Affine Combiners for Reed–Solomon Proximity Gaps

Alessandro Chiesa
EPFL
alessandro.chiesa@epfl.ch

Guy Weissenberg
EPFL
guy.weissenberg@epfl.ch

Eylon Yogev
Bar-Ilan University
eylon.yogev@biu.ac.il

Kai Zhe Zheng
Simons Institute for the Theory of Computing
kzzheng@mit.edu

November 13, 2026

Abstract

Reed–Solomon proximity gaps establish soundness guarantees for protocols that test received words through random combinations. For length-n codes over $\mathbb{F}_q$, classical scalar combiners achieve error $O(n/q)$ in the unique-decoding regime, with larger known bounds near the Johnson radius.

We prove smaller error bounds using affine combiners, which combine two received words f, g as $f(x) + (\alpha + \beta x)g(x)$, where α, β are independent uniform elements of $\mathbb{F}_q$. For output degree bound $d = \rho n$, fixed $\rho \in (0, 1)$ and $0 < \eta < 1 - \sqrt{\rho}$, and sufficiently large block length n, we prove mutual correlated agreement (MCA) with one coordinate of agreement loss at threshold $A = \lceil (\sqrt{\rho} + \eta)n \rceil$, with error

$$O_{\rho,\eta}\left(\frac{1}{q} + \frac{n^2}{q^2}\right).$$

The result holds over arbitrary evaluation domains in every finite field, including prime and binary extension fields. The agreement threshold can approach the Johnson threshold, beyond unique decoding.

For $q \geq n^2$, our error is $O_{\rho,\eta}(1/q)$, an asymptotic factor-n separation from scalar combiners: we prove that their worst-case MCA error is $\Omega_{\rho,\eta}(n/q)$ at the same agreement threshold and with the same one-coordinate loss. We also prove that affine combiners incur worst-case error $\Omega_{\rho,\eta}(n/q)$ when no agreement loss is allowed. Thus one coordinate of agreement loss is necessary for our guarantee; this relaxation of standard MCA reduces the guaranteed common agreement fraction by at most $1/n$. We also extend the result to k-variate Reed–Muller codes on subfield grids, obtaining MCA error $O_{\rho,\eta,k}(q^{-k})$ for fixed ρ, η, k, sufficiently large n, and $q \geq n^{k+1}$, with agreement loss $n^{1-1/k}$.

Keywords: Reed–Solomon codes, proximity gaps, affine combiners, mutual correlated agreement, Reed–Muller codes, low-degree testing.

Contents

1 Introduction **1**
1.1 Our results 1
1.2 Related Work and Further work 4

2 Techniques **5**

3 Preliminaries **9**
3.1 Algebraic Geometry 9

4 Proof of the proximity gap for affine combiners **13**
4.1 Interpolation along affine lines 13
4.2 Bounding excluded coefficient pairs 16
4.3 Bounding persistent coefficient pairs 17
4.4 Bounding nonpersistent coefficient pairs 19

5 Adjustments for Equal-degree, linear constraints, and interleaving **24**

6 Separation from scalar combiners and Necessity of Agreement Loss **29**

7 Extension to Reed–Muller codes **32**
7.1 Proof idea 33

References **35**

A Proof of the Reed–Muller extension **38**

1 Introduction

Reed–Solomon proximity testing is a central component of efficient proof systems, including FRI [BBHR18], STIR [ACFY24a], and WHIR [ACFY24b]. Many protocols test a batch of received words by checking a random linear combination of them for proximity to a Reed–Solomon code. A word is close to the code when it agrees with a low-degree polynomial on many evaluation points. The soundness of this approach rests on a *proximity gap* guarantee which is typically a statement of the following form: if at least one of the words in the combination is δ-far from the code, then the combination is still δ'-far from the code except with a small error probability ε. Typically one wants δ' either equal or at least close to δ and ε as small as possible.

Classical results establish such gaps for δ all the way up to the Johnson radius [BCI$^+$20, BCH$^+$25, DKT26], but the error, ε, that they achieve can sometimes be an issue. Indeed, for a length-n Reed–Solomon code over a field $\mathbb{F}_q$, the classical proximity-gap bounds are n/q, even in the unique-decoding regime [BCI$^+$20, Theorem 1.2], and this dependence on the block length costs $\log_2 n$ bits of soundness when used in proof systems. For many practical use cases, this can cost $20 - 30$ bits of soundness, forcing the proof system to compensate for this gap via other means such as working over a larger field or using more bits of grinding – both of which can lead to a less efficient prover.

In this paper, we propose a simple variant of proximity gaps which uses *affine combiners* and show that it can achieve error on the order of $1/q$, surpassing the $\Omega(n/q)$-error lower bound for the standard scalar combiners considered in classical proximity gaps. We describe our variant and main results in detail next.

1.1 Our results

Throughout this section fix an integer degree bound $1 \le d < n$, a domain $\mathcal{D} \subseteq \mathbb{F}_q$ of size n, and received words $f, g : \mathcal{D} \to \mathbb{F}_q$. Set $\rho \coloneqq d/n$ to be the rate. Our affine combiners for proximity gaps work as follows. Randomly sample independent uniform $\alpha, \beta \in \mathbb{F}_q$ and form the combined word

$$w_{\alpha,\beta}(x) \coloneqq f(x) + (\alpha + \beta x)g(x), \qquad x \in \mathcal{D}.$$

We call the coefficient $\alpha + \beta x$ an *affine combiner*. Unlike scalar mixing, multiplication by this coefficient can increase polynomial degree by one. To keep the combined polynomial's degree at most d, we therefore use different degree bounds for the two words: in the honest case, f is the evaluation of a polynomial of degree at most d, and g of one of degree at most $d - 1$.

