma \in \left[\frac{\delta}{3}, \frac{\delta}{2} - \frac{3}{\delta n}\right]$. If $S = \{z \in \mathbb{F}_q \mid \Delta(u_0 + z \cdot u_1, \mathcal{C}) \leq \gamma\}$ is of size*

$$a > \gamma \cdot n + 1,$$

then $\Delta([u_0, u_1], \mathcal{C}^2) \leq \gamma$.

For practical parameters, with $\delta = \Omega(1)$ and not too small word lengths n, the corollary improves [BCI+20, Theorem 4.1] from $a > n$ down to the [AHIV17, RZ18] bound $a > \gamma \cdot n + 1$. Moreover, since that bound is tight, the result is optimal on its range of validity.

1.3.2 Up to the Johnson radius

Up to the Johnson radius $J(\delta) = 1 - \sqrt{1-\delta}$, our main result is:

Theorem 1.5. *Let $\mathcal{C}$ be the code $\mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ with block-length $n = |\mathcal{D}|$ and minimum distance $\delta = 1 - \frac{k}{n}$. Denote $\rho = \frac{k}{n} = 1 - \delta$. For $\gamma \in \left(0, 1 - \sqrt{\rho}\right)$, let $\eta = 1 - \sqrt{\rho} - \gamma$, and $m = \max\left(\left\lceil \frac{\sqrt{\rho}}{2\eta} \right\rceil, 3\right)$. Suppose $u_0, u_1 : \mathcal{D} \to \mathbb{F}_q$ are functions such that $S = \{z \in \mathbb{F}_q \mid \Delta(u_0 + z \cdot u_1, \mathcal{C}) \leq \gamma\}$ is of size*

$$a > \frac{2(m+1/2)^5 + 3(m+1/2)\gamma\rho}{3\rho^{3/2}} \cdot n + \frac{m+1/2}{\sqrt{\rho}} = O_\rho\left(\frac{n}{\eta^5}\right). \tag{1}$$

Then

$$\Delta([u_0, u_1], \mathcal{C}^2) \leq \gamma.$$

Note that $\rho = 1 - \delta = \frac{k}{n}$ defined above is not the rate of the code, but off-by-$\frac{1}{n}$ from it. Also note that for $\eta \ll \sqrt{1-\delta}$, the asymptotics of the right-hand side of (1) is $O_\delta\left(\frac{n}{\eta^5}\right)$, as quoted in Table 2, with leading constant equal to $\frac{1}{48(1-\delta)^{3/2}}$. The theorem improves over [BCI+20, Theorem 5.1] by more than a factor n.

The newly achieved linear dependence on n is substantial for soundness proofs in certain real-world applications. In systems with parameters such that a soundness error of $O(n/q)$ yields reasonable security, but $O(n^2/q)$ error is too large, previously proven security could only be claimed for distances up to half the code's distance (via [BCI+20]), or the 1.5 Johnson bound (via [BGKS20]). The new result unlocks proven security with distances near the Johnson bound in such systems, possibly leading to significant improvements in performance.

1.3.3 Overview of techniques

The proofs of our positive results are based on the same approach as [BCI+20], which studies the Berlekamp–Welch and Guruswami–Sudan decoding algorithms in the rational function field $\mathbb{K} = \mathbb{F}_q(Z)$: Given that $u_0 + z \cdot u_1$ is γ-close to $\mathcal{C}$ for many $z \in \mathbb{F}_q$, the aim is to say something useful about the decodability of $w = u_0 + Zu_1$, to a word of the Reed–Solomon code extended by symbols in the rational function field $\mathbb{K} = \mathbb{F}_q(Z)$. By relating the decoding algorithm on w to the executions on $u_0 + z \cdot u_1$ for several substitutions $Z = z$, we show that it succeeds in finding a codeword $P(X) \in \mathbb{K}[X]$ that is γ-close to w. This $P(X)$ turns out to be of the form $v_0(X) + Zv_1(X)$, where $v_0, v_1 \in \mathcal{C}$, and this gives us the desired $[v_0, v_1] \in \mathcal{C}^2$ that is γ-close to $[u_0, u_1]$. Note that there is 0 proximity loss in this argument.

However, showing that the codeword $P(X)$ is of this specific form requires to keep track of the "algebraic complexity" of P's coefficients in the course of the algorithm, that is, their degree in Z. Our improvement comes from increasing the freedom in the first step of the algorithm, which solves a homogeneous linear system over $\mathbb{K}$ for a bivariate polynomial $Q(X, Y) \in \mathbb{K}[X, Y]$ that vanishes on the given word w.

To demonstrate the source of the improvement, we discuss a key idea in a simple situation. Suppose are given an $n \times (n+1)$ matrix A whose entries are all (at most) degree 1 polynomials in $\mathbb{F}_q[Z]$. We want to

find a vector $v \in (\mathbb{F}_q[Z])^{n+1}$ in the right kernel with all entries having low degree. How small can we take this degree? In general the kernel is only guaranteed to be at least 1 dimensional, and the formula for the vector in this kernel is given by Cramer's formula. Since it involves $n \times n$ determinants, we get that v can be taken to have all entries with degree at most n (and this happens to be the best upper bound on the degree that one can get).

Now suppose instead we are given an $n \times (n + \lambda n)$ matrix A whose entries are all degree 1 polynomials in Z. Again, we want to find a vector $v \in (\mathbb{F}_q[Z])^{n+\lambda n}$ in the kernel. How small can we take the degree? It turns out that now there is a vector v in the kernel with all entries having degree $O_\lambda(1)$! The proof is simple (dimension counting over $\mathbb{F}_q$). This is the phenomenon that we use – by making our linear systems have more slack than the bare minimum needed (which comes at a manageable cost to other parameters), we get solutions of much lower degree in Z. It can be viewed as a polynomial analogue of Siegel's Lemma from number theory, which is used in Diophantine approximation and transcendence.

Now we return to describing our proximity gap analysis.

1. In the case of Berlekamp-Welch decoding, we change the linear system from barely under-determined to significantly under-determined, by increasing the X-degree of $Q(X, Y) = A(X)Y + B(X)$. (That is, we increase the degree of what would have been the error-locator polynomial, $A(X)$.) However, a generously oversized error-locator polynomial yields a large distance loss for correlated agreement. We restore the distance afterwards, by applying a standard lemma for collinear proximates.

2. In the case of Guruswami–Sudan decoding, we discover that the linear system as set up in [BCI$^+$20] was already sufficiently under-determined. (This was a fluke: for the argument in [BCI$^+$20], it just needed to be barely under-determined.) Utilizing this slack, we find an interpolating polynomial $Q(X, Y) \in \mathbb{K}[X, Y]$ of Z-degree $O(1)$, compared to $O(n)$ as obtained therein. The smaller Z-degree directly translates to the improved bounds in the Guruswami–Sudan analysis of [BCI$^+$20].

1.4 Statements of Negative Results

1.4.1 The failure of n^τ-bounded proximity gaps

Our main negative result shows that the n^τ-bounded proximity gaps conjecture does not hold for any τ.

Theorem 1.6. *Let τ be a fixed positive integer, and $\lambda_\tau = 2^{-(\tau+2)}$. Let $\epsilon > 0$ be an arbitrary constant, and choose $\delta = 1 - \lambda_\tau$ and $\gamma = 1 - 4\lambda_\tau$.*

Then for all $\mathbb{F}_q$ of characteristic 2, there are Reed-Solomon codes $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, (1-\delta)n]$ over $\mathbb{F}_q$, domain $\mathcal{D}$ with $n = |\mathcal{D}| = O\left(q^{\frac{1}{\tau}(1+\epsilon)}\right)$, distance δ and words $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:

$$\left|\{z \in \mathbb{F}_q \mid \Delta(f + zg, \mathcal{C}) \leq \gamma\}\right| \geq (1 - o(1)) \cdot q \geq n^{\tau(1-\epsilon)},$$

for q large enough, and yet $\Delta([f, g], \mathcal{C}^2) \geq 1 - 2\lambda_\tau = \frac{2}{3}\delta + \frac{1}{3}\gamma$.

It seems plausible to us that analogues of this statement hold over any $\mathbb{F}_q$ with constant characteristic, but we have not checked this.

Versions of Conjecture 1.2 may still be true. For example, the conjecture may hold, even with $\tau = 1$, for fields of prime cardinality, or for well chosen evaluation domains $\mathcal{D}$ over fields of characteristic 2. We think this is a very exciting direction for future research.

The $\tau = 2$ case of the above theorem is particularly interesting: it shows that Theorem 1.5, on proximity gaps up to the Johnson radius, is tight for a particular value of δ. That theorem showed that proximity gaps hold for $\gamma < J(\delta)$ and with $a = O(n)$, with 0 proximity loss. What if we want to go to larger γ? The

following theorem shows that we cannot have proximity gaps at $\gamma = J(\delta)$, even with $\varepsilon^* = o(1)$, unless a is nearly quadratic, $\geq n^{2-\epsilon}$.

Corollary 1.7. *Let $\epsilon > 0$ be an arbitrary constant. Let $\delta = \frac{15}{16}$. Let $\gamma = 1 - \sqrt{1-\delta} = \frac{3}{4}$. Then for all $\mathbb{F}_q$ of characteristic 2, there are RS codes $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, (1-\delta)n]$ over $\mathbb{F}_q$, domain $\mathcal{D}$ with $n = |\mathcal{D}| = q^{\frac{1}{2}(1+\epsilon)}$, distance δ and functions $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:*

$$|\{z \in \mathbb{F}_q \mid \Delta(f + zg, \mathcal{C}) \leq \gamma\}| \geq (1 - o(1))q = n^{2(1-\epsilon)},$$

and yet $\Delta([f,g], \mathcal{C}^2) \geq \frac{7}{8} = \gamma + \frac{1}{8}$.

The proof of Theorem 1.6 is based on the $\mathbb{F}_2$-linear algebraic structure of $\mathbb{F}_q$. The evaluation domain $\mathcal{D}$ is essentially a uniformly random $\mathbb{F}_2$-subspace of $\mathbb{F}_q$ of size $\approx q^{(1+\epsilon)/\tau}$. The functions $f, g : \mathcal{D} \to \mathbb{F}_q$ will be the evaluation maps of well chosen monomials X^u, X^v (with $u > v > k$). Then $f + z \cdot g$ will be close to a polynomial $P_z(X)$ of degree at most k if the polynomial:

$$H(X) = X^u + zX^v - P_z(X)$$

has many roots in $\mathcal{D}$.

To create bad examples for proximity gaps using this framework, we need a special family of polynomials. Specifically, we need for many $z \in \mathbb{F}_q$ a polynomial $H_z(X)$ of the above form, with many roots in the small set $\mathcal{D}$. We find this special family through a combination of ideas from algebra and probability: using structural properties of the coefficients of *subspace polynomials*, and the second moment method to understand the distribution of these coefficients.

1.4.2 Limitations on proximity gaps at the list decoding radius

Our next result shows a strong connection with the extremely well studied question of list decoding Reed–Solomon codes. To talk concretely about these, we define:

Definition 1.8 ($\mathsf{LDR}_{\mathbb{F}_q,\mathcal{D},L}(\delta)$)**.** *Let $\mathcal{C}$ be the Reed–Solomon code $\mathsf{RS}[\mathbb{F}_q, \mathcal{D}, (1-\delta)|\mathcal{D}|]$. We define $\mathsf{LDR}_{\mathbb{F}_q,\mathcal{D},L}(\delta)$ to be the largest γ such that for all functions $c : \mathcal{D} \to \mathbb{F}_q$, we have:*

$$|\{v \in \mathcal{C} \mid \Delta(c, v) \leq \gamma\}| \leq L.$$

Further, for $n \leq q$, we define $\mathsf{LDR}_{\mathbb{F}_q,n,L}(\delta)$ to be the minimum, over all choices of $\mathcal{D} \subseteq \mathbb{F}_q$ with $|\mathcal{D}| = n$, of $\mathsf{LDR}_{\mathbb{F}_q,\mathcal{D},L}(\delta)$.

The Johnson bound applied to Reed–Solomon codes implies that $\mathsf{LDR}_{\mathbb{F}_q,\mathcal{D},n}(\delta) \geq J(\delta) = 1 - \sqrt{1-\delta}$. It is a well-known and easy fact that for $L \leq 2^{\min(\delta n, n - \delta n)}$, we have $\mathsf{LDR}_{\mathbb{F}_q,\mathcal{D},L}(\delta) \leq \delta$.

Determining $\mathsf{LDR}_{\mathbb{F}_q,n,\mathsf{poly}(n)}(\delta)$ is a very basic and well-studied problem in coding theory. In one direction, it is known [JH01, BSKR06] that some Reed–Solomon codes require large polynomial list size at radii strictly between $J(\delta)$ and δ in the $\delta = \Theta(1)$ regime. In the other direction, an exciting line of recent works [RW13, ST20, GLS+24, BGM23, GZ23, AGL24] showed that Reed–Solomon codes with randomly chosen evaluation domains are list-decodable almost all the way to radius δ with constant list size.

We show that the proximity gaps phenomenon with $o(1)$ proximity loss starts requiring large a once γ goes beyond the list-decoding radius for list size q:

Theorem 1.9. *Let $\mathcal{C}$ be the Reed–Solomon code $\mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, with $|\mathcal{D}| = n$ and $k = (1-\delta)n$. Let $\gamma = \mathsf{LDR}_{\mathbb{F}_q,\mathcal{D},q}(\delta) + \frac{2}{n}$. Then there exist functions $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:*

$$|\{z \in \mathbb{F}_q \mid \Delta(f + z \cdot g, \mathcal{C}) \leq \gamma\}| \geq \frac{q}{2n},$$

but with $\Delta([f,g], \mathcal{C}^2) \geq \delta - \frac{1}{n}$.

The value of a here means that the soundness error $\frac{a}{q}$ is at least $\frac{1}{2n}$, and thus cannot be reduced by increasing q.

The proof of the above theorem is quite roundabout: we stumbled across it while studying the soundness of STARKs and the so called "ETHSTARK toy problem" [Sta23]. It involves a clever choice of f and g based on the center of any bad list-decoding configuration and a well chosen rational function of the form $\frac{1}{X-\alpha}$. In particular, we do not know how to show any analogous theorem for general linear codes of distance δ.

Along the way, we prove a structural property of any bad list-decoding configuration, Lemma 6.1, which is also useful in our results on the limitations on the soundness of STARKs, discussed in 1.4.5.

1.4.3 Limitations on proximity gaps at radius near δ over prime fields?

Our next result gives mild limitations of proximity gaps over prime fields, which are of particular interest for practice. They are based on multiplicative subgroups, and they creates instances of proximity gaps with mild proximity loss. However, the existence of infinitely many of these instances is conjectural. It depends on a very clean and basic conjecture in additive number theory, which we now state.

Definition 1.10. *For $E \subseteq \mathbb{F}_q$, we define:*

$$E^{(+\ell)} = \{e_1 + \ldots + e_\ell \mid e_1, \ldots, e_\ell \in E \text{ and distinct}\}.$$

Definition 1.11. *For a prime power q and integers $a, b \leq q$, we say (q, a, b) is admissible if there exists a multiplicative subgroup $G \subseteq \mathbb{F}_q^*$ with:*

- $|G| = b$,
- *For $\ell = \lfloor \frac{b}{2} \rfloor$, we have that $|G^{(+\ell)}| \geq a$.*

Conjecture 1.12. *For infinitely many primes q, there exists $b \leq 10 \log q$ such that $(q, q/10, b)$ is admissible.*

Heuristically, if $|G| = b$, there are $\binom{b}{\ell}$ sums in $G^{(+\ell)}$, and if they are pseudorandomly distributed, we may very optimistically expect $|G^{(+\ell)}| \geq \Omega(\min(q, \binom{b}{\ell}))$. The above conjecture expresses a milder form of this heuristic.

If q is a Mersenne prime $2^p - 1$, we can take G to be the subgroup generated by -2: $G = \{\pm 1, \pm 2, \ldots, \pm 2^{p-1}\}$. Using binary expansions, every element of $\mathbb{F}_q$ turns out to be representable as a sum of exactly p elements of G (see Remark 7.3): this means that $(q, q, 2\log_2(q+1))$ is admissible. Thus Conjecture 1.12 is weaker than the well-known conjecture asserting the infinitude of Mersenne primes.

Sumsets of multiplicative groups in $\mathbb{F}_q$ have been extensively studied in additive combinatorics, especially recently in the context of the sum-product phenomeonon. However, existing results fall far short of the conjecture. The strongest unconditional result in this direction, of Glibichuk and Konyagin [GK07], implies that the ℓ-wise sumset of G (which includes nondistinct sums) has size at least $|G|^{\Omega(\log \ell)}$, which is far short of what the conjecture asks for.

The conjecture then gives infinitely many nontrivially bad instances for proximity gaps over prime fields, via the following theorem.

Theorem 1.13. *Suppose (q, a, b) is admissible with b even. Let G be the corresponding subgroup of $\mathbb{F}_q^*$ of cardinality b, and let $H \subseteq \mathbb{F}_q^*$ be any multiplicative subgroup containing G. Consider the Reed–Solomon code $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, with $\mathcal{D} = H$, $n = |\mathcal{D}|$, $k = \left(\frac{1}{2} - \frac{2}{b}\right)n$, and relative distance $\delta = \frac{1}{2} + \frac{2}{b}$.*

Then there exist functions $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:

$$\left| \left\{ z \in \mathbb{F}_q \mid \Delta(f + z \cdot g, \mathcal{C}) \leq \delta - \frac{2}{b} \right\} \right| \geq a,$$

but $\Delta([f, g], \mathcal{C}^2) \geq \delta - \frac{1}{b}$.

Note that n, the size of the evaluation domain, might be chosen as small as b or as large as $q-1$ (it just needs to satisfy $b \mid n$ and $n \mid (q-1)$).

With the latter choice, the conjecture implies that for Reed–Solomon codes over prime fields, for infinitely many n, when $\gamma = \delta - \Theta\left(\frac{1}{\log n}\right)$ and $a \geq q/10 \geq n/10$, we must have a proximity loss of $\Theta\left(\frac{1}{\log n}\right)$.

When instantiated with q being a Mersenne prime, the constants in the Θs above are explicit and small. So for $q = M_{31}$ (the Mersenne prime $2^{31}-1$) and $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathbb{F}_q^*, \frac{q-1}{2} + \frac{2}{\log_2 q}]$, this gives us functions $f, g : \mathbb{F}_q^* \to \mathbb{F}_q$ such that for all values of $z \in \mathbb{F}_q$ we have

$$\Delta(f + z \cdot g, \mathcal{C}) \leq \frac{1}{2},$$

and yet

$$\Delta([f,g], \mathcal{C}^2) \geq \frac{1}{2} + \frac{1}{62} \approx 0.516.$$

Another instantiation, closer to practice, is when $q = (M_{31})^4 \approx 2^{124}$. It turns out that (q, q, b) is admissible, where $b = 8 \cdot 31 \approx 2\log_2 q$, and we get an example of a Reed–Solomon code of any length n satisfying b and $n-1$, distance $\delta = \frac{1}{2} + \frac{2}{b} \approx 0.508$, such that proximity gaps at radius $\gamma = \delta - \frac{2}{b} = \frac{1}{2} \approx \delta - 0.008$ have proximity loss $\varepsilon^* = \frac{1}{b} \approx 0.004$ with $a = q$.

For other specific q that are used in practical applications, it is sometimes possible to experimentally check admissibility of a tuple (q, a, b) either by brute force or estimating collision probability of sums of random ℓ-tuples from a subgroup G of size b.

1.4.4 Some proximity gap phase transitions when $\delta = o(1)$

We now give some instances showing that for Reed–Solomon codes with vanishingly small relative distance, there are some sharp transitions in the behavior of proximity gaps. We view them as explaining why the simpler proof techniques from [AHIV17, RZ18, BKS18] for the $\gamma < \delta/3$ do not extend to $\gamma \geq \delta/3$, and that things will get tougher the closer we get to δ.

At $\delta/3$

Recall the basic proximity gap phenomenon at distance strictly less than $\delta/3$.

Theorem 1.14 ([AHIV17], [RZ18],[BKS18])**.** *Let $\mathcal{D} \subseteq \mathbb{F}_q$ with $|\mathcal{D}| = n$. Let $\epsilon \in [0,1]$. Let $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, and let $f, g : \mathcal{D} \to \mathbb{F}_q$. Let $\delta = 1 - k/n$ be the distance of $\mathcal{C}$. Let $\gamma < \delta/3$.*

Suppose:

$$|\{z \in \mathbb{F}_q \mid \Delta(f + zg, \mathcal{C}) \leq \gamma\}| \geq a.$$

Then $\Delta([f,g], \mathcal{C}^2) \leq \frac{a}{a-1} \cdot \gamma$.

In particular, for $a > \gamma n + 1$, this theorem gives us $\Delta([f,g], \mathcal{C}^2) \leq \gamma$, that is correlated agreement with 0 proximity loss. The following theorem shows that there are Reed-Solomon codes where the behavior exactly at $\gamma = \delta/3$ changes drastically: even when $a \gg \gamma n$, there is a substantial proximity loss.

Theorem 1.15. *Let $c > 0$ be an integer. There exist infinitely many q, RS codes $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ over $\mathbb{F}_q$, domain $\mathcal{D} = \mathbb{F}_q$ with $n = |\mathcal{D}| = q$, $k = n - c$ (so that the relative distance δ of $\mathcal{C}$ equals $\frac{c}{n}$), and functions $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:*

$$|\{z \in \mathbb{F}_q \mid \Delta(f + zg, \mathcal{C}) \leq \delta/3\}| \geq \frac{q-1}{c} = \frac{n}{c} = O\left(\frac{1}{\delta}\right),$$

and yet $\Delta([f,g], \mathcal{C}) \geq 2\delta/3$.

How is this consistent with our Theorem 1.3, which shows proximity gaps almost all the way up to $\delta/2$? Theorem 1.3 requires a at least as large as $\Omega(\frac{1}{\delta-2\gamma})$. Thus when $\delta = O(1/n)$ and $\gamma = \delta/3$ as in the above theorem, Theorem 1.3 needs $a \geq \Omega(\frac{1}{\delta}) = \Omega(n)$, which is consistent with Theorem 1.15.

The above example shows that there truly is a behavioral transition at $\delta/3$: the unrestricted proximity gap phenomonon that holds for $\gamma < \delta/3$ (which did not care about whether $\delta = o(1)$), does not hold in this form beyond $\delta/3$.

At $\delta/2$

We also see an interesting threshold occurring at radius $\gamma = \delta/2$. The following theorem shows that there are Reed-Solomon codes with prime field size q nearly quadratic in the block length n (but with vanishingly small relative distance $\delta = O(1/n^{1-\kappa})$), where proximity gaps at radius $\gamma = \delta/2$ with $\varepsilon^* = o(\gamma)$ needs a being nearly quadratic in n.

Theorem 1.16. *Let $0 < \epsilon < 1/4$ be fixed. There are infinitely many primes q, RS codes $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ over $\mathbb{F}_q$, domain $\mathcal{D}$ with $n = |\mathcal{D}| = O(q^{0.5+2\epsilon})$, $k = n - \Theta(n^\epsilon)$ (so that the relative distance is $\delta = \Theta(n^{-(1-\epsilon)})$), and functions $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:*

$$|\{z \in \mathbb{F}_q \mid \Delta(f + zg, \mathcal{C}) \leq \delta/2\}| \geq q - 1 = \Omega(n^{2/(1+4\epsilon)}) = \Omega(n^{2-8\epsilon}),$$

and yet $\Delta([f, g], \mathcal{C}^2) \geq 3\delta/4$.

We remark that both the negative results of this section in the $\delta = o(1)$ regime, about thresholds at $\delta/3$ and $\delta/2$, *cannot* hold for δ being $\Omega(1)$. This is because we do know positive proximity gaps results at radius $J_{1.5}(\delta) - \eta$ and $J(\delta) - \eta$ with a equal to $O_\eta(1)$ and $O_\eta(n)$ respectively, and when $\delta = \Omega(1)$ we have that $\delta/3 < J_{1.5}(\delta) - \Omega_\delta(1)$ and $\delta/2 < J(\delta) - \Omega_\delta(1)$.

1.4.5 Limits to the soundness of STARKs

Our final contribution is devoted to the soundness of STARKs.

The main theorem about STARKs [BSBHR18, BGKS20] shows how to interactively prove (in the IOPP model) the satisfiability of a constraint satisfaction problem presented in *Algebraic Intermediate Representation (AIR)*. The soundess of the resulting interactive protocol is a function of the complexity of the AIR.

There are two parts to this protocol.

1. The first part reduces the the satisfiablilty of the AIR to a statement of proximity to a Reed–Solomon code. Specifically, given an AIR A which the prover claims is satisfiable, the prover and verifier first interact and if the verifier does not reject, will end up with the prover writing down some functions $h_1, \ldots, h_c : \mathcal{D} \to \mathbb{F}_q$, where $\mathcal{D} \subseteq \mathbb{F}_q$ with $|\mathcal{D}| = n$. The prover now has to prove that $h_1, \ldots, h_c$ has high correlated agreement with the code $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$.
2. The second part now checks this claim of the prover – this is done using proximity gaps, to combine the c functions into one, and an IOPP protocol for Reed–Solomon codes, such as FRI, STIR or WHIR (which typically also rely on some form of proximity gaps).

In practical STARKs, this protocol is typically instantiated with the Reed–Solomon code $\mathcal{C}$ having distance δ some large constant like 0.75. Furthermore, q is typically very very large (eg. 2^{120}, for cryptographic security), and n is moderate (eg. 2^{20}). This will help interpret our results below: we view $\frac{1}{q}$ as negligible probability and $\frac{1}{n}$ as noticeable probability.

The first part of the protocol is based on a number of ideas: low-degree extension of the satisfying assignment, arithmetization of the constraints, checking that a polynomial vanishes on the set S using the vanishing polynomial of S, and the DEEP quotienting method. This protocol has the following completeness and soundness properties, expressed in terms of the list-decodability of $\mathcal{C}$:

- If the AIR is satisfiable, then there is a prover strategy for the protocol that will end with the prover writing down $h_1, \ldots, h_c$, such that $(h_1, \ldots, h_c) \in \mathcal{C}^c$.
- If the AIR is not satisfiable, then for any prover strategy, the probability that it will end with the prover writing down some $(h_1, \ldots, h_c)$ with

$$\Pr\left[\Delta([h_1, \ldots, h_c], \mathcal{C}^c) \leq \gamma_L\right] \leq O\left(\frac{L^2 \cdot k}{q}\right),$$

  where $\gamma_L = \mathsf{LDR}_{\mathbb{F}_q, \mathcal{D}, L}(\delta)$.

How are the parameters set? Note that for δ being a large constant near 1, and L being a large integer which is $O(1)$, then γ_L is also a large constant near 1, by the Johnson bound. In this case, the above "cheating probability" is bounded by $O(n/q)$, and this is made smaller than the tolerable soundness error by making q large enough.

We give an attack showing that the soundness analysis of this reduction cannot be improved much: it is truly governed by the list-decodability of $\mathcal{C}$ (for list size q).

Specifically, we give a simple unsatisfiable constraint satisfaction problem (specified by a simple AIR) and a prover strategy that makes the first part of the protocol produce $(h_1, \ldots, h_c)$ such that:

$$\Pr\left[\Delta([h_1, \ldots, h_c], \mathcal{C}^c) \leq \frac{1 + \gamma_q}{2}\right] \geq \Omega(1/n),$$

where $\gamma_q = \mathsf{LDR}_{\mathbb{F}_q, \mathcal{D}, q}(\delta)$. In more detail, given a "bad list decoding center" for $\mathcal{C}$, we find a prover strategy that uses it to get a probability $\Omega(1/n)$ of finding $(h_1, \ldots, h_c)$ that are much closer than trivial to $\mathcal{C}^c$.

The constraint satisfaction problem itself is very easy to describe:

CYCLE-SUM:

- Let a divide $q - 1$. Let $g \in \mathbb{F}_q^*$ have order a. Let $G = \{1, g, g^2, \ldots, g^{a-1}\}$ be the subgroup generated by g.
- We want a function $f : G \to \mathbb{F}_q$ such that for each $x \in G$,

$$f(gx) = f(x) + 1.$$

Observe that no such f exists. If there was such a function, we would have

$$f(1) = f(g^a) = f(g^{a-1}) + 1 = f(g^{a-2}) + 2 = \ldots = f(1) + a.$$

But a is relatively prime to q, and is thus nonzero in $\mathbb{F}_q$.

The prover strategy we give is based on bad list-decoding configurations for Reed–Solomon codes. We prove some structural statements about any such configuration. The structural statements enable us to handle the DEEP-quotienting challenge from the verifier – they tell us which nearby Reed–Solomon codeword to answer according to.

Theorem 1.17. *Consider the IOP protocol for the CYCLE-SUM constraint satisfaction problem given by the DEEP-ALI reduction, using the Reed–Solomon code $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, where:*

- $\mathcal{D}$ *is a union of* t *cosets of* G
- $|\mathcal{D}| = a \cdot t = n$ *and* $k = a = \frac{1}{t}n$.
- $\delta = 1 - \frac{1}{t}$ *is the distance of* $\mathcal{C}$.

Then there is a prover strategy that does not make the verifier reject, and produces $(h_1, \ldots, h_c)$ *such that:*

$$\Pr\left[\Delta([h_1, \ldots, h_c], \mathcal{C}^c) \leq \frac{1+\gamma_q}{2}\right] \geq \Omega(1/n),$$

where $\gamma_q = \mathsf{LDR}_{\mathbb{F}_q, \mathcal{D}, q}(\delta) + \frac{1}{n}$.

To understand the implication, it is good to think about settings where γ_q equals the Johnson radius $J(\delta) = 1 - \sqrt{1-\delta}$. (We know that this occurs for some Reed–Solomon codes [JH01, BSKR06], and does not happen for some Reed–Solomon codes [ST20, GLS$^+$24, BGM23, GZ23, AGL24]).

In this case, with noticeable probability the prover is producing $h_1, \ldots, h_c$ such that $[h_1, \ldots, h_c]$ is $\left(1 - \frac{1}{2}\sqrt{1-\delta}\right)$-close to $\mathcal{C}^c$. When δ is large enough (> 0.75), this proximity is a smaller than δ by a constant, and translates into noticeably larger success probability for a cheating prover in the final STARK proof. (For example, if $\delta = 0.84$ and $\gamma_q = J(\delta)$, then we have $(1 + \gamma_q)/2 = 0.8$).

1.5 Other Related Work

As mentioned earlier, the STIR [ACFY24] and WHIR [ACFY25] protocols improved the state of the art for theoretical and practial IOPP protocols for Reed–Solomon codes. The latter work introduced a key new notion, mutually correlated agreement, which is a strengthening of the property of having proximity gaps with 0 proximity loss.

Recently, Haböck [Hab25] showed that Reed–Solomon codes have the mutually correlated agreement property up to the Johnson radius, and gave a general method for establishing this property in any code whenever there is an underlying collinearity version of proximity gaps.

The results of [Zei24] and [GKL24] showed versions of the proximity gaps results of [BKS18] and [BGKS20] in the stronger mutual correlated agreement form.

The work of [GCXK25] gave a black box reduction showing that every code of distance δ has proximity gaps up to radius $J(\mathsf{LDR}(\delta))$, and observed that this gave interesting new proximity gaps results for some codes where we know very good bounds on the list decoding radius. Furthermore, they also achieved the stronger mutually correlated agreement form[2]. The results of [GCXK25] gave new proximity gaps for random Reed–Solomon codes up to the Johnson radius with $a = O(n)$. Our results now show the same for all Reed–Solomon codes.

Our Theorem 1.9 which gives a connection in the other direction (good proximity gaps results imply good list-decodability – but only for Reed–Solomon codes) arose while we were studying the ETHSTARK [Sta23] toy problem, and our methods can be used to give an attack on that problem. A different attack on the ETHSTARK toy problem, in a different setting of parameters, was given by Garreta, Gruen, and Manzur [GGM25] in upcoming work, using a more direct strategy studying the rank of an associated matrix.

Organization of this paper

The first part of the paper focuses on the positive results. In Section 2 we elaborate the aforementioned improved Berlekamp–Welch analysis and prove Theorem 1.3 as well as its lossless variant Corollary 1.4.

[2]We remark the the weaker form without mutually correlated agreement can also be proved by inspecting the double-Johnson radius proximity gap result of [BKS18], and noticing that one of the "Johnsons" could be replaced with the list-decoding radius.

Section 3 shows Theorem 1.5, using more careful analysis of the Guruswami–Sudan interpolant polynomial, as well as improved bookkeeping in other parts of the original proof from [BCI+20]. In Section 4 we then discuss how our improvements translate to more general proximity gaps statements, which are used in practice. We treat linear combinations of several words, the so-called weighted correlated agreement theorem, and we outline how our improvements impact the mutual correlated agreement theorem proven in [Hab25].

The rest of the paper is devoted to the negative results. In Section 5 we prove Theorem 1.6 and provide the aforementioned examples which falsify n^τ-bounded proximity gaps, Conjecture 1.2, for fields of characteristic 2. In Section 6 we prove Theorem 1.9 about the failure of proximity gaps beyond the list decoding radius of the code. Section 7 investigates the limits of Conjecture 1.2 in prime fields. We prove Theorem 1.13 which is the underlying tool for our counter examples over Mersenne primes, as well as the $\delta = o(1)$ phase transition Theorems 1.15 (at $\delta/3$) and 1.16 (at $\delta/2$). Finally, in Section 8 the attack on the elementary STARK for CYCLE-SUM is discussed and Theorem 1.17 is proven.

Concurrent Work

We just learnt about two concurrent and independent works that also address proximity gaps for Reed-Solomon codes. In [CS25], Crites and Stewart disproved the stronger form of Conjecture 1.2 written in [BCI+20] asking that a have a polynomial dependence on η (this result was also independently discovered by [DG25a] already mentioned above). [CS25] also independently discovered the reduction between list-decoding and proximity gaps (our Theorem 1.9). In [GG25], Goyal and Guruswami showed that random Reed-Solomon codes do have proximity gaps all the way to radius $\delta - \eta$, with a being $O_\eta(n)$. Random Reed-Solomon codes are the only family of Reed-Solomon codes known to have list decoding radius as large as $\delta - \eta$ for any $\eta > 0$, and because of Theorem 1.9, we would have suggested them as a natural candidate for improved proximity gaps.

Acknowledgements

We thank David Levit for valuable discussions and encouragement. We also thank Antonio Sanso for valuable discussions that motivated the improvements in Section 3.2.

2 Improved proximity gaps up to half the minimum distance

In this section we show the following improvement of [BCI+20, Theorem 4.1], together with its lossless variant, Corollary 1.4.

Theorem 1.3. *Let $\mathcal{C}$ be the code $\mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ of block-length $n = |\mathcal{D}|$ and minimum distance $\delta = 1 - \frac{k}{n}$. Let $\gamma \in \left[\frac{\delta}{3}, \frac{\delta}{2} - \frac{1}{n}\right]$. Suppose $u_0, u_1 : \mathcal{D} \to \mathbb{F}_q$ are functions such that $S = \{z \in \mathbb{F}_q \mid \Delta(u_0 + z \cdot u_1, \mathcal{C}) \leq \gamma\}$ is of size*

$$a \geq \left(\frac{\delta}{\gamma} - 1\right) \cdot \frac{1}{\delta - 2\gamma}$$

Then

$$\Delta([u_0, u_1], \mathcal{C}^2) \leq \left(1 + \frac{1}{a-1}\right) \cdot \gamma.$$

In other words, for distance loss ε^*, it suffices to take

$$a \geq \max\left(\left(\frac{\delta}{\gamma} - 1\right) \cdot \frac{1}{\delta - 2\gamma},\ 1 + \frac{\gamma}{\varepsilon^*}\right).$$

Let us give a quick outline of the proof, and hence the section. The proof goes along the lines of [BCI+20], which studies the Berlekamp–Welch decoder in the rational function field $\mathbb{K} = \mathbb{F}_q(Z)$ on the "received" word

$$u_0(x) + Z \cdot u_1(x).$$

In the first step, we search for polynomials $A(X, Z), B(X, Z) \in \mathbb{F}_q[X, Z]$ satisfying the interpolation constraints

$$A(x, Z) \cdot (u_0(x) + Z \cdot u_1(x)) = B(x, Z), \quad x \in \mathcal{D},$$

a homogeneous linear system with n equations, and the coefficients of the polynomials as unknowns. See Lemma 2.1 below. The crucial difference to [BCI+20] is that we consider the decoder in an unusual setting. Given $e = \lfloor \gamma \cdot n \rfloor$ we drastically oversize the error-locator polynomial,

$$\deg_X A = e + h,$$

with $h > 0$ chosen maximal so that divisibility is still assured, i.e. $\deg_X A = n - e - 1$. In this setting, the distance of the decoded word is only assured $< \delta - \gamma$ and typically beyond the unique decoding radius. In return, it allows us to find a solution of significantly lower Z-degree than in [BCI+20], see Lemma 2.1.

As a consequence of the smaller Z-degree, the divisibility step, which proves that

$$P(X, Z) = B(X, Z)/A(X, Z)$$

is again a polynomial with $\deg_X P \le k$ and $\deg_Z P \le 1$, goes through with a smaller bound for $a = |S|$. As in [BCI+20], this step uses the Polishchuk–Spielman lemma, concluding the desired divisibility from that

$$B(X, z)/A(X, z) = P_z(X),$$

for every $z \in S$, the promised proximate polynomial of degree at most k. A sensible choice of h achieves the bound on $|S|$ as stated in Theorem 1.3. We refer to Section 2.2 for the elaboration of this step.

Divisibilit the oversized A polynomial implies correlated agreement with a potientially huge distance loss. Nevertheless, the loss is bounded enough to show *collinearity*: With only few exceptions (where $P(X, z) = 0$) the γ-proximate polynomial $P_z(X)$ lies on the line spanned by the polynomials $P(0, X)$ and $P(1, X)$, that is

$$P_z(X) = P(X, z) = P(0, X) + z \cdot P(1, X).$$

It turns out that collinearity extends to all $z \in S$, and we may apply a well-known argument, Lemma 2.4, that proves the claimed distance of $[u_0, u_1]$ to the interleaved code $\mathcal{C}^2$ from collinearity, see Section 2.3.

2.1 Non-standard Berlekamp–Welch interpolant over $\mathbb{K}$

Let $u_0, u_1 \in \mathbb{F}_q^{\mathcal{D}}$, and let $e < \frac{n-k}{2}$ be an integer. For integer $h \ge 0$, which we call slack, we search for polynomials

$$A(X, Z) = \sum_{i=0}^{e+h} a_i(Z) \cdot X^i, \quad B(X, Z) = \sum_{j=0}^{k+e+h} b_j(Z) \cdot X^j, \tag{2}$$

with all $a_i(Z), b_j(Z) \in \mathbb{F}_q[Z]$, satisfying the interpolation constraints

$$A(x, Z) \cdot (u_0(x) + Z \cdot u_1(x)) = B(x, Z), \quad x \in \mathcal{D} \tag{3}$$

as polynomials, or in other words as elements of $\mathbb{K} = \mathbb{F}_q(Z)$. This is a homogeneous linear system of n equations with the polynomials $a_i(Z)$, $b_j(Z)$ as unknowns. We consider the decoder with very large slack h, so that

$$k + e + h = \deg_X(B) = n - e - 1, \tag{4}$$

or $h = n - k - 2e - 1$. This is always non-negative, even in the edge case $e = \frac{n-k-1}{2}$, where $h = 0$. In the following, we abbreviate statements such as "for bivariate polynomial $P(X, Z)$, its individual degrees are bounded as $\deg_X P \le d_X$ and $\deg_Z P \le d_Z$" by using the vector notation $\deg(P) \le (d_X, d_Z)$.