For many proof systems, we actually require a stronger notion of proximity gaps called *Mutual Correlated Agreement* (MCA) [ACFY24b]. At a high level, mutual correlated agreement says that, except with a small error probability (called MCA error), the combined word does not have *any* agreement domain I (with a degree d polynomial) which is not also a shared agreement domain of f (with degree d) and g (with degree $d - 1$). The asymmetry in our degree bounds also leads to a modified definition of MCA error which we define as follows.

Definition of the MCA error for affine combiners. Let $1 \le A \le n$ be an integer agreement threshold, and let the agreement loss $s \ge 0$ be an integer. Define

$$S_{\alpha,\beta}(P) \coloneqq \{x \in \mathcal{D} : P(x) = f(x) + (\alpha + \beta x)g(x)\}.$$

For (α, β), polynomials $F, G \in \mathbb{F}_q[X]$ of degrees at most d and $d - 1$, respectively, are admissible witnesses for $P = F + (\alpha + \beta X)G$ with agreement loss at most s if

$$\left|\{x \in S_{\alpha,\beta}(P) : F(x) \neq f(x) \text{ or } G(x) \neq g(x)\}\right| \le s.$$

We say that the affine combiner has MCA error at most ε with agreement loss at most s if, for every pair f, g,

$$\Pr_{\alpha,\beta \leftarrow \mathbb{F}_q}\left[\begin{array}{l}\exists P \in \mathbb{F}_q[X],\ \deg P \le d: \\ |S_{\alpha,\beta}(P)| \ge A \text{ and } P \text{ has no} \\ \text{admissible witnesses with agreement loss at most } s\end{array}\right] \le \varepsilon.$$

We denote the smallest such ε by $\varepsilon_{\mathrm{aff}}(n, \rho, A, s, q)$. Our main theorem shows that in the agreement loss $s = 1$ case, MCA error is $O_{\rho,\eta}(1/q + n^2/q^2)$. In this case we abbreviate the error as $\varepsilon_{\mathrm{aff}}$. For fixed ρ, η and $q \ge n^2$, the error is $O_{\rho,\eta}(1/q)$, removing the factor n which is necessary for scalar proximity gaps (see Section 6). Our results hold over all fields.

Theorem 1.1. *Fix $\rho \in (0,1)$ with ρn an integer, let $0 < \eta < 1 - \sqrt{\rho}$, and set $A \coloneqq \lceil (\sqrt{\rho} + \eta)n \rceil$. Suppose $n\eta^2 \ge 4/\rho$. Then, for every finite field $\mathbb{F}_q$ with $q \ge n$ and every fixed domain $\mathcal{D} \subseteq \mathbb{F}_q$ of size n, the MCA error with agreement loss $s = 1$ satisfies*

$$\varepsilon_{\mathrm{aff}}\Big(\underbrace{n}_{\substack{\text{block}\\ \text{length}}}, \underbrace{\rho}_{\substack{\text{degree}\\ \text{fraction}}}, \underbrace{A}_{\substack{\text{agreement}\\ \text{threshold}}}, \underbrace{1}_{\substack{\text{agreement}\\ \text{loss}}}, \underbrace{q}_{\substack{\text{field}\\ \text{size}}}\Big) \le \frac{4}{\eta q} + \frac{2^{14} n^2}{\rho^3 \eta^{10} q^2}. \tag{1.1}$$

Small error beyond unique decoding. Taking the 128-bit binary field $\mathbb{F}_{2^{128}}$, block length $n = 2^{20}$, rate $\rho = 1/4$, and agreement margin $\eta = 0.05$ (so $A = 576717$), the theorem gives

$$\varepsilon_{\mathrm{aff}}(2^{20}, 1/4, 576717, 1, 2^{128}) \le \frac{80}{2^{128}} + \frac{5^{10}}{2^{176}} \approx 2^{-122}.$$

Thus, with agreement loss $s = 1$, the MCA error bound is approximately 2^{-122}. In contrast, the classical scalar proximity-gap bound is $n/q = 2^{-108}$ even in the unique-decoding regime [BCI+20, Theorem 1.2].

Dealing with agreement loss in IOPPs. Moving from $s = 0$ to $s = 1$ reduces the guaranteed common agreement fraction by at most $1/n$ after the combination. Protocols can account for the agreement loss by making more queries. For example, suppose that, in a soundness argument, we bound the probability that t independent uniform queries all land in an agreement set by a^t, where $0 < a \le 1 - 1/n$. If the only change is agreement loss $s = 1$, the corresponding bound becomes $(a + 1/n)^t$, with multiplicative increase at most

$$\frac{(a + 1/n)^t}{a^t} = \left(1 + \frac{1}{an}\right)^t \le \exp\left(\frac{t}{an}\right).$$

For a constant agreement fraction and a query count much smaller than n, this has a negligible effect on the sampling contribution to soundness. For $n = 2^{20}$, $t = 128$, and $a \ge 1/2$, the soundness cost is less than 0.00036 bits.

Equal degree bound. For applications requiring equal input degree bounds, we give a simple modification. Sample an additional independent uniform $\theta \in \mathbb{F}_q$ and instead form

$$w_{\theta,\alpha,\beta}(x) \coloneqq g(x) + (\alpha + \beta x)(f(x) + \theta g(x)).$$

Both witness polynomials may now have degree at most $d - 1$, while the combined polynomial still has degree at most d. In Section 5, we show that this modification achieves essentially the same MCA error.

Application: smaller fields for Flock and leanVM. Flock [BRW26] uses the Ligerito polynomial commitment scheme over $\mathbb{F}_{2^{128}}$, where the block-length factor, n, in scalar MCA bounds creates a substantial soundness gap at the 128-bit security target. Moving to a larger field increases arithmetic and communication costs; leanVM uses a 192-bit binary field [lea26]. Affine combiners remove this factor n from the MCA error and could support efficient 128-bit arithmetic in these systems. For the representative parameters above, Remark 5.3 bounds the per-round MCA contribution below 2^{-128} with 3–5 grinding bits per fold challenge, depending on the variant. Establishing 128-bit security for a full system additionally requires accounting for composition across folds, query errors, and the security of the commitments and Fiat–Shamir transformation. Section 5 develops the adaptations for linear coefficient claims, degree growth, and shared agreement across interleaved rows needed to integrate affine combiners into these proof systems.