Lemma 2.1. *For any integer $e \in \left[0, \frac{n-k}{2}\right)$, let $D_X = k + e + h$, with $h = n - k - 2e - 1$, and*

$$D_Z \geq \left\lceil \frac{e+1}{h+1} \right\rceil.$$

Then there are $A(X,Z)$, $B(X,Z)$ with $\deg(A) \leq (D_X - k, D_Z - 1)$, $\deg(B) \leq (D_X, D_Z)$, A non-zero, and so that the equations (3) holds at every $x \in \mathcal{D}$, as an identity in $\mathbb{F}_q[Z]$.

Proof. For any integer $D_Z \geq 1$, roll out (3) as a homogeneous $\mathbb{F}_q$-linear system in the coefficients of the polynomials $a_i(Z), b_j(Z)$ with degree bounds $\deg a_i(Z) \leq D_Z - 1$, $\deg b_j(Z) \leq D_Z$. The number of unknowns is

$$(k + e + h + 1) \cdot (D_Z + 1) + (e + h + 1) \cdot D_Z = n \cdot (D_Z + 1) + (h + 1) \cdot D_Z - e,$$

whereas the number of equations is $n \cdot (D_Z + 1)$, one for each coefficient and each $x \in \mathcal{D}$. Therefore, if $(h + 1) \cdot D_Z - e \geq 1$, then the linear system has a non-trivial solution, regardless of the size of S. For such a non-trivial solution, $A(X,Z)$ must be a non-zero polynomial: Triviality of $A(x,Z)$ implies triviality of $B(x,Z)$ at more points than its degree in X hence if $A = 0$ then also $B = 0$, contradicting non-triviality. □

2.2 Dividing $B(X,Z)$ by $A(X,Z)$

Given two polynomials $A(X,Z)$, $B(X,Z)$ over an arbitrary field F, the Polishchuk–Spielman Lemma argues divisibility in $F[X,Z]$ from divisibility of sufficiently many univariate restrictions. Its original proof in [PS94] has a subtle gap noted and fixed by Ronald Cramer and Jade Nardi, and a fixed version appears in [BCI+20, Lemma 4.3].

Lemma 2.2 (Polishchuk–Spielman [PS94], see statement in [BCI+20])**.** *Let $A(X,Z), B(X,Z) \in F[X,Z]$ be polynomials of degrees $\deg(A) \leq (a_X, a_Z)$ and $\deg(B) \leq (b_X, b_Z)$, and suppose that $B(X,z)/A(X,z) \in F[X]^{\leq b_X - a_X}$ for at least $n_Z > 0$ different values of z, as well as $B(x,Z)/A(x,Z) \in F[Z]^{\leq b_Z - a_Z}$ for at least $n_X > 0$ different values of x, where*

$$\frac{b_X}{n_X} + \frac{b_Z}{n_Z} < 1. \tag{5}$$

Then $P(X,Z) = B(X,Z)/A(X,Z) \in F[X,Y]$ and its degrees satisfy $\deg(P) \leq (b_X - a_X, b_Z - a_Z)$.

Let us demonstrate how to use Lemma 2.1 in combination with Lemma 2.2 to obtain the following intermediate claim.

Claim 2.3. *For $0 < \gamma < \frac{\delta}{2}$. If $S = \{x \in \mathbb{F}_q \,:\, \Delta(u_0 + z \cdot u_1, \mathcal{C}) \leq \gamma\}$ is of size*

$$a \geq \frac{1}{\gamma} \cdot \frac{\delta - \gamma}{\delta - 2\gamma}$$

then there exist polynomials $p_0(X), p_1(X) \in \mathbb{F}_q[X]$ of degree at most k so that $\Delta([u_0, u_1], [p_0, p_1]) < \delta - \gamma$.

Proof. Let $e = \gamma n$, $h = n - k - 2e - 1$, and choose $D_X = k + e + h$ and $D_Z = \left\lceil \frac{e+1}{h+1} \right\rceil$. Lemma 2.1 states the existence of nontrivial polynomials $A(X,Z), B(X,Z)$ with $\deg(A) \leq (D_X - k, D_Z - 1)$ and $\deg(B) \leq (D_X, D_Z)$, which satisfy the Berlekamp–Welch equations (3).

We claim that this solution satisfies the conditions of Lemma 2.2. First, for every $z \in S$ the promised proximate $P_z(X) \in \mathbb{F}_q[X]$ of degree at most k agrees with $u_0 + u_1 \cdot z$ on a set of at least $n - e = D_X + 1$ points, and for these points $A(x,z) \cdot P_z(x) = B(x,z)$. Thus, by degree,

$$A(X,z) \cdot P_z(X) = B(X,z), \quad z \in S.$$

On the other hand, by the definition of the linear system we have

$$A(x,Z)\cdot(u_0(x)+Z\cdot u_1(x))=B(x,Z),\quad x\in\mathcal{D}.$$

Second, for that inequality (5) holds with $b_X=D_X$, $n_X=n$ and $b_Z=D_Z$, $n_Z=|S|$, we must have

$$\frac{D_Z}{|S|}<1-\frac{D_X}{n}=\gamma+\frac{1}{n},$$

since $n-D_X=e+1=\gamma n+1$. Using $D_Z\le\frac{e+h+1}{h+1}=\frac{\delta-\gamma}{\delta-2\gamma}$, we see that by our assumption on $|S|=a$, we even have

$$\frac{D_Z}{|S|}\le\gamma.$$

We thus can apply Lemma 2.2 and obtain a polynomial $P(X,Z)\in\mathbb{F}_q[X,Z]$ with $\deg P\le(k,1)$ which satisfies

$$\begin{aligned}B(X,Z)&=A(X,Z)\cdot P(X,Z),\\&=A(X,Z)\cdot(p_0(X)+Z\cdot p_1(X))\end{aligned}$$

with polynomials $p_1(X),p_2(X)$ of degree at most k. In particular,

$$A(x,Z)\cdot(u_0(x)+Z\cdot u_1(x))=A(x,Z)\cdot(p_0(x)+Z\cdot p_1(x)),$$

for all $x\in\mathcal{D}$. Dividing by $A(x,Z)$ in $\mathbb{K}$ whenever $A(x,Z)\neq0$, we conclude that

$$u_0(x)+Z\cdot u_1(x)=p_0(x)+Z\cdot p_1(x)$$

for every $x\in\mathcal{D}$, except of a set of at most $e+h$ points (for the potential zeros of A). Since $\frac{e+h}{n}=\delta-\gamma-\frac{1}{n}$, the proof of the claim is complete. □

2.3 From Claim 2.3 to Theorem 1.3

Despite the distance loss in Claim 2.3, all proximates $P_z(X)$, $z\in S$, lie on the line spanned by $p_0(X)$, $p_1(X)$. In fact, by the triangle equality

$$\begin{aligned}\Delta(P_z,p_0+z\cdot p_1)&\le\Delta(P_z,u_0+z\cdot u_1)+\Delta(u_0+z\cdot u_1,p_0+z\cdot p_1)\\&<\gamma+\delta-\gamma=\delta,\end{aligned}$$

and we conclude equality $P_z(X)=p_0(X)+z\cdot p_1(X)$, for every $z\in S$. With this in place, Theorem 1.3 follows from the following well-known lemma.

Lemma 2.4. *Assume $p_0,p_1\in\mathcal{C}$ satisfy $\Delta(u_0+z\cdot u_1,p_0+z\cdot p_1)\le\gamma$ for $a\ge2$ many values $z\in\mathbb{F}_q$. Then $\Delta([u_0,u_1],[p_0,p_1])\le\frac{a}{a-1}\cdot\lfloor n\cdot\gamma\rfloor$.*

Proof. Let $d=n\cdot\Delta([u_0,u_1],[p_0,p_1])$. We show that $d>\frac{a}{a-1}\cdot e$, with $e=\lfloor\gamma\cdot n\rfloor$, leads to a contradiction.

Consider the set of disagreement

$$E=\{x\in D\,:\,(p_0(x),p_1(x))\neq(u_0(x),u_1(x))\}.$$

Note that for each point $x\in E$ there is at most one $z\in\mathbb{F}_q$ so that $(p_0(x)-u_0(x))+z\cdot(p_1(x)-u_1(x))=0$, removing the disagreement at the point. By our assumption $d>e$, each of the claimed z with $\Delta(u_0+z\cdot u_1,p_0+z\cdot p_1)\le\gamma$ needs to remove at least $d-e$ disagreements, i.e.

$$A_z=\{x\in E\,:\,p_0(x)+z\cdot p_1(x)=u_0(x)+z\cdot u_1(x)\}$$

has size $|A_z| \geq d - e$. By the previous observation, different z produce disjoint A_z, and therefore

$$a \cdot (d - e) \leq |E|,$$

In other words $d \cdot (a - 1) \leq a \cdot e$, contradicting the assumption. This proves the claim of the lemma. □

Remark 2.5. The bound in Lemma 2.4 is tight whenever $d = a \cdot e/(a-1)$ is an integer. Too see this, start with arbitrary codewords $p_0, p_1 \in \mathcal{C}$, choose a set $E \subset \mathcal{D}$ of size d, and partition it into $a \geq 2$ disjoint subsets $E_1, \ldots E_a$, each of size $d - e$. (By assumption, $a \cdot (d - e) = d$.) Choose distinct roots $z_1, \ldots, z_a \in F$, and perturb $[p_0, p_1]$ on E via the piecewise defined function $(v_0(x), v_1(x)) = (z_i, 1)$ for $x \in E_i$, while leaving it unperturbed outside E, where we set $(v_0(x), v_1(x)) = (0, 0)$. By construction, the obtained word

$$(u_0, u_1) = [p_0, p_1] + (v_0, v_1)$$

has distance $\Delta([u_0, u_1], [p_0, p_1]) = d$, and

$$u_0(x) + z \cdot u_1(x) - (p_0(x) + z \cdot p_1(x)) = z \cdot v_1(x) - v_0(x) = \begin{cases} z - z_i & x \in E_i, \\ 0 & x \in \mathcal{D} \setminus \bigcup_i E_i. \end{cases}$$

Thus $z = z_i$ removes disagreement only on E_i, and nowhere else on E, meaning that $\Delta(u_0 + z_i \cdot u_1, p_0 + z_i \cdot p_1) = d - (d - e) = e$.

2.3.1 Proof of Corollary 1.4

We show that if $\delta \geq \frac{3\sqrt{2}}{\sqrt{n}}$ and $\frac{\delta}{3} \leq \gamma \leq \frac{\delta}{2} - \frac{3}{n\delta}$, then $a \leq \frac{1}{\delta/2 - \gamma} \leq \gamma \cdot n$ in Theorem 1.3, proving the corollary.

In terms of $\eta = \delta/2 - \gamma$ the desired inequality is $\frac{1}{\eta} \leq n \cdot (\frac{\delta}{2} - \eta)$, or

$$1 \leq n \cdot \eta \cdot \left(\frac{\delta}{2} - \eta \right),$$

which we wish to show for $\frac{3}{\delta n} \leq \eta < \frac{2\delta}{3}$. By concavity of the right-hand side on the interval $(0, \delta/2)$, it is sufficient to show the inequality at the boundaries of the interval. We have $\delta \geq 3\sqrt{2/n}$, that is $\delta^2 \geq 18/n$. At $\eta = \frac{3}{\delta n}$ we obtain

$$\frac{3}{2} - n \cdot \left(\frac{3}{\delta n} \right)^2 \geq \frac{3}{2} - \frac{9}{18} = 1,$$

and at $\eta = \frac{2\delta}{3}$ we simply get $n \cdot \frac{\delta}{3} \cdot \frac{\delta}{6} = \frac{n\delta^2}{18} \geq 1$. This completes the proof.

2.3.2 Further discussion about the statement of Theorem 1.3

In the edge case $\gamma = \frac{\delta}{2} - \frac{1}{n}$, the coarser bound yields zero distance loss for $a \geq n$, independent of δ, reproducing [BCI+20, Theorem 4.1]. In the other edge case $\gamma = \delta/3$, we obtain $a \geq 6/\delta$ for distance loss $\varepsilon^* = \frac{\delta}{6-\delta} \cdot \gamma \leq \frac{1}{5} \cdot \gamma$, whereas the example from Theorem 1.15 below has already loss $\varepsilon^* = \gamma$ (the common gap, in the light of [RVW13]) with $a = 1/\delta$. This indicates that the bound in Theorem 1.3 is not far from optimal.

3 Improved proximity gaps up to the Johnson bound

In this section we prove the following sharpening of [BCI+20, Theorem 5.1]. Recall that $\rho = 1 - \delta = \frac{k}{n}$ is the slightly reduced rate of the code.

Theorem 1.5. *Let $\mathcal{C}$ be the code $\mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ with block-length $n = |\mathcal{D}|$ and minimum distance $\delta = 1 - \frac{k}{n}$. Denote $\rho = \frac{k}{n} = 1 - \delta$. For $\gamma \in \big(0, 1 - \sqrt{\rho}\big)$, let $\eta = 1 - \sqrt{\rho} - \gamma$, and $m = \max\Big(\Big\lceil \frac{\sqrt{\rho}}{2\eta} \Big\rceil, 3\Big)$. Suppose $u_0, u_1 : \mathcal{D} \to \mathbb{F}_q$ are functions such that $S = \{z \in \mathbb{F}_q \mid \Delta(u_0 + z \cdot u_1, \mathcal{C}) \leq \gamma\}$ is of size*

$$a > \frac{2(m+1/2)^5 + 3(m+1/2)\gamma\rho}{3\rho^{3/2}} \cdot n + \frac{m+1/2}{\sqrt{\rho}} = O_\rho\left(\frac{n}{\eta^5}\right). \tag{1}$$

Then

$$\Delta([u_0, u_1], \mathcal{C}^2) \leq \gamma.$$

The proof of Theorem 1.5 studies the Guruswami–Sudan [GS99] list decoder on the word

$$u(x) = u_0(x) + Z \cdot u_1(x)$$

with values in the field of rational functions $\mathbb{K} = \mathbb{F}_q(Z)$. However, unlike in [BCI+20, Section 5], we consider the system of equations over $\mathbb{F}_q$ explicitly to get a low effective bound on the necessary degree D_Z, rather than considering equations over $\mathbb{K}$ and trying to investigate the degrees of the arbitrary solutions obtained. It turns out that the gap between the number of variables and equations is sufficient to ensure D_Z much smaller than previously obtained, see Section 3.1. This is the major source of our improvement. A further refinement on the final bound in Theorem 1.5, in terms of the degrees of the interpolant, is then discussed in Section 3.2. This part is independent of our choice of the interpolant, and applies to the regular one used in [BCI+20] as well.

3.1 Improved Guruswami–Sudan interpolant over $\mathbb{K}$

Given two functions $u_0, u_1 : \mathcal{D} \longrightarrow \mathbb{F}_q$ we consider the Guruswami–Sudan list decoder with input $u = u_0(x) + Z \cdot u_1(x)$, a word with values in the field of rational functions $\mathbb{K} = \mathbb{F}_q(Z)$. In the ordinary setting, and for integer $m \geq 1$, the decoder looks for a bivariate polynomial $Q(X, Y) \in \mathbb{K}[X, Y]$ of $(1, k)$-weighted degree

$$D_X = \sqrt{m \cdot (m+1) \cdot kn}$$

so that $Q(x, u(x)) = 0$ with multiplicity m,

$$\mathsf{mult}(Q, (x, u(x))) = m, \quad x \in \mathcal{D}. \tag{6}$$

With this setting, [BCI+20] argue a solution $Q(X, Y)$ with *polynomial* coefficients in Z, their degree bounded as $O_\delta(n)$, see e.g. Claim 5.4 therein. However, a more careful analysis of the system of equations allows for a significantly smaller Z-degree, as the following lemma shows.

Lemma 3.1. *Let $m \geq 3$ be an integer and set*

$$D_X = (m + \tfrac{1}{2}) \cdot \sqrt{nk}, \tag{7}$$

$$D_Y = (m + \tfrac{1}{2}) \cdot \sqrt{\tfrac{n}{k}}, \tag{8}$$

$$D_Z = \tfrac{1}{3}(m + \tfrac{1}{2})^2 \cdot \tfrac{n}{k}. \tag{9}$$

There exists non-zero $Q(X, Y, Z) \in \mathbb{F}_q[X, Y, Z]$ such that (6) holds for all $x \in \mathcal{D}$, and:

- *the $(1, k, 0)$-weighted degree of $Q(X, Y, Z)$ is less than D_X,*
- *the Y-degree of $Q(X, Y, Z)$ is less than D_Y, and*
- *the $(0, 1, 1)$-weighted degree of $Q(X, Y, Z)$ is less than D_Z.*

Note that the notation D_Z used here replaces the notation D_{YZ} used in [BCI+20]. The parameters chosen here are not fully optimized, but simplified to provide cleaner expressions. The $\frac{1}{2}$ in $m+\frac{1}{2}$ and the value of D_Z can be adjusted to obtain finer results. The condition $m \geq 3$ is only required to ensure $D_Z \geq D_Y$; the result can be generalized to smaller m by setting D_Z as the maximum between D_Y and the value given in Equation (9).

Proof. Given the degree bounds above, our polynomial has the form

$$Q(X,Y,Z) = \sum_{\substack{i+kj<D_X \\ j+h<D_Z}} Q_{i,j,h} X^i Y^j Z^h$$

with variables $Q_{i,j,h} \in \mathbb{F}_q$. Throughout the computations below we will make implicit use of the facts that $D_X \geq kD_Y$ (in fact equal in our choice), $D_Z \geq D_Y$ and $D_Y \geq m-1$, which ensure all enumerations made are of non-negative amounts. The Guruswami–Sudan equations (6) say that for each $x \in \mathcal{D}$ and integers $r, s \geq 0$ with $r+s<m$, we have the constraint

$$Q^{(r,s)}(x, u(x,Z), Z) = 0,$$

which can be explicitly written as

$$\sum_{\substack{i+kj<D_X \\ j+h<D_Z}} Q_{i,j,h} \cdot \binom{i}{r}\binom{j}{s} \cdot x^{i-r} \cdot (u_0(x) + Z \cdot u_1(x))^{j-s} \cdot Z^h = 0. \tag{10}$$

This is a polynomial equation in Z of degree at most $\lceil D_Z \rceil - 1 - s$, and therefore corresponds to $\lceil D_Z \rceil - s$ linear equations in the $Q_{i,j,h}$. For each $0 \leq s < m$ we have $m-s$ corresponding choices of $0 \leq r < m-s$ and n choices of $x \in \mathcal{D}$, thus the total number of equations is [3]

$$\mathrm{n_{eqs}} = n \cdot \sum_{s=0}^{m-1} (\lceil D_Z \rceil - s)(m-s) = n \cdot \left(\frac{m(m+1)}{2} \lceil D_Z \rceil - \frac{m^3-m}{6} \right). \tag{11}$$

We now count the number of variables $Q_{i,j,h}$. For each coefficient Y^j with $j < \lceil D_Y \rceil$, we have $\lceil D_X \rceil - kj$ corresponding X monomials and $\lceil D_Z \rceil - j$ corresponding Z monomials, so the number of variables $Q_{i,j,h}$ with fixed j is equal to $(\lceil D_X \rceil - kj)(\lceil D_Z \rceil - j)$, and in total we have[4]

$$\begin{aligned}\mathrm{n_{vars}} &= \sum_{j=0}^{\lceil D_Y \rceil - 1} (\lceil D_X \rceil - kj)(\lceil D_Z \rceil - j) \\ &= \left(\lceil D_X \rceil \lceil D_Y \rceil - k \frac{\lceil D_Y \rceil (\lceil D_Y \rceil - 1)}{2} \right) \lceil D_Z \rceil - \left(\frac{\lceil D_X \rceil \lceil D_Y \rceil (\lceil D_Y \rceil + 1)}{2} - k \frac{\lceil D_Y \rceil (\lceil D_Y \rceil - 1)(2\lceil D_Y \rceil - 1)}{6} \right).\end{aligned}$$

Noting that the expression for $\mathrm{n_{vars}}$ is monotone increasing in both $\lceil D_X \rceil$ and $\lceil D_Y \rceil$ (as long as the necessary inequalities are maintained), replacing them by the slightly smaller D_X and D_Y yields a lower bound on $\mathrm{n_{vars}}$, which can then be much simplified using $D_X = kD_Y$, yielding:

$$\mathrm{n_{vars}} \geq k \cdot \left(\frac{D_Y(D_Y+1)}{2} \lceil D_Z \rceil - \frac{D_Y^3 - D_Y}{6} \right). \tag{12}$$

[3] Using the identities $\sum_{j=1}^{t} j = \frac{t(t+1)}{2}$ and $\sum_{j=1}^{t} j(t-j) = \frac{t^3-t}{6}$.