Example: the toy protocol. To further demonstrate the benefits of affine combiners, we consider the toy protocol in [ABF26, Construction 6.2] with evaluation at zero as the linear constraint. The verifier receives $u, v \in \mathbb{F}_q$ and has oracle access to received words $f, g : \mathcal{D} \to \mathbb{F}_q$, where $|\mathcal{D}| = n$ and $0 \notin \mathcal{D}$. The honest prover knows polynomials $F, G \in \mathbb{F}_q[X]$ satisfying $F(0) = u$, $G(0) = v$, and

$$\deg F \le d, \quad \deg G \le d-1, \qquad F|_{\mathcal{D}} = f, \quad G|_{\mathcal{D}} = g.$$

The verifier samples independent uniform $\alpha, \beta \in \mathbb{F}_q$, and the prover sends the coefficients of a candidate polynomial P with $\deg P \le d$. The verifier then samples t independent uniform coordinates $x_1, \dots, x_t \in \mathcal{D}$, with replacement, and accepts if

$$P(0) = u + \alpha v, \qquad P(x_i) = w_{\alpha,\beta}(x_i) \quad \text{for every } i.$$

For the honest response $P = F + (\alpha + \beta X)G$, acceptance is certain.

We denote by ℓ the number of polynomial pairs with the stated degree bounds that agree with f, g on at least $A - 1$ common coordinates. If no such pair satisfies both evaluation claims, then by Theorem 1.1,

$$\Pr[\text{accept}] \le \varepsilon_{\text{aff}}(n, \rho, A, 1, q) + \frac{\ell}{q} + \left(\frac{A-1}{n}\right)^t. \tag{1.2}$$

Indeed, for each pair that does not satisfy both evaluation claims, $(F(0) - u) + \alpha(G(0) - v) = 0$ for at most q choices of (α, β). After excluding at most ℓq such challenges and the MCA exceptional set, $|S_{\alpha,\beta}(P)| \le A - 1$ whenever $P(0) = u + \alpha v$. By independence of the queries, we obtain (1.2), as in [ABF26, Lemmas 6.6 and 6.8].

Concretely, for $n = 2^{14}$, $d = n/4 - 1$, $A = \lceil 0.55n \rceil$, and $t = 160$, $\eta = 0.55 - \sqrt{\rho} > 0.05$ and $\ell \le 5$ by the Johnson bound. For scalar combiners, with degree bounds $\deg F, \deg G \le d$, agreement loss $s = 0$, and the same n, A, t, we obtain the soundness error bound $2^{30}/q + 0.55^{160}$ from [ABF26, Theorem 4.12]. Without grinding, the soundness error is less than 2^{-128} over $\mathbb{F}_{2^{136}}$ with affine combiners and over $\mathbb{F}_{2^{160}}$ with scalar combiners. The corresponding response sizes are 68 KiB and 80 KiB, a 15% reduction with affine combiners. For affine combiners, the verifier sends one additional field element and performs one additional multiplication per queried coordinate. We require $\deg G \le d - 1$ and allow agreement loss $s = 1$, keeping the output degree bound d and query count t in both protocols.

Extension to Reed–Muller codes. For Reed–Muller codes, which consist of evaluations of multivariate low-degree polynomials, we obtain MCA error *even below* $1/q$. Their multivariate structure gives us more degrees of freedom in the random combination. For k-variate codes, we

use an affine combiner $\alpha_0 + \alpha_1 x_1 + \cdots + \alpha_k x_k$, with independent uniform coefficients in $\mathbb{F}_q$. For fixed degree fraction ρ, Johnson slack η, and number of variables k, Section 7 gives error $O_{\rho,\eta,k}(q^{-k})$ for sufficiently large block length n and $q \geq n^{k+1}$. We use the agreement threshold of Theorem 1.1 and allow agreement loss $n^{1-1/k}$. Thus the number of variables becomes the error exponent, giving $o(1/q)$ error for every fixed $k \geq 2$. We believe that this phenomenon is related to the rich invariance structure of Reed–Muller codes and it would be interesting to see if proximity gaps (via some modified combination like here) below $1/q$ could be shown for any other codes with a large group of invariances.

1.2 Related Work and Further work

We discuss a few lines of work related to the questions we consider.

Progress and limits beyond Johnson. Recent work substantially extends the range of agreement thresholds in large characteristic. Building on the list-decoding breakthrough of Brakensiek, Chen, Putterman, Zhang, and Zheng [BCP$^+$26], Jeronimo [Jer26] proves exact-support MCA over prime fields on arbitrary evaluation domains, up to any fixed slack $\gamma > 0$ from capacity. For scalar lines, the error is $n^{O_\gamma(1)}/q$, with agreement loss $s = 0$. Dao, Kominers, and Thaler [DKT26, Theorem 1.1] give a first-order agreement threshold $a_1(\rho) < \sqrt{\rho}$ in sufficiently large characteristic. At agreement $a_1(\rho) + \eta_1$ with $\eta_1 > 0$, they bound the list size by $O_\rho(n/\eta_1^2)$ and the scalar MCA error by $O_\rho(n^2/(q\eta_1^4))$. They also give polynomial list-size and exceptional-challenge bounds up to any fixed gap from capacity [DKT26]. Gabizon [Gab26] proves a linear list-size bound at an agreement threshold strictly below Johnson over fields of sufficiently large characteristic. Harsha, Kumar, and Saptharishi [HKS26] give a unified exposition of the capacity list-decoding and proximity-gap results.