[4] Using $\sum_{j=1}^{t} j^2 = \frac{t(t+1)(2t+1)}{6}$.

Comparing equations (11) and (12), we find that $\mathrm{n_{vars}} > \mathrm{n_{eqs}}$ is satisfied if:

$$k \cdot \left(\frac{D_Y(D_Y+1)}{2} \lceil D_Z \rceil - \frac{D_Y^3 - D_Y}{6} \right) > n \cdot \left(\frac{m(m+1)}{2} \lceil D_Z \rceil - \frac{m^3 - m}{6} \right)$$

which is equivalent to

$$(\rho D_Y(D_Y+1) - m(m+1)) \lceil D_Z \rceil > \rho \cdot \frac{D_Y^3 - D_Y}{3} - \frac{m^3 - m}{3}.$$

Recalling now our choice of $D_Y = (m + \frac{1}{2})\sqrt{n/k} = \frac{m+1/2}{\sqrt{\rho}}$, the last inequality becomes

$$\big((m+\tfrac{1}{2})(m+\tfrac{1}{2}+\sqrt{\rho}) - m(m+1)\big) \lceil D_Z \rceil > \frac{1}{3}\left(\frac{(m+\frac{1}{2})^3}{\sqrt{\rho}} - (m+\tfrac{1}{2})\sqrt{\rho} - (m^3 - m) \right)$$

which is equivalent to

$$((m+\tfrac{1}{2})\sqrt{\rho} + \tfrac{1}{4}) \lceil D_Z \rceil > \frac{(m+\frac{1}{2})^3}{3\sqrt{\rho}} - \frac{1}{3}\big((m+\tfrac{1}{2})\sqrt{\rho} + (m^3 - m)\big),$$

which is clearly satisfied for our choice of $D_Z = \frac{(m+1/2)^2}{3}\frac{n}{k} = \frac{(m+1/2)^2}{3\rho}$.

Thus our choice of parameters guarantees $\mathrm{n_{vars}} > \mathrm{n_{eqs}}$, which ensures the homogeneous system of $\mathrm{n_{eqs}}$ equations (10) in the $\mathrm{n_{vars}}$ variables $Q_{i,j,h}$ has a non-trivial solution, i.e. a non-zero polynomial Q with the properties claimed exists. □

3.2 Improved bound on a in terms of D_X, D_Y, D_Z

Given a polynomial $Q(X, Y, Z)$ corresponding to the word $u_0 + Zu_1$ as in Lemma 3.1, the proof in [BCI+20, Section 5] establishes the proximity gap assuming $D_X \leq (1-\gamma)mn$ and $a > 2D_X D_Y^3 D_Z$ (see e.g. [BCI+20, Equation (5.8)]). In this section we show that the same methods, with slightly improved bookkeeping, actually yield the result assuming an improved bound of

$$a > 2D_X D_Y^2 D_Z + (\gamma n + 1) D_Y. \tag{13}$$

The condition $D_X \leq (1-\gamma)mn$ simplifies as

$$\begin{aligned} D_X \leq (1-\gamma)mn \quad &\Leftrightarrow \\ (m+\tfrac{1}{2})\sqrt{kn} = (m+\tfrac{1}{2})n\sqrt{\rho} \leq (\sqrt{\rho} + \eta)mn \quad &\Leftrightarrow \\ 1 + \frac{1}{2m} \leq 1 + \frac{\eta}{\sqrt{\rho}} \quad &\Leftrightarrow \\ m \geq \frac{\sqrt{\rho}}{2\eta}. \end{aligned}$$

We can thus take $m = \left\lceil \frac{\sqrt{\rho}}{2\eta} \right\rceil$, for which the degrees provided by Lemma 3.1 plugged into (13) yield Theorem 1.5.

Let $S = \{z \in \mathbb{F}_q \mid \Delta(u_0 + z \cdot u_1, \mathcal{C}) \leq \gamma\}$. The following is a brief summary of the steps from [BCI+20, Section 5] involving the manipulations on the bound $|S| \geq 2D_X D_Y^3 D_Z$ and the proof of its sufficiency. These steps are along the lines of the well-known procedure of factoring bivariate polynomials over finite fields via Hensel lifts. However, the underlying field is the field of rational functions $\mathbb{K} = \mathbb{F}_q(Z)$, and the required finite extension is an algebraic function field, which makes concrete computations quite technical.

1. *Decomposition into irreducible and separable factors.* The polynomial $Q(X,Y,Z)$ is decomposed as $C(X,Z)\prod_i R_i(X,Y^{p^{f_i}},Z)^{e_i}$ where each R_i is irreducible and separable in the Y variable. For simplicity we assume all $f_i = 0$; the case that some $f_i > 0$ was dealt within [BCI+20, Appendix C] and the arguments there remain applicable.

2. *Finding a uniform starting point for the Hensel lift.* An $x_0 \in \mathbb{F}_q$ is chosen such that all $R_i(x_0,Y,Z)$ remain separable in Y after the substitution. They are decomnposed in $\mathbb{F}_q[Y,Z]$ as $R_i(x_0,Y,Z) = C_i(Z)\prod_j H_{ij}(Y,Z)$ where H_{ij} are irreducible, non-constant in Y and separable in Y.

3. *Focusing on a single "useful" factor.* For each $z \in S$ (except maybe D_Z values)[5] the polynomial $P_z(X)$ satisfies $R_i(X,P_z(X),z) = 0$ and $H_{ij}(x_0,P_z(x_0),z) = 0$ for some pair i,j. The number of such pairs is at most D_Y, so for some pair $(R,H) = (R_i,H_{ij})$ the set

$$S_{x_0,R,H} = \{z \in S : R(X,P_z(X),z) \equiv 0 \text{ and } H(P_z(x_0),z) = 0\}$$

must have size at least $|S_{x_0,R,H}| \geq |S|/D_Y$.

4. *Determining a smooth solution via Hensel lift.* Let $D_Y^{(R)}, D_Y^{(H)}, D_Z^{(R)}$ denote the Y degree of R, the Y degree of H, and the $(1,1)$-weighted (Y,Z) degree of R, respectively. Suppose $|S_{x_0,R,H}| > 2D_X D_Y^{(R)} D_Y^{(H)} D_Z^{(R)}$. Then $R(X,Y,Z) = Y - P(X,Z)$ with a polynomial $P(X,Z) = v_0(X) + Z\cdot v_1(X)$ with $\deg_X P \leq k$, $\deg_Z P \leq 1$ and a subset $S' \subset S$ with

$$|S'| \geq |S_{x_0,R,H}| - D_Y^{(R)} D_Y^{(H)} D_Z^{(R)} > (2D_X - 1) D_Y^{(R)} D_Y^{(H)} D_Z^{(R)}$$

such that $P(X,z) = P_z(X)$ for each $z \in S'$. Reaching this conclusion is the major part of the proof, comprising sections 5.2.5 – 5.2.7 as well as Appendix A of [BCI+20]. This part computes an approximate solution $Y = P(X)$ of $R(X,Y) = 0$ in $\mathbb{K}[X,Y]/H(X,Y)$, a finite extension of $\mathbb{K}$, via Hensel lift. The approximate solution is subsequently shown to be exact from the fact that substitution by $z \in S'$ leads to the promised proximate polynomial $P_z(X)$.

5. *Correlated agreement from collinearity of the proximates.* This is the standard argument using Lemma 2.4. If $|S'|$ above is greater than $\gamma n + 1$, then $\Delta([u_0,u_1],[v_0,v_1]) \leq \gamma$.

The sufficient bound $|S| \geq 2D_X D_Y^3 D_Z$ is thus established by combining the obtained bound $|S_{x_0,R,H}| \geq |S|/D_Y$ with the sufficient condition $|S_{x_0,R,H}| \geq 2D_X D_Y^{(R)} D_Y^{(H)} D_Z^{(R)}$, using the naive bounds $D_Y^{(R)}, D_Y^{(H)} \leq D_Y$ and $D_Z^{(R)} \leq D_Z$. This naive approach to the bounds introduced an unnecessary extra factor of D_Y, which we now observe can be saved.

We tweak the definition of $D_Z^{(R_i)}$ to be the $(1,1)$ weighted (Y,Z) degree of the content-free part of $R(x_0,Y,Z)$ (i.e. not including $C_i(Z)$), rather than all of R_i. This degree enters play in the Hensel lifting and in bounds related directly to H_{ij}-s, and all appearances are unaffected by the content $C_i(Z)$. We also denote by $D_Z^{(C)}$ the Z degree of $C(X,Z)\prod_i C_i(Z)$.

Our aim is to show that there exist some (R_i,H_{ij}) for which both inequalities

$$|S_{x_0,R_i,H_{ij}}| \geq 2D_X D_Y^{(R_i)} D_Y^{(H_{ij})} D_Z^{(R_i)} \tag{14}$$

$$|S_{x_0,R_i,H_{ij}}| > D_Y^{(R_i)} D_Y^{(H_{ij})} D_Z^{(R_i)} + \gamma n + 1 \tag{15}$$

hold; the latter is required so that the final S' is not smaller than $\gamma n + 1$.

[5] This appears to be a previously unnoticed minor omission in [BCI+20]: some roots $(X,P_z(X),z)$ of $Q(X,Y,Z)$ could be due to $C(X,z)$ vanishing rather than any R_i, and similarly $C_i(z)$ might vanish instead of any H_{ij}, and so not every $z \in S$ is necessarily accounted for in the $S_{x_0,R,H}$-s. This can be accounted for without affecting the necessary bound on S, as will be described below. For the purpose of describing the proof in [BCI+20] we ignore this issue.

Suppose not; then for each i, j we have $|S_{x_0,R_i,H_{ij}}| \le 2D_X D_Y^{(R_i)} D_Y^{(H_{ij})} D_Z^{(R_i)} + \gamma n + 1$. We take the sum of this expression over all i, j, noting that each element of S (except at most $D_Z^{(C)}$ roots of $C(X,Z)\prod_i C_i(Z)$) must belong to some $S_{x_0,R_i,H_{ij}}$, and using the simple identities

$$\sum_j D_Y^{(H_{ij})} = D_Y^{(R_i)} \ ; \quad \sum_i D_Y^{(R_i)} = D_Y; \quad D_Z^{(R_i)} \le D_Z - D_Z^{(C)}$$

we get

$$\begin{aligned}
|S| &\le D_Z^{(C)} + \sum_i \sum_j |S_{x_0,R_i,H_{ij}}| \\
&\le D_Z^{(C)} + \sum_i \sum_j \Big(2D_X D_Y^{(R_i)} D_Y^{(H_{ij})} D_Z^{(R_i)} + \gamma n + 1 \Big) \\
&\le D_Z^{(C)} + (\gamma n + 1) D_Y + 2D_X D_Y (D_Z - D_Z^{(C)}) \sum_i D_Y^{(R_i)} \\
&= D_Z^{(C)} + (\gamma n + 1) D_Y + 2D_X D_Y^2 (D_Z - D_Z^{(C)}) \\
&\le 2D_X D_Y^2 D_Z + (\gamma n + 1) D_Y .
\end{aligned}$$

Thus, taking $|S| > 2D_X D_Y^2 D_Z + (\gamma n + 1) D_Y$ indeed suffices for Theorem 1.5, as claimed.

4 Generalizations

In this section we state generalizations of the positive Theorems 1.3 and 1.5.

4.1 More general linear combinations

In practice one desires proximity gaps for the interleaved code $\mathcal{C}^{M+1}$, often with large integer M. Given functions $u_0, \ldots, u_M : \mathcal{D} \longrightarrow \mathbb{F}_q$, where $M \ge 1$, one wishes to infer $\Delta([u_0, \ldots, u_M], \mathcal{C}^{M+1}) \le \gamma$ from the fact that sufficiently many linear combinations of $u_0, \ldots, u_M$ are close to $\mathcal{C}$. The multilinear case,

$$u_{z_1,\ldots,z_m} = \sum_{i=0}^{M} z_0^{i_0} \cdots z_{m-1}^{i_{m-1}} \cdot u_i,$$

where $i_0, \ldots, i_{m-1}$ are the $m = \lceil \log M \rceil$ bits of i, can be derived from elementary line case covered by Theorem 1.3 and 1.5, with a factor l larger bounds [DG25b]. However, of particular importance in practice is the case

$$u_z = u_0 + u_1 \cdot z + \ldots + u_L \cdot z^M,$$

since it requires only a single randomness z. Both Section 2 and Section 3 can be generalized to this case, by considering the word $u(x) = u_0(x) + u_1(x) \cdot Z + \ldots + u_M(x) \cdot Z^M$ as input of the decoder.

We only state the lossless results.

Theorem 4.1 (Correlated agreement for curves, up to $\delta/2$)**.** *Let $\mathcal{C}$ be the code $\mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ with block-length $n = |\mathcal{D}|$, dimension $k+1$, and minimum distance $\delta \ge \frac{3\cdot\sqrt{2}}{\sqrt{n}}$. Then for any $\gamma \in \left[\frac{\delta}{3}, \frac{\delta}{2} - \frac{3}{\delta n}\right]$ and functions $u_0, \ldots, u_M : \mathcal{D} \to \mathbb{F}_q$ the following holds. If*

$$\left| \{ z \in \mathbb{F}_q \mid \Delta(u_0 + z \cdot u_1 + \ldots + z^M \cdot u_M, \mathcal{C}) \le \gamma \} \right| > M \cdot (\gamma n + 1),$$

then

$$\Delta([u_0, \ldots, u_M], \mathcal{C}^{M+1}) \le \gamma.$$

In words, there exist polynomials $p_0(X), \ldots, p_M(X) \in \mathbb{F}_q[X]$ of degree at most k, which agree with $u_0, \ldots, u_M$ on a joint set A of density $|A|/n \geq 1 - \gamma$.

In a nutshell, the proof of Theorem 4.1 adapts the interpolant Lemma 2.1 to the M times higher Z-degree, which translates to the same blow-up of the bound for a in Claim 2.3. The collinearity argument from Lemma 2.4 extends to degree M curves with distance loss $\frac{M}{a-M} \cdot \gamma$ instead of $\frac{1}{a-1} \cdot \gamma$. Zero distance loss is obtained by requiring $\frac{M}{a-M} \cdot \gamma < \frac{1}{n}$, that is $a > M \cdot (\gamma n + 1)$.

Likewise, the oversized Guruswami–Sudan interpolant from Lemma 3.1 requires an M times larger bound for D_Z, which is now the $(0, M, 1)$-weighted degree of Q, and the final bound for a again scales by the same factor.

Theorem 4.2 (Correlated agreement for curves, up to Johnson bound)**.** *Let $\mathcal{C}$ be the code $\mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ with block-length $n = |\mathcal{D}|$, dimension $k+1$ and minimum distance δ. Denote $\rho = 1 - \delta$, the slightly reduced rate of the code. Then for any $\gamma \in \left(0, 1 - \sqrt{\rho}\right)$ and functions $u_0, \ldots, u_M : \mathcal{D} \to \mathbb{F}_q$ the following holds. If*

$$\left|\{z \in \mathbb{F}_q \mid \Delta(u_0 + z \cdot u_1 + \ldots + z^M \cdot u_M, \mathcal{C}) \leq \gamma\}\right| > M \cdot \left(\frac{2(m+1/2)^5 + 3(m+1/2)\gamma\rho}{3\rho^{3/2}} \cdot n + \frac{m+1/2}{\sqrt{\rho}}\right),$$

where $m = \max\left(\left\lceil \frac{\sqrt{\rho}}{1-\sqrt{\rho}-\gamma} \right\rceil, 3\right)$, then

$$\Delta([u_0, \ldots, u_M], \mathcal{C}^{M+1}) \leq \gamma.$$

In words, there exist polynomials $p_0(X), \ldots, p_M(X) \in \mathbb{F}_q[X]$ of degree at most k, which agree with $u_0, \ldots, u_M$ on a joint set A of density $|A|/n \geq 1 - \gamma$.

4.2 Constrained agreement

The following strengthened statements on correlated agreement are useful in the soundness analysis of interactive proofs of proximity such as FRI [BBHR18, BCI+20], Basefold [ZCF23], FRI-Binius [DP24] and WHIR [ACFY25]. These refined properties are a direct consequence of collinearity (or, in the light of Section 4.1, co-curvilinearity) of proximate polynomials, the core property behind correlated agreement, and they are obtained by minor strengthenings of Lemma 2.4.

For brevity, we restrict to the curve case.

Theorem 4.3 (Correlated agreement with given sets, up to Johnson bound)**.** *For $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ and δ, ρ, γ as in Theorem 4.2, and words $u_0, \ldots, u_M : \mathcal{D} \longrightarrow \mathbb{F}_q$. Suppose that*

$$S = \{z \in \mathbb{F}_q \mid \Delta(u_0 + z \cdot u_1 + \ldots + z^M \cdot u_M, \mathcal{C}) \leq \gamma\}$$

is as large as in Theorem 4.2, and let

$$\{A_z\}_{z \in S}$$

be any choice of agreement sets with density $1 - \gamma$. (That is, for each $z \in S$ there exists a proximate $p_z \in \mathcal{C}$ such that $u_0 + z \cdot u_1 + \ldots + z^M \cdot u_M = p_z$ on A_z, and $|A_z|/|\mathcal{D}| \geq 1 - \gamma$.) Then there exists $z_0 \in S$ and $[p_0, \ldots, p_M] \in \mathcal{C}^{M+1}$ such that

$$[u_0, \ldots, u_M]\big|_{A_{z_0}} = [p_0, \ldots, p_M]\big|_{A_{z_0}}.$$

The theorem is obtained by plugging in the improved bounds in the proof of [Sta25, Theorem 22]. Notably, the theorem immediately implies the *weighted correlated agreement theorem* in a simpler form than stated by [BCI+20, Theorem 7.2], and with no restriction on the weights. Assume a sub-probability measure on $\mathcal{D}$ of the form

$$\mu(\{x\}) = \frac{w(x)}{|\mathcal{D}|},$$

for some function w with values in $[0, 1]$.

Corollary 4.4 (Weighted correlated agreement, up to Johnson bound)**.** *Given $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ and δ, ρ, γ as in Theorem 4.2, and words $u_0, \ldots, u_M : \mathcal{D} \longrightarrow \mathbb{F}_q$, and let μ be a sub-probability measure as above. Suppose that*

$$S = \left\{ z \in F \ : \ \exists p_z \in \mathcal{C}, \mu(\{x : u_z(x) = p_z(x)\}) \geq 1 - \gamma \right\}$$

is as large as in Theorem 4.2, where $u_z = u_0 + u_1 \cdot z + \ldots + u_L \cdot z^M$. Then there exists $[p_0, \ldots, p_M] \in \mathcal{C}^{M+1}$, and a set A of weight $\mu(A) \geq 1 - \gamma$ such that $[p_0, \ldots, p_M] = [u_0, \ldots, u_M]$ on A.

Proof. Since μ is dominated by the regular density, the agreement sets A_z for the scalars $z \in S$ of weight $\mu(A_z) \geq 1 - \gamma$ are also of regular density $|A_z|/|D| \geq 1 - \gamma$. We thus may apply Theorem 4.3 to see that, over one of these agreement sets we have correlated agreement. □

The error bound of Corollary 4.4 replaces the $O(n^2)$-bound from [BCI+20], used in any round-by-round analysis of FRI in the list-decoding regime ([BCI+20], see also [HLP24] or [Sta25]) yielding significantly smaller soundness errors.

Of similar importance is the following agreement theorem where the proximate polynomials are subject to given linear constraints. This is the situation met in Basefold [ZCF23], and a non-linear generalization which we do not cite here is useful for WHIR and FRI-Binius, see [Hab24]. Again, its proof is obtained by plugging in the improved bounds in the proof of [Hab24, Theorem 3].

Theorem 4.5. *(Weighted correlated agreement for subspaces) For $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ and δ, ρ, γ as in Theorem 4.2, and words $u_0, \ldots, u_M : \mathcal{D} \longrightarrow \mathbb{F}_q$, and μ a sub-probability measure as above. Suppose that $\mathcal{C}'$ is a linear subspace of $\mathcal{C}$, and that*

$$S = \left\{ z \in F \ : \ \exists p_z \in \mathcal{C}', \mu(\{x : u_z(x) = p_z(x)\}) \geq 1 - \gamma \right\},$$

with $u_z = u_0 + u_1 \cdot z + \ldots + u_L \cdot z^M$, is of size

$$a > M \cdot \left(\frac{2(m + 1/2)^5 + 3(m + 1/2)\gamma\rho}{3\rho^{3/2}} \cdot n + \frac{m + 1/2}{\sqrt{\rho}} \right).$$

Then there exists $[p_0, \ldots, p_M] \in (\mathcal{C}')^{M+1}$ and a set A of weight $\mu(A) \geq 1 - \gamma$, such that $[p_0, \ldots, p_M] = [u_0, \ldots, u_M]$ on A.

4.3 List correlated agreement

List correlated agreement, also called *strong correlated agreement* in [Zei24] or *mutual correlated agreement* in [ACFY25], is a "global" variant of regular correlated agreement. Given two functions $u_0, u_1 : \mathcal{D} \longrightarrow \mathbb{F}_q$ and proximity parameter γ, the list correlated agreement property claims that, except for few scalars z, every γ-proximate $p_z \in \mathcal{C}$ of $u_0 + zu_1$ stems from a γ-proximate $[p_0, p_1] \in \mathcal{C}^2$ of $[u_0, u_1]$. The property has been shown for general linear codes up to the double-Johnson bound in [Zei24], and up to one-and-a-half Johnson bound in [GKL24]. For Reed–Solomon codes it has been conjectured to hold up to the Johnson bound [ACFY25, Conjecture 4.12]; a proof of it, which generalizes the decoder analysis from [BCI+20] is discussed in [Hab25].

We state the result with the improved bounds.

Theorem 4.6 (List correlated agreement, up to Johson bound)**.** *Let $\mathcal{C} = \mathsf{RS}_k[\mathbb{F}_q, \mathcal{D}, k]$ be the Reed–Solomon code over $\mathbb{F}_q$ with domain of definition $\mathcal{D}$ of size $|\mathcal{D}| = n$, and dimension $k+1$. Denote $\rho = k/n$, the slightly reduced rate of the code.*

Then for any $u_0, \ldots, u_M : \mathcal{D} \longrightarrow \mathbb{F}_q$, and $\gamma \in (0, 1 - \sqrt{\rho})$, the size of

$$E = \left\{ z \in \mathbb{F}_q \ : \ \begin{array}{c} \exists A \subset \mathcal{D}, \\ |A| \geq (1 - \gamma) \cdot n, \end{array} \ s.t. \ \begin{array}{c} (u_0 + zu_1 + \ldots + z^M u_M)\big|_A \in \mathcal{C}\big|_A, \ but \\ [u_0, \ldots, u_M]\big|_A \notin \mathcal{C}^{M+1}\big|_A \end{array} \right\}$$

is bounded as

$$|E| \le M \cdot \left(\frac{2(m+1/2)^5 + 3(m+1/2)\gamma\rho}{3\rho^{3/2}} \cdot n + \frac{m+1/2}{\sqrt{\rho}} \right),$$

where $m = \max\left(\left\lceil \frac{\sqrt{\rho}}{1-\sqrt{\rho}-\gamma} \right\rceil, 3 \right)$.

The proof of Theorem 4.6 generalizes Section 3.2 as follows. Instead focusing on a *single* "useful" factor of

$$Q(X,Y,Z) = C(X,Z) \cdot \prod_i R_i(X,Y,Z)$$

(for simplicity we restrict to the separable case) one looks at *all* such useful factors $R_i(X,Y,Z)$, where useful means that $R_i(X, p_z(X), Z) = 0$ for sufficiently many of the claimed γ-proximate $p_z(X)$. (Again, each useful factor is decomposed into irreducible and separable components $R_i(x_0, Y, Z) = C_i(Z) \cdot \prod_j H_{i,j}(Y,Z)$.) Each such good factor can only cause few proximates which correct an error of $[u_0, \ldots, u_M]$ (since most of their proximates are co-curvilinear), and the remaining non-useful factors cover few proximates by definition.

5 The failure of n^τ-bounded proximity gaps

In this section we prove Theorem 1.6 showing that a has to be superpolynomial if we want proximity gaps at radius $\delta - o(1)$. This also immediately gives us Corollary 1.7, showing that a has to be nearly quadratic in n if we want proximity gaps at the Johnson radius.

Recall the statement of theorem:

Theorem 1.6. *Let τ be a fixed positive integer, and $\lambda_\tau = 2^{-(\tau+2)}$. Let $\epsilon > 0$ be an arbitrary constant, and choose $\delta = 1 - \lambda_\tau$ and $\gamma = 1 - 4\lambda_\tau$.*

Then for all $\mathbb{F}_q$ of characteristic 2, there are Reed-Solomon codes $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, (1-\delta)n]$ over $\mathbb{F}_q$, domain $\mathcal{D}$ with $n = |\mathcal{D}| = O\left(q^{\frac{1}{\tau}(1+\epsilon)}\right)$, distance δ and words $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:

$$\left|\{z \in \mathbb{F}_q \mid \Delta(f + zg, \mathcal{C}) \le \gamma\}\right| \ge (1 - o(1)) \cdot q \ge n^{\tau(1-\epsilon)},$$

for q large enough, and yet $\Delta([f,g], \mathcal{C}^2) \ge 1 - 2\lambda_\tau = \frac{2}{3}\delta + \frac{1}{3}\gamma$.

We will need a special family of polynomials with many roots and well-behaved coefficients. We will find this special family through a combination of ideas from algebra and probability: using structural properties of the coefficients of *subspace polynomials*, and the second moment method to understand the distribution of these coefficients. This is done in Theorem 5.1 below.

Throughout this section, we assume that $q = 2^m$, for some integer $m \ge 1$. We will work with $\mathbb{F}_q$ and the $\mathbb{F}_2$-linear structure of it. For any $\mathbb{F}_2$-linear subspace $W \subseteq \mathbb{F}_q$ of dimension b, recall that the monic vanishing polynomial $P_W(X)$ of W, called the *subspace polynomial* of W, is a linearized polynomial of degree 2^b:

$$P_W(X) = X^{2^b} + c_1 X^{2^{b-1}} + c_2 X^{2^{b-2}} + \ldots + c_b X.$$

Define $\Lambda(W)$ to be the coefficient c_1 of $X^{2^{b-1}}$ in $P_W(X)$.

For an $\mathbb{F}_2$-linear subspace (in short $\mathbb{F}_2$-subspace) $V \subseteq \mathbb{F}_q$, we define $\mathsf{supp}_b(V)$ by:

$$\mathsf{supp}_b(V) = \left|\{\lambda \in \mathbb{F}_q \mid \exists\ \mathbb{F}_2\text{-subspace } W \subseteq V \text{ with } \dim_{\mathbb{F}_2}(W) = b, \Lambda(W) = \lambda)\}\right|.$$

We will informally call $\mathsf{supp}_b(V)$ the *diversity* of V.

Theorem 5.1 (Subspaces with large diversity)**.** *Let $a, b \in [m]$ with $b \cdot (a-b) \geq (1+\epsilon)m$. Then there exists an a-dimensional $\mathbb{F}_2$-subspace $V \subseteq \mathbb{F}_q$ with*

$$\mathsf{supp}_b(V) \geq \left(1 - O_\epsilon(q^{-\epsilon})\right) \cdot q.$$

Heuristically: there are about $2^{b\cdot(a-b)} \geq 2^{(1+\epsilon)m} = q^{1+\epsilon}$ different b-dimensional $\mathbb{F}_2$-subspaces W of V. (In V, the number of ordered linear independent sets of b vectors is $\prod_{i=0}^{b-1}(2^a - 2^i)$, and each b-dimensional $\mathbb{F}_2$-subspace has $\prod_{i=0}^{b-1}(2^b - 2^i)$ ordered bases.) If the values of $\Lambda(W)$ as W varies over b-dimensional $\mathbb{F}_2$-subspaces of V are roughly uniformly distributed over $\mathbb{F}_q$, then all values in $\mathbb{F}_q$ ought to appear as $\Lambda(W)$ for some subspace W of V.

Note that if V is contained in a subfield K of $\mathbb{F}_q$, then for any subspace W of V, $\Lambda(W)$ also lies in K. So not every V satisfies the conclusion of Theorem 5.1 (but most V do, as our proof will show).

Using this theorem, we can proceed.

Proof. Proof of Theorem 1.6. Let $q = 2^m$ be an arbitrary power of 2. Choose $b = \left\lceil \frac{(1+\epsilon)m}{\tau} \right\rceil$, where $\epsilon > 0$ is the arbitrary small constant in the statement of the theorem, and let $a = b + \tau$. By Theorem 5.1, there exists an a-dimensional $\mathbb{F}_2$-subspace $V \subseteq \mathbb{F}_q$ with $\mathsf{supp}_b(V) \geq (1 - q^{-\epsilon}) \cdot q$. Let $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, where $\mathcal{D} = V$ and $k = 2^{b-2}$. Then $\mathcal{C}$ has blocklength

$$n = |V| = 2^a = O(q^{\frac{1+\epsilon}{\tau}}),$$

so that $n^{\tau(1-\epsilon)} = q \cdot (1 + O(q^{-\epsilon^2}))$ which smaller than any $q \cdot (1 - o(1))$ for q large enough. The minimum distance of the code is

$$\delta = 1 - \frac{k}{n} = 1 - \frac{2^{b-2}}{2^a} = 1 - 2^{-(\tau+2)} = 1 - \lambda_\tau.$$

Take $f : V \to \mathbb{F}_q$ to be the function $f(x) = x^{2^b}$. Take $g : V \to \mathbb{F}_q$ to be the function $g(x) = x^{2^{b-1}}$. By Theorem 5.1, for least $(1 - o(1))q$ choices of $\lambda \in \mathbb{F}_q$, there exists some $\mathbb{F}_2$-subspace $W \subseteq V$ of dimension $b = a - \tau$ with a subspace polynomial $P_W(X)$ of the form:

$$P_W(X) = X^{2^b} + \lambda X^{2^{b-1}} + (\text{ some polynomial of degree} \leq 2^{b-2}).$$

This means that $f + \lambda g$ agrees with some polynomial of degree at most $2^{b-2} = k$ on all the points of W, in particular on at least 2^b points. Thus for such $\lambda \in \mathbb{F}_q$, we have

$$\Delta(f + \lambda g, \mathcal{C}) \leq 1 - \frac{2^b}{2^a} = 1 - 2^{-\tau} = 1 - 4\lambda_\tau.$$

On the other hand, since g is a polynomial of degree at most 2^{b-1}, it has at most $2^{b-1} = 2 \cdot \lambda_\tau \cdot n$ agreement evaluation points with polynomials of degree at most $k = 2^{b-2}$ – which proves that $\Delta(g, \mathcal{C}) \geq 1 - 2 \cdot \lambda_\tau$, and thus $\Delta([f, g], \mathcal{C}^2) \geq 1 - 2 \cdot \lambda_\tau$. This completes the proof of Theorem 1.6. □

The rest of this section is devoted to proving Theorem 5.1.

5.1 Finding a subspace with large diversity

For an $\mathbb{F}_2$-subspace $V \subseteq \mathbb{F}_q$ of dimension $\dim(V) = a$ and an integer $b \in [0, a]$, we let $\begin{bmatrix} V \\ b \end{bmatrix}$ denote the set of all b-dimensional subspaces of V.

Proof of Theorem 5.1. For an $\mathbb{F}_2$-subspace $V \subseteq \mathbb{F}_q$, we introduce a parameter $\mathsf{score}_b(V)$. Let W be a uniformly random b-dimensional $\mathbb{F}_2$-subspace of V. Then we define:

$$\mathsf{score}_b(V) = \sum_{\lambda \in \mathbb{F}_q} \left(\Pr_W[\Lambda(W) = \lambda] \right)^2.$$

This measures how well spread the values of $\Lambda(W)$ are as W varies over all b-dimensional $\mathbb{F}_2$-subspaces of V. Observe that we also have an alternative interpretation of $\mathsf{score}_b(V)$ as follows:

$$\mathsf{score}_b(V) = \Pr_{W,W'}[\Lambda(W) = \Lambda(W')],$$

where W and W' are independent, uniformly chosen b-dimensional $\mathbb{F}_2$-subspaces of V. By the Cauchy–Schwarz inequality, we get a simple relation between $\mathsf{supp}_b(V)$ and $\mathsf{score}_b(V)$:

$$\mathsf{supp}_b(V) \geq \frac{\left(\sum_{\lambda \in \mathbb{F}_q} \Pr_W[\Lambda(W) = \lambda] \right)^2}{\mathsf{score}_b(V)} = \frac{1}{\mathsf{score}_b(V)}.$$

Thus to prove the theorem, we need to find a V with small $\mathsf{score}_b(V)$. We do this by averaging. Pick V uniformly at random from amongst all a-dimensional $\mathbb{F}_2$-subspaces of $\mathbb{F}_q$.

$$\mathbf{E}_V[\mathsf{score}_b(V)] = \mathbf{E}_V \left[\Pr_{W,W' \in \begin{bmatrix} V \\ b \end{bmatrix}} [\Lambda(W) = \Lambda(W')] \right] = \Pr_{(W,W') \in \mu} [\Lambda(W) = \Lambda(W')],$$

where W and W' are chosen uniformly and independent from all b-dimensional subspaces of the randomly selected V, and μ is the distribution of (W, W') created by the process.

Let us a closer look at the distribution μ. Let $\tau = a - b$, and K be the random variable $b - \dim(W \cap W')$. Observe that K takes values between 0 and τ. By symmetry, it is clear that the distribution of (W, W') can be equivalently generated by first choosing $K \in [0, \tau]$, then picking a uniformly random $b - K$ dimensional $\mathbb{F}_2$-subspace U of $\mathbb{F}_q$, and then picking a pair of $\mathbb{F}_2$-subspaces (W, W') of $\mathbb{F}_q$, each of dimension b, with $W \cap W' = U$, uniformly random amongst all such pairs.

The following lemma, proved in the next subsection, shows that conditioned on K being nonzero, the probability of $\Lambda(W) = \Lambda(W')$ is very small:

$$\Pr[\Lambda(W) = \Lambda(W') \mid K \neq 0] \leq \frac{1}{q} + O\left(\frac{2^{2\tau}}{q^2} \right).$$

Lemma 5.2 (Λ-collisions lemma)**.** *Let $0 \leq c < b \leq m$. Consider the following experiment. Let U be a uniformly random c-dimensional $\mathbb{F}_2$-linear subspace of $\mathbb{F}_q$. Let (W, W') be a pair of $\mathbb{F}_2$-linear subspaces of dimension b, picked uniformly at random amongst all pairs of b-dimensional $\mathbb{F}_2$-subspaces, the intersection of which is exactly U. Then*

$$\Pr[\Lambda(W) = \Lambda(W')] \leq \frac{1}{q} + O\left(\frac{2^{2(b-c)}}{q^2} \right).$$

Furthermore, the probability that $K = 0$ is simply the probability that $W = W'$, and this is also small: Since $\left| \begin{bmatrix} V \\ b \end{bmatrix} \right| = \prod_{i=0}^{b-1} \frac{2^a - 2^i}{2^b - 2^i} \geq 2^{b \cdot (a-b)}$, we obtain

$$\Pr[K = 0] = \Pr[W = W'] \leq \frac{1}{\left| \begin{bmatrix} V \\ b \end{bmatrix} \right|} \leq \frac{1}{2^{\tau \cdot (a-\tau)}}.$$

Putting these together, we get:

$$\mathbf{E}_V[\mathsf{score}_b(V)] = \Pr[\Lambda(W) = \Lambda(W')] \leq \Pr[K = 0] + \Pr[\Lambda(W) = \Lambda(W') \mid K \neq 0]$$

$$\leq \frac{1}{2^{\tau\cdot(a-\tau)}} + \frac{1}{q} + O\left(\frac{2^{2\tau}}{q^2}\right).$$

By our assumptions on τ, a, the right hand side is at most $(1 + O(q^{-\epsilon})) \cdot \frac{1}{q}$. Therefore, by the probabilistic method, there exists a V with $\mathsf{score}_b(V) \leq (1 + O(q^{-\epsilon})) \cdot \frac{1}{q}$. For any such V, we have

$$\mathsf{supp}_b(V) \geq \frac{1}{\mathsf{score}_b(V)} \geq \left(1 - O\left(q^{-\epsilon}\right)\right) \cdot q,$$

as desired. □

5.2 Preparations for the proof of the Λ-collisions lemma

We will need two ingredients: Berlekamp duality, and the relationship between $\Lambda(U)$ and $\Lambda(W)$ for $\mathbb{F}_2$-linear subspaces $U \subseteq W$.

Definition 5.3 (Berlekamp Dual)**.** *Let U be an $\mathbb{F}_2$-linear subspace of $\mathbb{F}_q$, and $P_U(X)$ the vanishing polynomial of U. The Berlekamp dual of U is defined as*

$$U^* = P_U(\mathbb{F}_q).$$

Lemma 5.4 (Berlekamp Duality)**.** *Let U be an $\mathbb{F}_2$-linear subspace of $\mathbb{F}_q$ with $\dim(U) = c$. Then:*

1. *U^* is an $\mathbb{F}_2$-linear subspace of $\mathbb{F}_q$ with $\dim(U^*) = m - c$.*
2. *$(U^*)^* = U$.*

The proof of the lemma can be found in [Ber15] (Theorem 11.35). For another (totally unrelated) application of this machinery, see [BSK12] (the above lemma is also reproved there).

Lemma 5.5 ($\Lambda(W)$ and $\Lambda(U)$)**.** *Suppose $U \subseteq W$ are $\mathbb{F}_2$-linear subspaces of $\mathbb{F}_q$ with dimension c, b respectively. Then*

$$\Lambda(W) = \Lambda(U)^{2^{b-c}} + \Lambda(L),$$

where $L = P_U(W)$.