Recent work establishes limits beyond the Johnson radius over fields of small characteristic [DKTZ26, AGG26]. In particular, Dao, Kominers, Thaler, and Zheng rule out a general $\mathrm{poly}(n)/q$ proximity-gap error bound on additive binary domains. These results highlight the distinction between extending the decoding radius and reducing the error within a given radius. Our focus is on the latter: we remain in the Johnson regime and prove a small, explicit error bound over every finite field by using affine combiners with agreement loss $s = 1$.

We remark that we can also extend our analysis to achieve slightly better MCA error beyond Johnson using the one derivative interpolator, analogous to the (slightly) beyond Johnson proximity gap results in [DKT26, Gab26], but leave it to future work to give a clean writeup.

Concurrent work. Frigo and shelat's *Polygero* [Fs26] consider a version of proximity gaps that uses polynomial coefficients. Instead of one scalar for each received word w_i, the verifier samples an independent polynomial B_i of degree at most r and tests the combination $\sum_i B_i w_i$. Roughly speaking, their analysis gives MCA error $O(1/q)$ with agreement loss up to r coordinates using degree $r = \Omega(n/\log_2 q)$ challenges. As a concrete example, [Fs26, Remark 2] uses degree-74 challenges for a code of length 25724 and rate about $1/16$ over $\mathbb{F}_{2^{128}}$, with 64 queried coordinates and soundness error at most 2^{-120}.

Our work shows that *degree-one challenges already give small error with agreement loss only one.* At fixed rate and fixed slack from Johnson, we bound the probability of an exceptional combiner by $O(1/q + n^2/q^2)$, and hence by $O(1/q)$ when $q \geq n^2$.

2 Techniques

In this section, we overview our techniques and give some intuition on why affine combiners can achieve error on the order of $O(1/q)$. We begin with the intuition, and our inspiration for considering affine combiners in the first place. For the remainder of the section, fix the evaluation domain $\mathcal{D}$, a degree d, received words $f, g : \mathcal{D} \to \mathbb{F}_q$ and set the agreement parameter to be $A = \lceil (\sqrt{\rho} + \eta)n \rceil$, where $\eta > 0$ is a slack parameter and $\rho = d/n$ is the rate.

We were motivated by the following view of proximity gaps as a low-degree test for a special type of input function, which we think may be instructive to share. Starting from, f, g, form the table

$$H(x, z) \coloneqq f(x) + zg(x), \qquad (x, z) \in \mathcal{D} \times \mathbb{F}_q.$$

In the honest case, when f, g are evaluations of polynomials F, G with $\deg F \leq d$ and $\deg G \leq d-1$, this table agrees with $F(X) + ZG(X)$, a bivariate polynomial of total degree at most d that is linear in Z. An important fact about low degree bivariate polynomials is that their restrictions to affine lines are low degree as well. Specifically, restricting H to the affine line intersected with $\mathcal{D} \times \mathbb{F}_q$,

$$\ell_{\alpha,\beta} \coloneqq \{(x, \alpha + \beta x) : x \in \mathcal{D}\},$$

gives the univariate function

$$x \longmapsto H(x, \alpha + \beta x) = f(x) + (\alpha + \beta x)g(x),$$

which has degree at most d in the honest case. Hence, these restrictions give exactly our affine combinations. On the other hand, scalar combinations correspond to restricting only to the axis-parallel lines $z = \alpha$.

Viewed this way, affine combiners are analogous to the affine line-versus-point test of Arora and Sudan [AS03], whereas scalar combiners are analogous to the axis-parallel test of Polishchuk and Spielman [PS94]. The former admits analyses in the list-decoding regime [HKSS24, KTZ26], whereas the classical Polishchuk–Spielman analysis applies only in the unique-decoding regime. This gives some motivation for affine combiners potentially achieving smaller MCA error.

A natural question is then the following: if at least $\Omega(1/q)$ fraction of the affine-line restrictions of H above agree with degree-at-most-d polynomials on at least A coordinates, does H have comparably large agreement with a global total-degree-d polynomial that is linear in Z? Such a polynomial would have the form $F(X) + ZG(X)$ with F, G of degrees at most $d, d-1$ respectively, and it is not hard to see that concluding agreement of H with such a polynomial would lead to correlated agreement of f, g with F, G respectively (up to a one coordinate loss).

To this end, Minzer and Zheng [MZ25] give a motivating result for total-degree Reed–Muller codes. For an absolute constant $\tau > 0$, let $\varepsilon \geq (d/q)^\tau$ and $u : \mathbb{F}_q^3 \to \mathbb{F}_q$. Suppose that for more than a $100/(\varepsilon^2 q)$ fraction of affine planes $P \subseteq \mathbb{F}_q^3$, the restriction $u|_P$ agrees with a polynomial of total degree at most d on at least a 1.2ε fraction of P. Then u agrees with a trivariate polynomial of total degree at most d on more than an ε fraction of $\mathbb{F}_q^3$. Thus agreement on more than $O_\varepsilon(1/q)$ of the plane restrictions forces global agreement, giving evidence that global structure can potentially be concluded from such a small fraction of restrictions. One can also check that a typical choice of $\Omega(1/q)$-fraction of affine lines intersected with $\mathcal{D} \times \mathbb{F}_q$ covers a constant fraction of the domain $\mathcal{D} \times \mathbb{F}_q$ whereas $\Omega(1/q)$-fraction of the axis parallel lines in $\mathcal{D} \times \mathbb{F}_q$ covers only a negligible fraction of the domain.

Proof overview. The proof follows the same high-level strategy as in the scalar case, developed in [BCI+20], requiring mostly minor adaptations to handle the affine-combiners we consider. We heavily borrow tools from prior works [BCI+20, BCH+25, DKT26].