Proof. Note that the quotient W/U has dimension $b - c$. Thus, since P_U is $\mathbb{F}_2$-linear with kernel U, we get that $L = P_U(W)$ has dimension exactly $b - c$, and that $P_U^{-1}(L) = W$. Therefore the monic polynomial

$$P_L(P_U(X)) = \prod_{\alpha \in L} (P_U(X) - \alpha),$$

which has degree $|L| \cdot \deg(P_U(X)) = 2^b$, vanishes at all points of W. Thus it must equal $P_W(X)$. Expanding $P_L(P_U(X))$, we obtain:

$$\begin{aligned} P_L(P_U(X)) &= P_U(X)^{2^{b-c}} + \Lambda(L) P_U(X)^{2^{b-c-1}} + \left(\deg \leq 2^{b-2}\right) \\ &= \left(X^{2^c} + \Lambda(U) X^{2^{c-1}} + (\deg \leq 2^{c-2})\right)^{2^{b-c}} + \Lambda(L) \left(X^{2^c} + (\deg \leq 2^{c-1})\right)^{2^{b-c-1}} + \left(\deg \leq 2^{b-2}\right) \\ &= X^{2^b} + \left(\Lambda(U)^{2^{b-c}} + \Lambda(L)\right) X^{2^{b-1}} + \left(\deg \leq 2^{b-2}\right), \end{aligned}$$

where $(\deg \leq 2^r)$ is short for a polynomial of degree at most 2^r. This proves the claimed formula for $\Lambda(W)$. □

5.3 Proof of Λ-collisions lemma

We can now prove Lemma 5.2.

Proof. Our proof relies on an alternate way to generate the distribution of (U, W, W'). The key point is that specifying U is equivalent to specifying U^*, and that specifying a superspace $W \supseteq U$ is equivalent to specifying the subspace $P_U(W)$ of $U^* = P_U(\mathbb{F}_q)$.

First pick a uniformly random $(m-c)$-dimensional subspace $\widetilde{U}$. Then pick a pair of $(b-c)$-dimensional subspaces (L, L') of $\widetilde{U}$ with $L \cap L' = \{0\}$, uniformly random among all such pairs. We then set

$$\begin{aligned} U &= (\widetilde{U})^*, \\ W &= P_U^{-1}(L), \\ W' &= P_U^{-1}(L'). \end{aligned}$$

Note that the distribution of (U, W, W') generated in this manner is exactly as in the hypothesis. Indeed, from (U, W, W') we can recover $(\widetilde{U}, L, L')$ as $(U^*, P_U(W), P_U(W'))$.

Now by Lemma 5.5,

$$\begin{aligned} \Lambda(W) &= \Lambda(U)^{2^{b-c}} + \Lambda(L), \\ \Lambda(W') &= \Lambda(U)^{2^{b-c}} + \Lambda(L'). \end{aligned}$$

Thus $\Lambda(W) = \Lambda(W')$ if and only if $\Lambda(L) = \Lambda(L')$. Thus, the probability that we want to bound, $\Pr_{(U,W,W')}[\Lambda(W) = \Lambda(W')]$, is given by the exact expression:

$$\Pr_{(U,W,W')}[\Lambda(W) = \Lambda(W')] = \Pr_{(\widetilde{U},L,L')}[\Lambda(L) = \Lambda(L')].$$

This right hand side only depends on the marginal distribution of (L, L'), which we now investigate.

By symmetry, picking the uniformly random $(m-c)$-dimensional $\mathbb{F}_2$-subspace $\widetilde{U}$ and then picking L, L' of dimension $b-c$ contained in $\widetilde{U}$ with $L \cap L' = \{0\}$ is exactly the same as *first* picking $\mathbb{F}_2$-subspaces L, L' of the entire space $\mathbb{F}_q$ of dimension $b-c$ with $L \cap L' = \{0\}$, uniformly at random from among all such pairs, and *then* picking $\widetilde{U}$ to be a $(m-c)$-dimensional space uniformly at random from among all such spaces that contain both L and L'. Thus the marginal distribution of (L, L') is uniform over all pairs of $(b-c)$-dimensional subspaces of $\mathbb{F}_q$ with intersection $\{0\}$.

The following lemma, proved next, gives a strong bound on the probability that $\Lambda(L) = \Lambda(L')$ under this distribution.

Lemma 5.6. *Let (L, L') be a pair of e-dimensional subspaces of $\mathbb{F}_q$ with intersection $\{0\}$, picked uniformly at random from amongst all such pairs.*

Then:

$$\Pr[\Lambda(L) = \Lambda(L')] \leq \frac{1}{q} + O\left(\frac{2^{2e}}{q^2}\right).$$

This completes the proof of Lemma 5.2. □

5.4 Proof of Lemma 5.6

We now prove Lemma 5.6.

Proof. We can equivalently generate the distribution (L, L') as follows. Pick $\alpha_1, \ldots, \alpha_e$, $\alpha'_1, \ldots, \alpha'_e \in \mathbb{F}_q$ independently and uniformly at random. In the sequel we condition on the event I, that all these $2e$ elements are $\mathbb{F}_2$-linearly independent. Then define $L = \mathsf{span}(\alpha_1, \ldots, \alpha_e)$ and $L' = \mathsf{span}(\alpha'_1, \ldots, \alpha'_e)$. The event I implies that L and L' both have dimension e, while $L \cap L' = \{0\}$. By symmetry, the distribution of (L, L') is uniform over all such pairs, matching the distribution in the hypothesis of the lemma.

Notice that $\Lambda(L)$ can be expressed as the 2^{e-1}-th elementary symmetric polynomial of all the elements of L. Since the elements of L are precisely $\{\sum_{i \in S} \alpha_i \mid S \subseteq [e]\}$, we get that

$$\Lambda(L) = H_e(\alpha_1, \ldots, \alpha_e),$$

and likewise

$$\Lambda(L') = H_e(\alpha'_1, \ldots, \alpha'_e),$$

where

$$H_e(Z_1, \ldots, Z_e) = \sum_{J \subseteq \mathcal{P}([e]), |J| = 2^{e-1}} \prod_{S \in J} \left(\sum_{i \in S} Z_i \right).$$

Note that $H_e(Z_1, \ldots, Z_e)$ is a non-zero, homogeneous polynomial of degree 2^{e-1}. The latter is obvious from its explicit form, and the first fact follows from that the coefficient of $Z_1^{2^{e-1}}$ equals 1: this coefficient comes from the unique subset $J \subseteq \mathcal{P}([e])$ of size 2^{e-1} that contains all subsets S of $[e]$ with $1 \in S$.

Thus we want to bound:

$$\Pr_{\alpha_1, \ldots, \alpha_e, \alpha'_1, \ldots, \alpha'_e \in \mathbb{F}_q} [H_e(\alpha_1, \ldots, \alpha_e) = H_e(\alpha'_1, \ldots, \alpha'_e) \mid I].$$

For that, we first bound

$$\Pr_{\alpha_1, \ldots, \alpha_e, \alpha'_1, \ldots, \alpha'_e \in \mathbb{F}_q} [H_e(\alpha_1, \ldots, \alpha_e) = H_e(\alpha'_1, \ldots, \alpha'_e)],$$

that is the collision probability of the distribution ν of $H_e(\alpha_1, \ldots, \alpha_e)$ for uniformly random $(\alpha_1, \ldots, \alpha_e) \in \mathbb{F}_q^e$. From the properties of H_e, we see that:

1. $H_e(tZ_1, \ldots, tZ_e) = t^{2^{e-1}} H_e(Z_1, \ldots, Z_e)$ for all $t \in \mathbb{F}_q$. Since $t \mapsto t^{2^{e-1}}$ is a bijection on $\mathbb{F}_q$, we get that all non-zero elements of $\mathbb{F}_q$ have equal probability under ν.
2. $\Pr_{\alpha_1, \ldots, \alpha_e \in \mathbb{F}_q}[H_e(\alpha_1, \ldots, \alpha_e) = 0] \leq \frac{2^e}{q}$. This is by the Schwartz–Zippel lemma. Thus 0 has probability at most $\frac{2^e}{q}$ under ν.

Putting these together, we get that

$$\Pr_{\alpha_1, \ldots, \alpha_e, \alpha'_1, \ldots, \alpha'_e \in \mathbb{F}_q} [H_e(\alpha_1, \ldots, \alpha_e) = H_e(\alpha'_1, \ldots, \alpha'_e)] \leq \frac{1}{q-1} + \frac{2^{2e}}{q^2} = \frac{1}{q} + O\left(\frac{2^{2e}}{q^2}\right).$$

Finally, using the fact that $\Pr[I] \geq 1 - O\left(\frac{2^{2e}}{q}\right)$, we get that:

$$\Pr_{\alpha_1, \ldots, \alpha_e, \alpha'_1, \ldots, \alpha'_e \in \mathbb{F}_q} [H_e(\alpha_1, \ldots, \alpha_e) = H_e(\alpha'_1, \ldots, \alpha'_e) \mid I] \leq \frac{\frac{1}{q} + O\left(\frac{2^{2e}}{q^2}\right)}{1 - O\left(\frac{2^{2e}}{q}\right)} \leq \frac{1}{q} + O\left(\frac{2^{2e}}{q^2}\right).$$

This completes the proof of Lemma 5.6. □

6 Limitations on the proximity gaps at the list-decoding radius

In this section we prove Theorem 1.9.

6.1 Structural lemma: Many values at a random point

The basic phenomenon that we will use is that any large collection of far-apart functions take many different values at a random point. Namely, at a random point α, there are many different β for which some function in the collection has $f(\alpha) = \beta$.

Formally, let $\mathcal{L}$ be a set of functions from S to $\mathbb{F}_q$. For $\alpha \in S$, define:

$$\mathcal{L}(\alpha) = \{h(\alpha) \mid h \in \mathcal{L}\}.$$

Lemma 6.1. *Let $\mathcal{L}$ be a set of functions from S to $\mathbb{F}_q$. Suppose any two elements of $\mathcal{L}$ have at most A agreements. Then:*

$$\mathbf{E}_{\alpha \in S}\left[|\mathcal{L}(\alpha)|\right] \geq \frac{1}{2}\min\left(|\mathcal{L}|, \frac{|S|}{A}\right).$$

Proof. Let $L = |\mathcal{L}|$. For $\alpha \in S$, $y \in \mathbb{F}_q$, define

$$W(\alpha, y) = \frac{1}{L} \cdot |\{h \in \mathcal{L} \mid h(\alpha) = y\}|.$$

Then $W(\alpha, \cdot)$ is the probability distribution of $h(\alpha)$ when h is chosen uniformly at random from $\mathcal{L}$.

Observe that $\mathcal{L}(\alpha) = \{y \mid W(\alpha, y) > 0\}$ is the support of the distribution $W(\alpha, \cdot)$.

We bound the size of $\mathcal{L}(\alpha)$ from below using the second moment by Cauchy–Schwarz:

$$|\mathcal{L}(\alpha)| \geq \frac{\left(\sum_y W(\alpha, y)\right)^2}{\sum_y W(\alpha, y)^2} = \frac{1}{\sum_y W(\alpha, y)^2}.$$

By definition,

$$\sum_y W(\alpha, y)^2 = \frac{1}{L^2} \cdot |\{(h_1, h_2) \in \mathcal{L} \mid h_1(\alpha) = h_2(\alpha)\}|.$$

By linearity of expectation,

$$\begin{aligned}
\mathbf{E}_{\alpha \in S}\left[\sum_y W(\alpha, y)^2\right] &= \frac{1}{L^2} \cdot \mathbf{E}_\alpha\left[|\{(h_1, h_2) \in \mathcal{L} \mid h_1(\alpha) = h_2(\alpha)\}|\right] \\
&= \frac{1}{L^2} \cdot \left(L + \sum_{h_1 \neq h_2} \mathsf{agree}(h_1, h_2)\right) \\
&\leq \frac{1}{L^2} \cdot \left(L + L \cdot (L-1) \cdot \frac{A}{|S|}\right).
\end{aligned}$$

Thus,

$$\mathbf{E}_\alpha\left[|\mathcal{L}(\alpha)\right] \geq \mathbf{E}_\alpha\left[\frac{1}{\sum_y W(\alpha, y)^2}\right]$$

$$\geq \frac{1}{\mathbf{E}_\alpha\left[\sum_y W(\alpha,y)^2\right]} \qquad \text{since } \mathbf{E}\left[\tfrac{1}{X}\right] \geq \tfrac{1}{\mathbf{E}[X]} \text{ for nonnegative r.v.s } X$$
$$\geq \frac{L^2}{L + L(L-1)\frac{A}{|S|}}$$
$$\geq \frac{1}{2} \cdot \frac{2}{\frac{1}{L} + \frac{A}{|S|}} \geq \frac{1}{2}\min\left(|\mathcal{L}|, \frac{|S|}{A}\right),$$

where the last inequality is just the fact that the harmonic mean of L, $\frac{|S|}{A}$ is at least the minimal value. This completes the proof. □

In general, the term $\frac{|S|}{A}$ cannot be improved (for example, if $\mathcal{L}$ consists of perfect A'th powers of degree 1 polynomials evaluated on $S = \mathbb{F}_q$, and A divides $q-1$).

6.2 Proximity gaps stop at the list decoding radius

We now prove Theorem 1.9:

Theorem 1.9. *Let $\mathcal{C}$ be the Reed–Solomon code $\mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, with $|\mathcal{D}| = n$ and $k = (1-\delta)n$. Let $\gamma = \mathsf{LDR}_{\mathbb{F}_q,\mathcal{D},q}(\delta) + \frac{2}{n}$. Then there exist functions $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:*

$$|\{z \in \mathbb{F}_q \mid \Delta(f + z \cdot g, \mathcal{C}) \leq \gamma\}| \geq \frac{q}{2n},$$

but with $\Delta([f,g], \mathcal{C}^2) \geq \delta - \frac{1}{n}$.

Since $\mathsf{LDR}_{\mathbb{F}_q,\mathcal{D},q}(\delta) = \gamma - \frac{2}{n}$, we know that there is some Hamming ball of radius $\gamma - \frac{1}{n}$ which contains $> q$ elements of $\mathcal{C}$. Let $c : \mathcal{D} \to \mathbb{F}_q$ be the center of this ball, and let $\mathcal{L}$ be the set of polynomials $H(X) \in \mathbb{F}_q[X]$ of degree at most k with $\Delta(c, H) \leq \gamma - \frac{1}{n}$. We know that $|\mathcal{L}| \geq q$.

We will now view $\mathcal{L}$ as a set of functions from $\mathbb{F}_q \to \mathbb{F}_q$ ($\mathcal{L}$ was defined based only on the values in $\mathcal{D}$). By Lemma 6.1 applied with domain $\mathbb{F}_q$, we get that there exists $\alpha \in \mathbb{F}_q$, for which the set of values

$$\mathcal{L}(\alpha) = \{H(\alpha) : H(X) \in \mathcal{L}\}$$

has size at least $\min(\frac{q}{2}, \frac{q}{2k}) \geq \frac{q}{2n}$. Define $f : \mathcal{D} \to \mathbb{F}_q$ and $g : \mathcal{D} \to \mathbb{F}_q$ by

$$f(x) = \frac{c(x)}{x - \alpha},$$
$$g(x) = \frac{-1}{x - \alpha},$$

for each $x \in \mathcal{D}$. Then the function $f + z \cdot g : \mathcal{D} \to \mathbb{F}_q$ satisfies

$$(f + z \cdot g)(x) = \frac{c(x) - z}{x - \alpha}.$$

Claim 6.2. *If $z \in \mathcal{L}(\alpha)$, then there exists a polynomial $P(X)$ of degree at most $k-1$ with:*

$$\Delta(f + z \cdot g, P) \leq \gamma.$$

Proof. Let $z \in \mathcal{L}(\alpha)$, and let $H(X) \in \mathcal{L}$ be a polynomial in $\mathcal{L}$ of degree $\leq k$ with $H(\alpha) = z$. We have that the quotient $H_{\alpha,z}(X) = \frac{H(X)-z}{X-\alpha}$ is in fact a polynomial of degree $\leq k-1$. For any $x \in \mathcal{D} \setminus \{\alpha\}$ with $H(x) = c(x)$, we have

$$H_{\alpha,z}(x) = \frac{H(x) - z}{x - \alpha} = \frac{c(x) - z}{x - \alpha} = f(x) + z \cdot g(x).$$

Thus $\Delta(f + z \cdot g, H_{\alpha,z}) \leq \Delta(H, c) + \frac{1}{n} \leq \gamma$, which proves the claim. □

Finally, we note that $\Delta(g, \mathcal{C}) \geq \delta - \frac{1}{n}$. Indeed, if $P(X) \in \mathbb{F}_q[X]$ is of degree $\leq k$, then whenever $x \in \mathcal{D} \setminus \{\alpha\}$ satisfies $g(x) = P(x)$, we also have:

$$P(x) \cdot (x - \alpha) + 1 = 0.$$

Since the polynomial $P(X) \cdot (X - \alpha) + 1$ has degree at most $k + 1$, there can be at most $k + 1$ such x. Thus $\Delta(g, P) \geq 1 - \frac{k+1}{n} = \delta - \frac{1}{n}$, and so $\Delta([f, g], \mathcal{C}^2) \geq \delta - \frac{1}{n}$, as desired.

7 Limits to proximity gaps over prime fields

In this section, we develop some machinery to show limitations to proximity gaps over prime fields. We use this machinery to prove Theorem 1.15 and Theorem 1.16, which exhibit proximity gap thresholds and $\delta/3$ and $\delta/2$ in the regime $\delta = o(1)$. We also show limitations on proximity gaps over prime fields in the setting $\delta = \Omega(1)$ for primes with a certain special number theoretic property. For specific primes of interest, this hypothesis could be checked experimentally, and it seems plausible to us that there are infinitely many such primes (see Conjecture 1.12).

We begin with a general construction based on multiplicative subgroups G of $\mathbb{F}_q$, which we will instantiate with a good choice of $E \subseteq G$. For the sake of a wider range of evaluation domains, we directly state construction lifted to a larger subgroup $H \supseteq G$.

Theorem 7.1. *Let $G \subseteq H \subseteq \mathbb{F}_q^*$ be multiplicative subgroups, and denote $\Phi : H \to G$ the (onto) map sending $x \mapsto x^c$, where $c = |H|/|G|$. Given a subset $E \subseteq G$ and integers $\ell, a \geq 1$, so that*

$$|E^{(+\ell)}| = \left| \left\{ \sum_{i=1}^{\ell} e_i \mid e_1, \ldots, e_\ell \in E \text{ are distinct} \right\} \right| \geq a,$$

let $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, n - (\ell + 2)c]$ be the Reed–Solomon code with evaluation domain $\mathcal{D} = \Phi^{-1}(E)$ of size n, and minimum distance $\delta = \frac{(\ell+2)c}{n}$. Then, for $\gamma = \frac{\ell c}{n} = \frac{\ell}{\ell+2} \cdot \delta$, there exist $f, g : \mathcal{D} \to \mathbb{F}_q$ such that

$$|\{z \in \mathbb{F}_q \mid \Delta(f + zg, \mathcal{C}) \leq \gamma\}| \geq a,$$

and yet $\Delta([f, g], \mathcal{C}^2) \geq \frac{(\ell+1)c}{n} = \frac{\ell+1}{\ell+2}\delta = \frac{\ell+1}{\ell}\gamma$.

Proof. Take $f, g : \mathcal{D} \to \mathbb{F}_q$ defined as

$$f(x) = x^{n-\ell \cdot c},$$
$$g(x) = x^{n-(\ell+1) \cdot c}.$$

Since g is the evaluation of a polynomial of degree $n - (\ell + 1)c$, we have $\Delta(g, \mathcal{C}) \geq \frac{(\ell+1)c}{n} = \frac{\ell+1}{\ell}\gamma$, and therefore

$$\Delta([f, g], \mathcal{C}^2) \geq \frac{\ell + 1}{\ell}\gamma.$$

We want to show that for at least a choices of $z \in \mathbb{F}_q$, there is a polynomial $P_z(X) \in \mathbb{F}_q[X]$ of degree at most $n - (\ell + 2) \cdot c$ such that $f + zg - P_z$ has $n - \ell \cdot c$ zero evaluations in $\mathcal{D}$. This will hold if the polynomial

$$X^{n-\ell \cdot c} + zX^{(n-\ell+1)c} - P_z(X)$$

is the vanishing polynomial of some subset of $\mathcal{D}$. Let $s = \sum_{\alpha \in E} \alpha$. For every $z \in \mathbb{F}_q$ of the form

$$z = s - \sum_{i=1}^{\ell} e_i$$

with $e_1, \ldots, e_\ell$ distinct elements of E, consider the vanishing polynomial of $S = \Phi^{-1}(S')$, where $S' = E \setminus \{e_1, \ldots, e_\ell\}$. The size of S' is $\frac{n}{c} - \ell$, and

$$P_S(X) = \prod_{\alpha \in S'} (X^c - \alpha).$$

Expanding:

$$\begin{aligned} P_S(X) &= X^{n-\ell c} - \left(\sum_{\alpha \in S'} \alpha\right) X^{n-(\ell+1)c} + \Big(\text{Some polynomial of } \deg \leq n - (\ell+2)c\Big) \\ &= X^{n-\ell c} + zX^{n-(\ell+1)c} + \Big(\text{Some polynomial of } \deg \leq n - (\ell+2)c\Big). \end{aligned}$$

Thus $P_S(X)$ has the desired form. By our hypothesis on $\mathcal{D}$, there are at least a choices of z for which $f + zg$ has agreement $n - \ell c$ with some $P_z(X) \in \mathcal{C}$. This proves the theorem. □

7.1 Proof of Theorem 1.15

Recall the theorem statement:

Theorem 1.15. *Let $c > 0$ be an integer. There exist infinitely many q, RS codes $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ over $\mathbb{F}_q$, domain $\mathcal{D} = \mathbb{F}_q$ with $n = |\mathcal{D}| = q$, $k = n - c$ (so that the relative distance δ of $\mathcal{C}$ equals $\frac{c}{n}$), and functions $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:*

$$|\{z \in \mathbb{F}_q \mid \Delta(f + zg, \mathcal{C}) \leq \delta/3\}| \geq \frac{q-1}{c} = \frac{n}{c} = O\left(\frac{1}{\delta}\right),$$

and yet $\Delta([f, g], \mathcal{C}) \geq 2\delta/3$.

Proof. Fix c. Let q be a prime with $q - 1$ divisible by c. (By Dirichlet's theorem on primes in arithmetic progressions, there are infinitely many such q.) Instantiating Theorem 7.1 with $\ell = 1$, $E = G$, $H = \mathbb{F}_q^*$, and $a = |E| = |G| = \frac{q-1}{c}$, exactly gives us the content of Theorem 1.15: examples with constant absolute distance showing that proximity gaps behavior for $\gamma < \delta/3$ changes drastically at $\gamma = \delta/3$. □

To state the difference starkly, when $\delta = O(1/n)$ and $\gamma = \delta/3$, even when a is as large as $\Omega(1/\delta) = \Omega(n)$, we must have distance loss $\varepsilon^* \geq \gamma$. While if $\delta = O(1/n)$ and $\gamma < \delta/3$, already when $a = \gamma n + 1 = O(1)$, we have $\varepsilon^* = 0$.

Observe that the evaluation domain $\mathcal{D}$ of these examples equals $\mathbb{F}_q^*$, and thus has $n = q - 1$.

7.2 Proof of Theorem 1.16

Recall the theorem statement:

Theorem 1.16. *Let $0 < \epsilon < 1/4$ be fixed. There are infinitely many primes q, RS codes $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$ over $\mathbb{F}_q$, domain $\mathcal{D}$ with $n = |\mathcal{D}| = O(q^{0.5+2\epsilon})$, $k = n - \Theta(n^\epsilon)$ (so that the relative distance is $\delta = \Theta(n^{-(1-\epsilon)})$), and functions $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:*

$$|\{z \in \mathbb{F}_q \mid \Delta(f + zg, \mathcal{C}) \leq \delta/2\}| \geq q - 1 = \Omega(n^{2/(1+4\epsilon)}) = \Omega(n^{2-8\epsilon}),$$

and yet $\Delta([f, g], \mathcal{C}^2) \geq 3\delta/4$.

The theorem is obtained by instantiating Theorem 7.1 with $\ell = 2$ and a suitable choice of E, supported by the following lemma.

Lemma 7.2. *Let θ, β be constants with $\beta < 1/4$, $\theta > 1/2$ and $\theta + \beta < 1$. Suppose $\mathbb{F}_q$ is a prime field, and $G \subseteq \mathbb{F}_q^*$ is a multiplicative subgroup of size $\frac{q-1}{c}$, where $c \in \Theta(q^\beta)$. Then there exists a subset E of G of size $O(q^\theta)$ such that*

$$|\{e_1 + e_2 : e_1, e_2 \in E \text{ distinct }\}| \geq q - 1.$$

Proof. We use the probabilistic method.

Set $p = \frac{q^\theta}{|G|}$ (this is at most 1 since $\theta + \beta < 1$). Pick a random set E in Bernoulli manner, by independently deciding for each element in G whether to be included (probability p) or not (probability $1 - p$). We will show that with overwhelming probability the sampled set E has the desired property.

The expected size of E is $p|G|$, thus by the Chernoff bound, $|E| \leq 2p|G| = 2q^\theta$ with probability at least

$$1 - \exp(-p|G|) = 1 - o(1).$$

In this case $|E| = O(q^\theta)$, as required.

Let us investigate the second property. By the finite field Waring theorem (e.g., see [LN96]) every nonzero element of $\mathbb{F}_q$ can be written as $x + y$ for $(x, y) \in G^2$ in $|G|^2/q \pm O(\sqrt{q}) = \Omega(q^{1-2\beta})$ ways. Moreover, asking for x, y to be distinct can change this quantity by at most 1. Fix a nonzero $\alpha \in G$. The probability that none of the $(x, y) \in G^2$ with $x + y = \alpha$ are such that $\{x, y\} \subseteq E$ is at most:

$$(1 - p^2)^{\Omega(q^{1-2\beta})} = e^{-\Omega(p^2 q^{1-2\beta})},$$

since the events $\{x, y\} \subseteq E$ for different unordered pairs $\{x, y\} \subseteq G$ with $x + y = \alpha$ are independent. Taking a union bound over all $\alpha \in \mathbb{F}_q^*$, the probability that there exists some $\alpha \in \mathbb{F}_q^*$ which cannot be written as $e_1 + e_2$ for some distinct $e_1, e_2 \in E$ is at most:

$$q \cdot e^{-\Omega(p^2 q^{1-2\beta})} = q \cdot e^{-q^{2\theta-1}} \leq o(1),$$

where the last inequality holds because $\theta > 1/2$.

Taking the intersection of the two events of large probability yields the claim of the lemma. □

Proof of Theorem 1.16. Let $\epsilon \in (0, 1/4)$ be arbitrary. Choose constants $\beta = \epsilon$ and $\theta = 1/2 + \epsilon$.

By a standard application of the Bombieri–Vinogradov theorem [Dav00], there are infinitely many primes q for which $q - 1$ has a factor c in the range $(q^\beta, 2q^\beta)$. (In other words, we seek infinitely many integers c such that the arithmetic progression $\{n \mid n \equiv 1 \mod c\}$ contains a prime between $(c/2)^{1/\beta}$ and $c^{1/\beta}$. It is known that the Generalized Riemann Hypothesis implies this for all large enough c.)

Take such a q. Let G be the set of c-th powers in $\mathbb{F}_q^*$, and let $E \subseteq G$ be a subset of G given by the above lemma. Using $\ell = 2$, c, E as above, and applying Theorem 7.1, we get proximity loss $\varepsilon^* = \frac{1}{4}\delta$ at radius $\gamma = \delta/2$ for the Reed–Solomon code $\mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, with domain $\mathcal{D}$ satisfying:

$$n = |\mathcal{D}| = c \cdot |E| = O(q^\beta \cdot q^\theta) = O(q^{1/2+2\epsilon}),$$

$k = n - 4c = n - \Theta(q^\beta) = n - \Theta(n^\epsilon)$, and the number of exceptional z's:

$$a = q - 1.$$

This completes the proof of Theorem 1.16. □

7.3 Limits on proximity gaps over prime fields in the $\delta = \Omega(1)$ regime?

We now prove Theorem 1.13:

Theorem 1.13. *Suppose (q, a, b) is admissible with b even. Let G be the corresponding subgroup of $\mathbb{F}_q^*$ of cardinality b, and let $H \subseteq \mathbb{F}_q^*$ be any multiplicative subgroup containing G. Consider the Reed–Solomon code $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, with $\mathcal{D} = H$, $n = |\mathcal{D}|$, $k = \left(\frac{1}{2} - \frac{2}{b}\right)n$, and relative distance $\delta = \frac{1}{2} + \frac{2}{b}$.*

Then there exist functions $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:

$$\left| \left\{ z \in \mathbb{F}_q \mid \Delta(f + z \cdot g, \mathcal{C}) \leq \delta - \frac{2}{b} \right\} \right| \geq a,$$

but $\Delta([f, g], \mathcal{C}^2) \geq \delta - \frac{1}{b}$.

Proof. As mentioned in Section 1.4.3, this combined with Conjecture 1.12 provides infinitely many n for which there are Reed–Solomon codes over prime fields $\mathbb{F}_q$ with length $n = q - 1$ and relative distance $\delta = 1/2$, such that any proximity gaps result with $\gamma = \delta - \Theta(1/\log n)$ and $a \geq n$ has a proximity loss of $\Theta(1/\log n)$.

Let (q, a, b) be admissible. Let G be the subgroup of $\mathbb{F}_q^*$ of size b. Let $\mathcal{D}$ be a multiplicative subgroup of $\mathbb{F}_q^*$ containing G. Set $c = |\mathcal{D}|/b$. Then the c-th power map $\Phi : \mathcal{D} \to G$ is c-to-1 and onto. We apply Theorem 7.1 with this c, $G = E$, $H = \mathcal{D}$, and $\ell = b/2$. By our admissibility assumption, we have $|E^{(+\ell)}| \geq a$. We get that for parameters $\delta = \frac{(\ell+2)c}{n} = \frac{1}{2} + \frac{2}{b}$, $\gamma = \frac{\ell c}{n} = \frac{1}{2}$ and code $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, n - (\ell + 2)c]$ (which has distance δ), there are functions $f, g : \mathcal{D} \to \mathbb{F}_q$ such that:

$$|\{z \in \mathbb{F}_q \mid \Delta(f + z \cdot g, \mathcal{C}) \leq \gamma\}| \geq a,$$

and yet $\Delta([f, g], \mathcal{C}^2) \geq \gamma + \frac{1}{b}$. This completes the proof of Theorem 1.13. □

Remark 7.3. Let $q = 2^p - 1$ be a Mersenne prime, with p an odd prime. Then $(q, q, 2p)$ is admissible. Indeed, let G be the subgroup of $\mathbb{F}_q^*$ generated by -2. Then $G = \{\pm 1, \pm 2, \pm 2^2, \ldots, \pm 2^{p-1}\}$, and $|G| = 2p$. For any $u \in \mathbb{F}_q$, consider the binary representation of $u/2 = \sum_{i=0}^{p-1} b_i 2^i$ with $b_i \in \{0, 1\}$. Note that for each i, $2b_i - 1 = \pm 1$ and consider the p distinct elements $e_i = (2b_i - 1) \cdot 2^i \in G$. We get

$$\sum_{i=0}^{p-1} e_i = 2 \sum_{i=0}^{p-1} b_i 2^i - \sum_{i=0}^{p-1} 2^i \equiv u \pmod{q},$$

which shows that $u \in G^{(+p)}$. This proves the claim about admissibility and Mersenne primes.

8 Attacks on STARKs near the list decoding radius

In this section, we prove Theorem 1.17:

Theorem 1.17. *Consider the IOP protocol for the CYCLE-SUM constraint satisfaction problem given by the DEEP-ALI reduction, using the Reed–Solomon code $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, where:*

- *$\mathcal{D}$ is a union of t cosets of G*
- *$|\mathcal{D}| = a \cdot t = n$ and $k = a = \frac{1}{t}n$.*
- *$\delta = 1 - \frac{1}{t}$ is the distance of $\mathcal{C}$.*

Then there is a prover strategy that does not make the verifier reject, and produces $(h_1, \ldots, h_c)$ such that:

$$\Pr\left[\Delta([h_1, \ldots, h_c], \mathcal{C}^c) \leq \frac{1+\gamma_q}{2}\right] \geq \Omega(1/n),$$

where $\gamma_q = \mathsf{LDR}_{\mathbb{F}_q, \mathcal{D}, q}(\delta) + \frac{1}{n}$.

We begin by recalling the basic STARKs IOP protocol, instantiated for the CYCLE-SUM CSP.

8.1 The basic STARK

The basic STARK IOPP, from [BSBHR18], is as follows.

1. We first pick an evaluation domain $\mathcal{D} \subseteq \mathbb{F}_q$ which is a union of t cosets of G. So $n = |\mathcal{D}| = a \cdot t$.
2. We ask the prover to write down a function $f : \mathcal{D} \to \mathbb{F}_q$ (which is supposed to be evaluations of a low degree polynomial $F(X)$ interpolating the satisfying assignment $f : G \to \mathbb{F}_q$).

The key observation: if F is truly as intended, then the polynomial $F(gX) - F(X) - 1$ will vanish on G, and thus be divisible by the vanishing polynomial of G, namely $X^a - 1$. Thus we expect $B^f : \mathcal{D} \to \mathbb{F}_q$ to be a low degree polynomial, where for a general function $h : \mathcal{D} \to \mathbb{F}_q$, the function $B^h : \mathcal{D} \to \mathbb{F}_q$ is defined to be:

$$B^h(x) = \frac{h(gx) - h(x) - 1}{x^a - 1},$$

where f and B^f are the output of this first phase of the protocol. Observe that there was no interaction in this protocol after the prover wrote down f, and that oracle access to B^f can be simulated from oracle access to f. The next phase is:

3. Check the low-degreeness of f and B^f by running a low-degree test such as FRI, to confirm that they are both of degree at most $k = a$.

This is the problem of checking proximity to $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$. Let $\rho = k/n$, and note that the distance δ of $\mathcal{C}$ equals $1 - \rho = 1 - \frac{1}{t}$.

8.1.1 Attack on the basic STARK

We now show that it is possible to choose f so that both f and B^f have distance $\leq \delta/2$ from $\mathcal{C}$, and furthermore it can be correlated: $\Delta([f, B^f], \mathcal{C}^2) \leq \delta/2$. (This is the unique decoding radius of $\mathcal{C}$.)

The idea is that the value of B^f at a single element of $\mathcal{D}$ can be set to any desired value by modifying a single value of f. Thus taking half the values of f consistent with some low degree polynomial $F(X)$ and the other half of the values of f so that B^f is consistent with some low degree polynomial $C(X)$, we get that both f and B^f have agreement at least $1/2$ with low degree polynomials.

Explicitly, let $F(X), C(X) \in \mathbb{F}_q[X]$ be arbitrary polynomials of degree at most k. For each coset $yG \subseteq D$, define:

$$\begin{aligned} f(y \cdot g^{2i}) &= F(y \cdot g^{2i}) \\ f(y \cdot g^{2i+1}) &= C(y \cdot g^{2i}) \left((y \cdot g^{2i})^a - 1\right) + 1 + f(y \cdot g^{2i}). \end{aligned}$$

This ensures that for all i,

$$f(y \cdot g^{2i}) = F(y \cdot g^{2i}),$$

$$B^f(y \cdot g^{2i}) = C(y \cdot g^{2i}).$$

This gives us a correlated agreement of at least 1/2.

By choosing $F(X)$ arbitrarily and $C(X)$ to be the low degree extension of B^F on some $k = \rho \cdot n$ points, we can get some extra agreements. It is easy to check that this lets us increase the agreement to $\frac{1+\rho}{2} = \frac{\delta}{2}$.

8.2 STARK with DEEP queries

Now we consider the STARK protocol with DEEP[6] queries [BGKS20]. This is the version of the protocol with the best known soundness, and our main result about attacking STARKs is for this.

Again we have an evaluation domain $\mathcal{D} \subseteq \mathbb{F}_q$ which is a union of t cosets of G. So $n = |\mathcal{D}| = a \cdot t$.

We ask the prover to write down a function $f : \mathcal{D} \to \mathbb{F}_q$ (which is supposed to be evaluations of a low degree polynomial $F(X)$ interpolating the satisfying assignment $f : G \to \mathbb{F}_q$). Thus f is supposed to be an element of $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, where $k = a$. Let $\rho = k/n$ and $\delta = 1 - \rho$, the distance of $\mathcal{C}$.

The main technical notion we need is the quotient of a function $h : \mathcal{D} \to \mathbb{F}_q$ by some *claims*: "$h(\alpha_1) = \beta_1$", ..., "$h(\alpha_r) = \beta_r$". This is a function $h' : \mathcal{D} \to \mathbb{F}_q$ given by:

$$h'(x) = \frac{h(x) - V(x)}{Z(x)}$$

where:

- $V(X) \in \mathbb{F}_q[X]$ is the unique polynomial of degree at most $r-1$ with $V(\alpha_i) = \beta_i$ for each i,
- $Z(X)$ is the vanishing polynomial of $\{\alpha_1, \ldots, \alpha_r\}$.

The definition has the following property: suppose h is the evaluation table of some polynomial $H(X)$ of degree at most d such that $H(\alpha_i) = \beta_i$ – then h' is the evaluation table of some polynomial $H'(X)$ of degree at most $d - t$.

In this protocol, after the prover has committed to values of f on $\mathcal{D}$, the verifier then asks for values of (the polynomial underlying) f at values $\alpha, g\alpha$ outside the domain $\mathcal{D}$, for a uniformly random $\alpha \in \mathbb{F}_q$. If the prover responds with values u, v, the verifier then *quotients* f by the claims "$f(\alpha) = u$" and "$f(g\alpha) = v$", and quotients B^f by the claim "$B^f(\alpha) = \frac{v-u-1}{\alpha^a - 1}$". These quotiented functions are then checked for low-degreeness using a low degree test like FRI.

8.2.1 Attack on the DEEP STARK

The main lemma from [BGKS20] about quotienting is below.

Lemma 8.1. *Given a function $h : \mathcal{D} \to \mathbb{F}_q$, and some claimed values $h(\alpha_i) = \beta_i$ for some $\alpha_1, \ldots, \alpha_r \in \mathbb{F}_q \setminus \mathcal{D}$, let h' be the quotient of h by these claimed values. Then the following are equivalent:*

- *There exists $H'(X)$ of degree at most $d - r$ such that $\Delta_{\mathcal{D}}(h', H') \leq \delta$.*
- *There exists $H(X)$ of degree at most d such that $\Delta_{\mathcal{D}}(h, H) \leq \delta$ and $H(\alpha_i) = \beta_i$ for each i.*

Let us first see why the previous attack does not work.

[6] DEEP is short for Domain Extension for Eliminating Pretenders, and the main idea is to make queries outside the domain of definition of what the prover has sent.

Suppose we (the attacking provers) chose f so that f is close to F and B^f is close to C. Now when the verifier asks us for the values of f at α and $g\alpha$, we can either answer consistent with F or not. If we do not answer consistent with F, then by the quotienting lemma there will not exist a low degree polynomial close to the quotient of f, and the low degree test will catch us. If we do answer consistent with F, then we will quotient B^f by the claim "$B^f(\alpha) = \frac{F(g\alpha)-F(\alpha)-1}{\alpha^a-1}$", which is very unlikely to be consistent with $C(\alpha)$ (since $C(X)$ and $\frac{F(gX)-F(X)-1}{X^a-1}$ are distinct low-degree rational functions). This means, by the quotienting lemma again, that the quotiented B^f will be very far from low degree polynomials, and the low degree test will catch us.