We start by interpolating a low-degree "explainer" polynomial $Q(X, Y, Z)$ with the following property. For every $\alpha, \beta \in \mathbb{F}_q$ and polynomial P of degree at most d whose agreement domain $I = S_{\alpha,\beta}(P)$ has size at least A, we require

$$Q(X, \alpha + \beta X, P(X)) = 0.$$

In particular, this identity holds for every *exceptional triple*, i.e. a triple for which I contains at least two points outside the correlated agreement domain of every polynomial pair F, G satisfying $\deg F \le d$, $\deg G \le d-1$, and $P = F + (\alpha + \beta X)G$. We construct Q by adapting the interpolation argument of [BCH+25].

Once such a Q is found, we count the pairs (α, β) for which this identity has a polynomial solution P forming an exceptional triple. This counting is again done via standard techniques, this time via adaptations of the factorization and root lifting arguments from [BCI+20, BCH+25, DKT26]. For simplicity, assume that Q is irreducible; otherwise we factor Q and apply the argument to each irreducible factor. For factors of positive Z-degree, we also describe the case $Q_Z \not\equiv 0$, where $Q_Z = \partial Q / \partial Z$. The proof handles factors with identically zero derivative in small characteristic by a change of variables. We partition the exceptional triples into three classes.

- *Triples over excluded coefficient pairs.* In this step, we first count the coefficient pairs (α, β) occurring in roots (α, β, P) such that either

$$\begin{aligned} Q(X, \alpha + \beta X, Z) &\equiv 0, \qquad \text{or} \\ Q_Z(X, \alpha + \beta X, P(X)) &\equiv 0. \end{aligned}$$

  The reason for filtering out these roots first is that our later arguments need both expressions above to be nonzero. If the first is zero, then the equation above holds for every P and gives no information on which polynomial P can be. For irreducible Q, this can happen only if Q is independent of Z. Each such pair gives a linear factor $Y - \alpha - \beta X$ of Q. Distinct pairs give distinct factors of Y-degree one, so there are at most $\deg_Y Q$ such pairs.

  For the second, we will need a nonzero derivative in order to perform Hensel lifting. Specifically, once α, β are fixed, we use the known equation $Q(X, \alpha + \beta X, P(X)) = 0$, together with one value of P, to recover the coefficients of P one at a time. Each step requires dividing by the Z-derivative at the starting root. Thus we need to start at a point where this derivative is nonzero, which we can choose when $Q_Z(X, \alpha + \beta X, P(X)) \not\equiv 0$. To count this second case, observe that P is a common root of $Q(X, \alpha + \beta X, Z)$ and its Z-derivative. Hence their resultant vanishes along the challenge line:

$$\mathrm{Res}_Z(Q, Q_Z)|_{Y=\alpha+\beta X} \equiv 0.$$

  This resultant is nonzero because Q is irreducible and Q_Z is nonzero of smaller Z-degree. Each such pair therefore gives a linear factor $Y - \alpha - \beta X$ of the resultant, and we bound their number by its Y-degree, just as in the first case. This also counts challenges at which the restriction has any repeated root, whether or not that root is a polynomial P.

  Straightforward degree counts then bound the total number of excluded pairs by D^2, after summing over all factors. Here we use that the interpolated explainer satisfies $\deg_{Y,Z} Q \le D = O(1/(\rho\eta^2))$, independent of n.

- *Persistent triples.* In this step, we count exceptional triples whose agreement can be explained by correlated agreement with a pair of rational functions. That is, (α, β, P) such that we can write
$$P = F + (\alpha + \beta X)G,$$
where F, G have a common denominator of degree at most one and numerators of degree at most $d+1$ and d, respectively, and have correlated agreement with f, g on a large set $I \subseteq \mathcal{D}$. Such an explanation need not be an MCA witness.

  Applying standard list-decoding up to the Johnson bound to such rational pairs gives at most $4/\eta$ pairs of rational functions that can have correlated agreement with the received words. For each fixed pair of rational functions F, G, a straightforward counting argument bounds the number of exceptional coefficient pairs (α, β) associated with it by $\max\{\binom{n}{2}, q\}$. This gives an MCA error of
$$\frac{4}{\eta q^2} \max\{\tbinom{n}{2}, q\} = O_\eta(1/q + n^2/q^2),$$
coming from persistent triples.

- *Nonpersistent triples.* We count these triples by covering them with algebraic sets of bounded degree and dimension at most two and then using tools from [DKT26] to bound the number of nonpersistent triples in each algebraic set. This step relies on Hensel lifting to construct the algebraic sets, which is why we first filtered out the triples with Z-derivative equal to zero (after plugging in to Q) in the excluded coefficient case. The absence of a rational explanation (F, G) as in the persistent case is used when bounding the number of nonpersistent triples in each algebraic set.

Altogether, this gives MCA error at most $\frac{4}{\eta q} + O_{\rho,\eta}(n^2/q^2)$.

Where the gain in the proof is. Here we give some more detail on where in the proof affine combiners have an advantage over scalar combiners. At a high level, prior proximity gaps results [BCI+20, BCH+25, DKT26] follow the same partition when counting exceptional candidates. The error due to persistent and nonpersistent candidates is $O(n/q)$ for each. In our case, both give an $O(n^2/q^2)$ term, with an additional $O(1/q)$ term for persistent candidates.