Thus, the previous attack does not work as is. This is for a good reason, since the analysis in [BGKS20] proves soundness upto the list-decoding radius of Reed–Solomon codes.

Instead we need to use bad list-decoding configurations, which only exist for radii beyond the list-decoding radius of Reed–Solomon codes.

Let $\gamma = \mathsf{LDR}_{\mathbb{F}_q,\mathcal{D},q}(\delta) + \frac{1}{n}$. By definition, there exists a function $c : \mathcal{D} \to \mathbb{F}_q$ such that

$$\mathcal{L} = \{H(X) \in \mathbb{F}_q[X] \mid \deg(H) \leq k, \Delta(c, H) \leq \gamma\} \tag{16}$$

has $|\mathcal{L}| \geq q$.

We know that $\gamma \geq 1 - \sqrt{1-\delta}$ and $\gamma \leq \delta$.

Theorem 8.2. *There is a function f which the prover can write down such that probability at least $\Omega(\frac{1}{n})$ over the choice of $\alpha \in \mathbb{F}_q$, the prover can answer values for $f(\alpha)$ and $f(g\alpha)$ so that the quotiented versions of f and B^f, denote h_1 and h_2, satisfy:*

$$\Delta([h_1, h_2], \mathcal{C}^2) \leq \frac{1+\gamma}{2}.$$

Proof. The idea is to use the bad list-decoding center c in place of the polynomial $C(X)$ in the original attack on the STARK without DEEP.

Before we describe the prover's strategy, we first some preliminary adjustments to c and $\mathcal{L}$.

Suppose Y is a set of representatives for the cosets of G that make up $\mathcal{D}$. (Thus $\mathcal{D}$ is a disjoint union of yG for $y \in Y$).

Let $G^2 = \{g^{2i} \mid i \in \mathbb{N}\}$. Define $\mathcal{D}^* = \bigcup_{y\in Y}(yG^2)$. Note that $|\mathcal{D}^*| = |\mathcal{D}|/2$.

Each $H(X)$ in $\mathcal{L}$ has $\Delta_{\mathcal{D}}(c, H) \leq \gamma$. This means that either $\Delta_{\mathcal{D}^*}(c, H) \leq \gamma$ or $\Delta_{\mathcal{D}\setminus\mathcal{D}^*}(c, H) \leq \gamma$. In the latter case, defining $c'(x) = c(gx)$ and $H'(X) = H(gX)$, we have $\Delta_{\mathcal{D}^*}(c', H') \leq \gamma$.

Thus, by replacing c with c' if needed, we may assume that there is a large number of polynomials $H(X)$ that are close to c on $\mathcal{D}^*$. Concretely, define:

$$\mathcal{L}^* = \{H(X) \in \mathbb{F}_q[X] \mid \deg(H) \leq k, \Delta_{\mathcal{D}^*}(c, H) \leq \gamma\},$$

and we have:

$$|\mathcal{L}^*| \geq q/2.$$

Let $F(X)$ be a low degree polynomial that will be chosen carefully later. For each coset $yG \subseteq D$, define as before for each i:

$$\begin{aligned} f(y \cdot g^{2i}) &= F(y \cdot g^{2i}) \\ f(y \cdot g^{2i+1}) &= c(y \cdot g^{2i}) \left((y \cdot g^{2i})^a - 1\right) + 1 + f(y \cdot g^{2i}). \end{aligned}$$

This will be the f that the prover writes down.

Our definition of f ensures that for all $x \in \mathcal{D}^*$:

$$f(x) = F(x)$$
$$B^f(x) = c(x)$$

This gives us perfect agreement between $[f, B^f]$ and $[F, c]$ on $\mathcal{D}^*$, (and thus correlated agreement at least $1/2$ between (f, B^f) and (F, c) on $\mathcal{D}$).

Now when the verifier asks for $f(\alpha)$ and $f(g\alpha)$, the prover answers $F(\alpha)$ and $F(g\alpha)$. By the quotienting lemma, this ensures that the quotiented F, which is a low degree polynomial of degree at most $k-2$, has perfect agreement with the quotiented f on $\mathcal{D}^*$, (and agreement $\geq 1/2$ with the quotiented f on $\mathcal{D}$).

What about the quotiented B^f? B^f was quotiented by the claim "$B^f(\alpha) = \frac{F(g\alpha)-F(\alpha)-1}{\alpha^a-1}$". Does there exist a polynomial of degree at most $k-1$ that is close to the quotiented B^f on $\mathcal{D}^*$? By the quotienting lemma, this will be true if and only if there is a polynomial of degree at most k that is close to B^f on $\mathcal{D}^*$ (recall that $B^f = c$ on $\mathcal{D}^*$), such that $H(\alpha) = \frac{F(g\alpha)-F(\alpha)-1}{\alpha^a-1}$.

We can now choose the univariate polynomial $F(X)$; it will be so that the probability (over α) of such an H existing is noticeable. First, a lemma.

Lemma 8.3. *Let $\mathcal{L}$ be the set of polynomials close to c as given by* (16)*.*

Define $\mathcal{L}' = \left\{ \frac{H(X)\cdot(X^a-1)+1}{(g-1)X} : H(X) \in \mathcal{L}^* \right\}$. Then there exists some $\lambda \in \mathbb{F}_q$ such that:

$$\Pr_{\alpha \in \mathbb{F}_q}[\lambda \in \mathcal{L}'(\alpha)] \geq \frac{1}{2}\frac{1}{k+a}.$$

Proof. Observe that any two elements of $\mathcal{L}'$ have agreement at most $k+a$. Lemma 6.1 tells us that the average $\mathcal{L}'(\alpha)$ is big:

$$\mathbf{E}_{\alpha \in \mathbb{F}_q}[|\mathcal{L}'(\alpha)|] \geq \frac{1}{2}\frac{q}{k+a}.$$

Thus by the probabilistic method, there exists $\lambda \in \mathbb{F}_q$ that lies in $\frac{1}{2}\frac{1}{k+a}$ fraction of all the $\mathcal{L}'(\alpha)$. This λ satisfies the requirements of the lemma. □

Define $F(X) = \lambda X$. Then $\frac{F(g\alpha)-F(\alpha)-1}{\alpha^a-1} = \frac{\lambda(g-1)\alpha-1}{\alpha^a-1}$. Thus $H(X) \in \mathcal{L}^*$ will have

$$H(\alpha) = \frac{F(g\alpha) - F(\alpha) - 1}{\alpha^a - 1}$$

if and only if:

$$\frac{H(\alpha)\cdot(\alpha^a-1)+1}{(g-1)\cdot\alpha} = \lambda.$$

By choice of λ, such an $H(X) \in \mathcal{L}$ will exist with probability at least $\frac{1}{2}\frac{1}{k+a}$ over the choice of α.

If this happens, we get that $[f, B^f]$ is close to the pair of low degree polynomials $[F, H]$:

$$\Delta_{\mathcal{D}}([f, B^f], [F, H]) \leq \frac{1+\gamma}{2}. \tag{17}$$

Indeed, by our choice of f, we already know that for all $x \in \mathcal{D}^*$,

$$f(x) = F(x).$$

We also know that $H(X) \in \mathcal{L}^*$: thus $\Delta_{\mathcal{D}^*}(H, c) \leq \gamma$. This means that $\Delta_{\mathcal{D}^*}([f, B^f], [F, H]) \leq \gamma$. Taking into account the fact that $|\mathcal{D}^*| = |\mathcal{D}|/2$, we get Equation (17) above. Thus

$$\Pr\left[\Delta([f, B^f], \mathcal{C}^2) \leq \frac{1+\gamma}{2}\right] \geq \frac{1}{2(k+a)},$$

as desired. □

8.3 Attack on the Ethstark Toy Problem

The Ethstark Toy Problem [Sta23] asked what is the best prover success probability that one could have for the following protocol. We fix a Reed-Solomon code $\mathcal{C} = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k]$, and let $n = |\mathcal{D}|$, and its distance $\delta = 1 - \frac{k}{n}$.

- The prover writes down $f : \mathcal{D} \to \mathbb{F}_q$
- The verifier picks α, β uniformly at random, and sends them to the prover.
- The prover will try to convince the verifier that f is in $\mathcal{C}$, and the polynomial $F(X)$ underlying it has $F(\alpha) = \beta$. Concretely, define the function $h[f, \alpha, \beta] : \mathcal{D} \to \mathbb{F}_q$ by:

$$h^{f,\alpha,\beta}(x) = \frac{f(x) - \beta}{x - \alpha}.$$

  (Access to h can be simulated using access to f).

  The prover and verifier now run the FRI protocol with repetition parameter t to prove that h lies in $\mathcal{C}' = \mathsf{RS}[\mathbb{F}_q, \mathcal{D}, k-1]$.

In [Sta23], it was noted that the best known success probability for the prover was $\frac{1}{q} + \left(\frac{k}{n}\right)^t = \frac{1}{q} + (1-\delta)^t$.

We prove:

Lemma 8.4. *Let $\gamma_q = \mathsf{LDR}_{\mathbb{F}_q,\mathcal{D},q}(\delta) + \frac{1}{n}$. Then there is a function $f : \mathcal{D} \to \mathbb{F}_q$ such that:*

$$\Pr_{\alpha,\beta}[\Delta(h^{f,\alpha,\beta}, \mathcal{C}') \leq \gamma_q] \geq \frac{1}{2n}.$$

Using the fact that the FRI protocol accepts functions that are β-close to $\mathcal{C}'$ with probability $(1-\beta)^t$, this lemma translates into a prover strategy with success probability

$$\frac{1}{2n} + (1 - \gamma_q)^t.$$

If the Reed-Solomon code $\mathcal{C}$ has a small list decoding radius for list size q, then this is potentially a much larger success probability. Examples of (infinite families of) Reed-Solomon codes with $\gamma_q < \delta - \Omega(1)$ are known, for example in [BSKR06].

Proof of Lemma 8.4

We take $f : \mathcal{D} \to \mathbb{F}_q$ to be a function with many nearby codewords of $\mathcal{C}$. Specifically, we take f such that:

$$\mathcal{L} := \{P(X) \in \mathcal{C} \mid \Delta_{\mathcal{D}}(f, P) \leq \gamma_q\}$$

satisfies:

$$|\mathcal{L}| \geq q.$$

Such an f is guaranteed to exist by definition of γ_q.

We now view $\mathcal{L}$ as a collection of functions from $\mathbb{F}_q \to \mathbb{F}_q$, and apply Lemma 6.1 (with $S = \mathbb{F}_q$) to it. We get that for random $\alpha \in \mathbb{F}_q$:

$$\mathbf{E}_{\alpha \in \mathbb{F}_q}[|\mathcal{L}(\alpha)|] \geq \frac{1}{2} \min(q, \frac{q}{k}) \geq \frac{q}{2n}.$$

Thus for uniformly random $\beta \in \mathbb{F}_q$,

$$\Pr_{\alpha,\beta}[\beta \in \mathcal{L}(\alpha)] \geq \frac{1}{2n}.$$

In other words, with probability at least $\frac{1}{2n}$ over the choice of $\alpha, \beta \in \mathbb{F}_q$, there exists some polynomial $F(X) \in \mathcal{L}$ such that $F(\alpha) = \beta$.

If this happens, we show that $h^{f,\alpha,\beta}$ is close to $\mathcal{C}'$. For this, we use the main fact about quotienting: Suppose $F(X)$ is a polynomial such that:

- $\Delta_{\mathcal{D}}(f, F) \leq \gamma_q$
- $F(\alpha) = \beta$

Then defining $H(X)$ to be the degree $\leq (k-1)$ *polynomial* $\frac{F(X)-\beta}{X-\alpha}$, we have:

$$\Delta_{\mathcal{D}}(h^{f,\alpha,\beta}, H) \leq \gamma_q.$$

This completes the proof of Lemma 8.4.

References

[ACFY24] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. STIR: Reed–Solomon proximity testing with fewer queries. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024*, 2024. Full paper: `https://eprint.iacr.org/2024/390`.

[ACFY25] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. WHIR: Reed–Solomon proximity testing with super-fast verification. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology – EUROCRYPT 2025*, 2025. Full paper: `https://eprint.iacr.org/2024/1586`.

[ACY23] Gal Arnon, Alessandro Chiesa, and Eylon Yogev. IOPs with inverse polynomial soundness error. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, 2023. full paper `https://eprint.iacr.org/2023/1062`.

[AGL24] Omar Alrabiah, Venkatesan Guruswami, and Ray Li. Randomly punctured Reed–Solomon codes achieve list-decoding capacity over linear-sized fields. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, 2024.

[AHIV17] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkitasubramaniam. Ligero: Lightweight sublinear arguments without a trusted setup. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, New York, NY, USA, 2017. Full paper: `https://eprint.iacr.org/2022/1608`.

[BBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast Reed–Solomon interactive oracle proofs of proximity. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, 2018. Full paper: `https://eccc.weizmann.ac.il/report/2017/134/`.

[BCI+20] Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty, and Shubhangi Saraf. Proximity gaps for Reed–Solomon codes. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, 2020. Full paper: `https://eprint.iacr.org/2020/654`.

[Ber15] Elwyn R. Berlekamp. Algebraic coding theory. 2015. `https://doi.org/10.1142/9407`.

[BGKS20] Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty, and Shubhangi Saraf. DEEP-FRI: Sampling outside the box improves soundness. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, 2020. Full paper: https://eprint.iacr.org/2019/336.

[BGM23] Joshua Brakensiek, Sivakanth Gopi, and Visu Makam. Generic Reed–Solomon codes achieve list-decoding capacity. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, page 1488–1501, 2023.

[BKS18] Eli Ben-Sasson, Swastik Kopparty, and Shubhangi Saraf. Worst-case to average case reductions for the distance to a code. In *Proceedings of the 33rd Computational Complexity Conference*, CCC '18, 2018. Full paper: https://eccc.weizmann.ac.il/report/2018/090/.

[BSBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity. Cryptology ePrint Archive, Paper 2018/046, 2018. https://eprint.iacr.org/2018/046.

[BSK12] Eli Ben-Sasson and Swastik Kopparty. Affine dispersers from subspace polynomials. *SIAM Journal on Computing*, 41(4), 2012. https://doi.org/10.1137/110826254.

[BSKR06] Eli Ben-Sasson, Swastik Kopparty, and Jaikumar Radhakrishnan. Subspace polynomials and list decoding of Reed–Solomon codes. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 207–216, 2006.

[CS25] Elizabeth Crites and Alistair Stewart. On reed–solomon proximity gaps conjectures. Cryptology ePrint Archive, Paper 2025/2046, 2025. https://eprint.iacr.org/2025/2046.

[Dav00] Harold Davenport. *Multiplicative Number Theory*, volume 74 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, third edition, 2000.

[DG25a] Benjamin E. Diamond and Angus Gruen. On the distribution of the distances of random words. Cryptology ePrint Archive, Paper 2025/2010, 2025. https://eprint.iacr.org/2025/2010.

[DG25b] Benjamin E. Diamond and Angus Gruen. Proximity gaps in interleaved codes. *IACR Communications in Cryptology*, 1(4), 2025. https://cic.iacr.org/p/1/4/8/pdf.

[DP24] Benjamin E. Diamond and Jim Posen. Polylogarithmic proofs for multilinears over binary towers. In *IACR preprint archive 2024/504*, 2024. https://eprint.iacr.org/2024/504.

[GCXK25] Yiwen Gao, Dongliang Cai, Yang Xu, and Haibin Kan. From list-decodability to proximity gaps. Cryptology ePrint Archive, Paper 2025/870, 2025. https://eprint.iacr.org/2025/870.

[GG25] Rohan Goyal and Venkatesan Guruswami. Optimal proximity gaps for subspace-design codes and (random) Reed-Solomon codes, 2025. (Personal communication).

[GGM25] Albert Garreta, Angus Gruen, and Igancio Manzur. Attacking FRI and the STARK toy problem. 2025. (Personal communication).

[GK07] A. A. Glibichuk and S. V. Konyagin. Additive properties of product sets in fields of prime order. 2007. https://arxiv.org/abs/math/0702729.

[GKL24] Yiwen Gao, Haibin Kan, and Yuan Li. Linear proximity gap for linear codes within the 1.5 Johnson bound. Cryptology ePrint Archive, Paper 2024/1810, 2024. https://eprint.iacr.org/2024/1810.

[GLS+24] Zeyu Guo, Ray Li, Chong Shangguan, Itzhak Tamo, and Mary Wootters. Improved list-decodability and list-recoverability of Reed–Solomon codes via tree packings. *SIAM Journal on Computing*, 53(2):389–430, 2024.

[GS99] Venkatesan Guruswami and Madhu Sudan. Improved decoding of Reed–Solomon and algebraic-geometry codes. In *IEEE Trans. on Information Theory*, volume 45(6), 1999.

[GZ23] Zeyu Guo and Zihan Zhang. Randomly punctured Reed–Solomon codes achieve the list decoding capacity over polynomial-size alphabets. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 164–176, 2023.

[Hab24] Ulrich Haböck. Basefold in the list decoding regime. Cryptology ePrint Archive, Paper 2024/1571, 2024. `https://eprint.iacr.org/2024/1571`.

[Hab25] Ulrich Haböck. A note on mutual correlated agreement. 2025. (Personal communication).

[HLP24] Ulrich Haböck, David Levit, and Shahar Papini. Circle STARKs. In *IACR preprint archive*, 2024. `https://eprint.iacr.org/2024/278`.

[JH01] J. Justesen and T. Hoholdt. Bounds on list decoding of MDS codes. *IEEE Transactions on Information Theory*, 47(4):1604–1609, 2001.

[Joh62] Selmer M. Johnson. A new upper bound for error-correcting codes. In *IRE Transactions on Information Theory*, volume 8, pages 203–207, 1962.

[LN96] Rudolf Lidl and Harald Niederreiter. *Finite Fields*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2 edition, 1996.

[MZ25] Dor Minzer and Kai Zhe Zheng. Improved round-by-round soundness IOPs via Reed–Muller codes. Cryptology ePrint Archive, Paper 2025/600, 2025. `https://eprint.iacr.org/2025/600`.

[PS94] Alexander Polishchuk and Daniel A. Spielman. Nearly-linear size holographic proofs. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '94, page 194–203. Association for Computing Machinery, 1994. `https://doi.org/10.1145/195058.195132`.

[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*, STOC '13, page 793–802, 2013.

[RW13] Atri Rudra and Mary Wootters. Every list-decodable code for high noise has abundant near-optimal rate puncturings. *CoRR*, abs/1310.1891, 2013.

[RZ18] Ron Roth and Gilles Zémor. Personal communcation to the authors of [AHIV17]. 2018. Cited in the full version of [AHIV17], `https://eprint.iacr.org/2022/1608`.

[ST20] Chong Shangguan and Itzhak Tamo. Combinatorial list-decoding of Reed–Solomon codes beyond the Johnson radius. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, page 538–551, New York, NY, USA, 2020. Association for Computing Machinery.

[Sta23] StarkWare Team. ethSTARK documentation – version 1.2. In *IACR preprint archive 2021/582*, 2023. `https://eprint.iacr.org/2021/582`.

[Sta25] StarkWare Team. S-two whitepaper. 2025. (Personal communication).

[WB86] Lloyd R. Welch and Elwyn R. Berlekamp. Error correction for algebraic block codes. US Patent 4633470, 1986. `https://patents.google.com/patent/US4633470A`.

[ZCF23] Hadas Zeilberger, Binyi Chen, and Ben Fisch. BaseFold: efficient field-agnostic polynomial commitment schemes from foldable codes. In *IACR preprint archive 2023/1705*, 2023. `https://eprint.iacr.org/2023/1705`.

[Zei24] Hadas Zeilberger. Khatam: Reducing the communication complexity of code-based SNARKs. In *IACR preprint archive 2024/1843*, 2024. `https://eprint.iacr.org/2024/1843`.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il