Persistent triples. Here the gain of affine combiners is relatively easy to see, and one can consider the following toy scenario for simplicity. Suppose f, g have correlated agreement with polynomials F, G of degrees at most $d, d-1$, respectively, and let I be their full common agreement domain. For every challenge (α, β), the polynomial
$$P_{\alpha,\beta}(X) = F(X) + (\alpha + \beta X)G(X)$$
agrees with the combined word on I. An agreement at a coordinate $x \in \mathcal{D} \setminus I$ requires cancellation of the discrepancies, so
$$(F(x) - f(x)) + (\alpha + \beta x)(G(x) - g(x)) = 0.$$
We call this an accidental agreement. With one coordinate of agreement loss, an exceptional candidate needs at least two such agreements and at any accidental agreement we have $g(x) \neq G(x)$, so writing $r(x) = (F(x) - f(x))/(g(x) - G(x))$, two distinct such coordinates give
$$\alpha + \beta x_1 = r(x_1), \qquad \alpha + \beta x_2 = r(x_2).$$

These equations determine exactly one (α, β). Taking a union bound over pairs of coordinates outside I gives

$$\Pr_{\alpha,\beta}[\text{at least two accidental agreements}] \leq \frac{\binom{n-|I|}{2}}{q^2}.$$

For scalar combiners, the same two coordinates give the equations $\alpha = r(x_1)$ and $\alpha = r(x_2)$. We can arrange the coordinates outside I into disjoint pairs and choose the values of f, g so that the two equations in each pair have the same solution, with distinct solutions across pairs. This gives $\lfloor (n - |I|)/2 \rfloor$ scalar challenges with two accidental agreements, out of q possible challenges. The construction in Section 6 makes such challenges exceptional and gives the scalar $\Omega(n/q)$ lower bound.

Nonpersistent triples. Here the second challenge coefficient β gives us q^2 possible challenges and increases the dimension of the algebraic sets used in the count. We cover the triples (α, β, P) by algebraic sets of dimension at most two and total degree $O(n^2)$. The scalar proof of Dao, Kominers, and Thaler [DKT26] covers pairs (α, P) by algebraic sets of dimension at most one and total degree $O(n)$. Ultimately, this gives $O(n^2/q^2)$ error for affine combiners and $O(n/q)$ error for scalar combiners.

3 Preliminaries

Throughout this paper, we work over a finite field $\mathbb{F}_q$, and fix a size n evaluation domain $\mathcal{D} \subseteq \mathbb{F}_q$. Let $d \in \mathbb{N}$ be a degree bound and set $\rho = d/n$ as the fractional rate along with the following parameters

$$\sqrt{\frac{4}{d}} \leq \eta < 1 - \sqrt{\rho}, \qquad A \coloneqq \lceil (\sqrt{\rho} + \eta)n \rceil . \tag{3.1}$$

Above, η is a slack parameter and A is an absolute agreement parameter. We fix two received words $f, g : \mathcal{D} \to \mathbb{F}_q$. When working with proximity gaps, it will be helpful to define the notion of exceptional triples (α, β, P) which are essentially triples where $\deg(P) \leq d$ and P and $w_{\alpha,\beta}$ agree on at least A points over $\mathcal{D}$, but this agreement is not explained by correlated agreement between f, g with degree $d, d-1$ polynomials.

Definition 3.1. A triple $(\alpha, \beta, P) \in \mathbb{F}_q^2 \times \mathbb{F}_q[X]$ is *exceptional* if $\deg P \leq d$, $|S_{\alpha,\beta}(P)| \geq A$, and P has no admissible witnesses with agreement loss at most 1, as defined in Section 1.1.

In this language, bounding the MCA error ε_{aff} from Theorem 1.1 is equivalent to bounding the number of pairs (α, β) that can occur in an exceptional triple. The proof in the remainder of the paper will take this viewpoint.

We will require the following form of the Johnson bound.

Lemma 3.2 ([MS77, Ch. 17, Theorem 3])**.** *Let $0 < a \leq 1$ and $0 \leq b < a^2$. Any family of subsets of an n-element set, each of size at least an and with pairwise intersections of size at most bn, has at most $\frac{a-b}{a^2-b}$ members.*

3.1 Algebraic Geometry

In this section, we introduce some notions and results from algebraic geometry that will be useful. Let $\mathbb{K}$ be an algebraically closed field containing $\mathbb{F}_q$.

Definition 3.3. An *algebraic set* in $\mathbb{K}^k$ is a common zero set

$$V = \{x \in \mathbb{K}^k : F_1(x) = \cdots = F_s(x) = 0\}, \qquad F_1, \ldots, F_s \in \mathbb{K}[X_1, \ldots, X_k].$$

A *hypersurface* is the zero set of one nonconstant polynomial; a *hyperplane* is the zero set of a nonconstant affine linear polynomial. A nonempty algebraic set is *irreducible* if it cannot be written as the union of two proper algebraic subsets. Every algebraic set is a finite union of its maximal irreducible algebraic subsets, called its *irreducible components*.

Definition 3.4. The *Zariski closure* of $U \subseteq \mathbb{K}^k$, denoted $\overline{U}$, is the smallest algebraic set containing U.

Definition 3.5. For a nonempty algebraic set V, define

$$\dim V \coloneqq \max\{r : C_0 \subsetneq C_1 \subsetneq \cdots \subsetneq C_r \subseteq V\},$$

where the C_i are nonempty irreducible algebraic sets. The dimension of V is the largest dimension of its irreducible components. We set $\dim \varnothing \coloneqq -1$.

Definition 3.6. For an irreducible algebraic set $C \subseteq \mathbb{K}^k$ of dimension r, let $L \subseteq \mathbb{K}^k$ be a generic affine linear subspace of dimension $k - r$, and define

$$\deg C \coloneqq |C \cap L| \qquad \text{(counted with multiplicity)}.$$

Here *generic* means that L is chosen outside a fixed proper algebraic subset of the parameter space of such subspaces. For an algebraic set with irreducible components $C_1, \ldots, C_s$, define

$$\deg V \coloneqq \sum_{i=1}^{s} \deg C_i,$$

including components of smaller dimension. A zero-dimensional algebraic set is finite and satisfies $\deg V = |V|$. We set $\deg \varnothing \coloneqq 0$.

Definition 3.7. Let $V \subseteq \mathbb{K}^k$ be algebraic and let $N_1, \ldots, N_t, S \in \mathbb{K}[X_1, \ldots, X_k]$. On $U \coloneqq \{x \in V : S(x) \neq 0\}$, these polynomials define a *rational map*

$$\phi : U \longrightarrow \mathbb{K}^t, \qquad x \longmapsto \left(\frac{N_1(x)}{S(x)}, \ldots, \frac{N_t(x)}{S(x)} \right).$$

Its *image closure* is $\overline{\phi(U)}$.

The next two lemmas are standard facts from algebraic geometry.

Lemma 3.8. *A hypersurface in $\mathbb{K}^k$ defined by a polynomial of degree $b \geq 1$ has dimension $k - 1$ and total degree at most b.*

Lemma 3.9. *If $C \subseteq \mathbb{K}^k$ is an irreducible algebraic set and a hyperplane $H \subseteq \mathbb{K}^k$ does not contain C, then*

$$\dim(C \cap H) \leq \dim C - 1, \qquad \deg(C \cap H) \leq \deg C. \tag{3.2}$$

Proof. Since $C \cap H$ is a proper closed subset of the irreducible set C, its dimension is strictly smaller than $\dim C$. The degree bound is from the hyperplane case of Bézout's inequality [Hei83, Theorem 1], which gives $\deg(C \cap H) \leq \deg C \cdot \deg H = \deg C$. □

Lemma 3.10 ([DKT26, Lemma A.4])**.** *Let V be an algebraic set of dimension at most $r \geq 0$, and let $\phi \coloneqq (N_1/S, \ldots, N_t/S)$, where all numerators and the common denominator have degree at most $b \geq 1$. If S is not identically zero on any irreducible component of V, the image closure $W \coloneqq \overline{\phi(\{x \in V : S(x) \neq 0\})}$ satisfies*

$$\dim W \leq r, \qquad \deg W \leq (\deg V) b^r. \tag{3.3}$$

Components contained in $S = 0$ can first be discarded without increasing the dimension or degree of V.

Hensel lifting. Suppose we have a polynomial equation $R(X, \alpha + \beta X, Z) = 0$ and at a particular setting, $X = t_0$, it has a root at $Z = z$, i.e. $R(t_0, \alpha + \beta t_0, z) = 0$. Then, if the derivative of R with respect to Z is nonzero at $(t_0, \alpha + \beta t_0, z)$, Hensel lifting constructs a formal power series $z + E(X - t_0)$ with $E(0) = 0$ that solves the equation and captures its polynomial roots with initial value z. More precisely, if a polynomial P satisfies $R(X, \alpha + \beta X, P(X)) = 0$ and $P(t_0) = z$, then

$$z + E(X - t_0) = P(X).$$

The following lemma expresses the coefficients of E as rational functions of α, β, z and bounds their degrees.

Lemma 3.11. *Let $\mathbb{K}$ be a field, let $R \in \mathbb{K}[X, Y, Z]$, and fix $t_0, \alpha, \beta, z \in \mathbb{K}$. Suppose z is a root of $R(t_0, \alpha + \beta t_0, Z) = 0$ and $R_Z(t_0, \alpha + \beta t_0, z) \neq 0$, where $R_Z = \partial R/\partial Z$. Then there is a unique formal power series $z + E(X - t_0)$ with $E(0) = 0$ satisfying*

$$R(X, \alpha + \beta X, z + E(X - t_0)) = 0 \tag{3.4}$$

In particular, any polynomial root $P(X)$ with $P(t_0) = z$ equals $z + E(X - t_0)$.

Moreover, writing $E(X - t_0) = \sum_{j \geq 1} c_j(\alpha, \beta, z)(X - t_0)^j$, we can express each coefficient in the power series as

$$c_j(\alpha, \beta, z) = \frac{N_j(\alpha, \beta, z)}{S(\alpha, \beta, z)^{2j-1}}, \tag{3.5}$$

where $S(Y, V, Z) = R_Z(t_0, Y + Vt_0, Z)$, $h = \deg_{Y,Z} R$, and $N_j \in \mathbb{K}[Y, V, Z]$ has total degree at most $1 + (2j - 1)(h - 1)$ for all $j \geq 1$.

The proof essentially follows [BCI+20, Claim A.2], with the notation adapted to the affine term in the second variable. We include a brief proof for completeness.

Proof. Set $W = X - t_0$. The polynomial $R(t_0 + W, \alpha + \beta(t_0 + W), z + U)$ vanishes at $W = U = 0$ and its U-derivative there is nonzero by assumption. Then Hensel lifting over $\mathbb{K}$ (see [KTZ26, Lemma 2.7]) gives existence and uniqueness. If $P(X)$ is a polynomial root with $P(t_0) = z$, then $P(X) - z$ has zero constant term in $X - t_0$ and satisfies (3.4) in place of $E(X - t_0)$. Uniqueness gives $P(X) = z + E(X - t_0)$.

For the coefficient bounds, we take the given identity $R(X, \alpha + \beta X, Z) = 0$, shift the variable $Z = z + U$, and write out the expansion in these variables as

$$R(t_0 + W, \alpha + \beta(t_0 + W), z + U) = \sum_{a,b \geq 0} A_{a,b}(\alpha, \beta, z) W^a U^b,$$

where $A_{a,b} \in \mathbb{K}[Y, V, Z]$ is the coefficient of $W^a U^b$ in $R(t_0 + W, Y + V(t_0 + W), Z + U)$. Since $h = \deg_{Y,Z} R$, we have $\deg A_{a,b} \leq h - b$. The coefficient of U at $W = 0$ in this expansion is $S(\alpha, \beta, z)$, and $\deg S \leq h - 1$. We write S for $S(\alpha, \beta, z)$ in the coefficient comparison below.

Write $E(W) = \sum_{j \geq 1} c_j W^j$. Substituting $U = E(W)$ into the expansion gives

$$0 = \sum_{a,b \geq 0} A_{a,b}(\alpha, \beta, z) W^a \left(\sum_{j \geq 1} c_j W^j \right)^b.$$

To prove (3.5), we compare coefficients of W^j on both sides of the equation one by one to sequentially "reconstruct" all of them. Then, the degree bound on N_j can be proved inductively.

First, for $j = 1$, the coefficient of W on the right handside must cancel, so we get $A_{1,0}(\alpha, \beta, z) + Sc_1 = 0$ and

$$c_1 = -\frac{A_{1,0}(\alpha, \beta, z)}{S}.$$

Then, the $j = 1$ case of (3.5) holds with $N_1 = -A_{1,0}$,and indeed $\deg N_1 \leq h$.

For $j \geq 2$, the coefficient of W^j is Sc_j plus terms of the form

$$A_{a,b} c_{j_1} \cdots c_{j_b}, \qquad a + j_1 + \cdots + j_b = j, \quad 1 \leq j_1, \ldots, j_b < j,$$

c_j is $-1/S$ times the sum of these terms. By induction, we can write each c_{j_ℓ} in the form $N_{j_\ell}/S^{2j_\ell-1}$. $c_{j_\ell} = N_{j_\ell}/S^{2j_\ell-1}$ where N_{j_ℓ} satisfies the degree bound in (3.5). Substituting and using $j_1 + \cdots + j_b = j - a$, each summand of c_j becomes

$$\begin{aligned} -\frac{1}{S} \cdot A_{a,b} \prod_{\ell=1}^{b} \frac{N_{j_\ell}}{S^{2j_\ell-1}} &= -\frac{1}{S} \cdot \frac{A_{a,b} N_{j_1} \cdots N_{j_b}}{S^{2(j-a)-b}} \\ &= -\frac{1}{S} \cdot \frac{A_{a,b} S^{2a+b-2} N_{j_1} \cdots N_{j_b}}{S^{2j-2}}. \end{aligned}$$

If $a = 0$, then $b \geq 2$ since the indices j_ℓ sum to j and each is less than j. Thus $2a + b - 2 \geq 0$ and each exponent of S is positive. This shows that we can write c_j with denominator S^{2j-1} as desired. Then N_j is the sum of the numerators above and it remains to check its degree bound. Using $\deg A_{a,b} \leq h - b$, $\deg S \leq h - 1$, and the inductive bounds, each numerator satisfies

$$\deg\Big(A_{a,b} S^{2a+b-2} N_{j_1} \cdots N_{j_b}\Big) \leq h + (2j-2)(h-1) = 1 + (2j-1)(h-1)$$

as desired $\square$

4 Proof of the proximity gap for affine combiners

We now prove Theorem 1.1 by bounding the number of challenges (α, β) which are part of an exceptional triple.

Lemma 4.1. *For every $f, g : \mathcal{D} \to \mathbb{F}_q$, the number of pairs $(\alpha, \beta) \in \mathbb{F}_q^2$ for which (α, β, P) is exceptional for some P is at most*

$$\frac{4q}{\eta} + \frac{2^{14} n^2}{\rho^3 \eta^{10}}.$$

Dividing the bound above by q^2 proves Theorem 1.1. The remainder of this section is dedicated to the proof of Lemma 4.1. We follow the interpolation and root-counting approach of Guruswami and Sudan [GS99] for list decoding, extended to proximity gaps by Ben-Sasson, Carmon, Ishai, Kopparty, and Saraf [BCI+20].

Specifically, from f and g, we interpolate a nonzero polynomial $Q \in \mathbb{F}_q[X, Y, Z]$ such that every exceptional triple (α, β, P) satisfies

$$Q(X, \alpha + \beta X, P(X)) = 0 \qquad \text{in } \mathbb{F}_q[X].$$

Once Q is fixed, to prove Lemma 4.1, we must bound the number of coefficient pairs (α, β) for which this there exists P forming an exceptional triple with it.

4.1 Interpolation along affine lines

We construct one polynomial relation satisfied by every candidate with at least A agreements. For $H(X, Y, Z) = \sum_{i,j,k} h_{ijk} X^i Y^j Z^k \neq 0$, define its $(1, 1, d)$-weighted degree by

$$\deg_{1,1,d}(H) \coloneqq \max_{h_{ijk} \neq 0} (i + j + dk).$$

For the interpolation, define

$$m \coloneqq \lceil 2/\eta \rceil, \qquad D \coloneqq \left\lceil \frac{(m + 1/2)^2}{3\rho} \right\rceil - 1,$$

and

$$M \coloneqq \min \left\{ D, \left\lfloor \frac{mA - D - 1}{d} \right\rfloor \right\}. \tag{4.1}$$

Lemma 4.2. *For every pair of received words $f, g : \mathcal{D} \to \mathbb{F}_q$, there exists a nonzero polynomial $Q \in \mathbb{F}_q[X, Y, Z]$ such that:*

- $\deg_{Y,Z} Q \leq D$ *and* $1 \leq \deg_Z Q \leq M$.
- $\deg_{1,1,d} Q < mA$.
- *For every $\alpha, \beta \in \mathbb{F}_q$ and every $P \in \mathbb{F}_q[X]$ with $\deg P \leq d$ and $|S_{\alpha,\beta}(P)| \geq A$,*

$$Q(X, \alpha + \beta X, P(X)) = 0. \tag{4.2}$$

Proof. We adapt the interpolation argument of [BCH+25, Lemma 3.1]. Consider polynomials of the form

$$Q(X, Y, Z) = \sum_{\substack{i,t,j \geq 0 \\ t + j \leq D,\ i + dj < mA - D}} Q_{i,t,j} X^i Y^t Z^j,$$

where the coefficients $Q_{i,t,j} \in \mathbb{F}_q$ are treated as unknowns. Each monomial has $j \le M$, $t + j \le D$, and $i + t + dj < mA$, so the required degree upper bounds hold for all polynomials we are considering. Let us view Q as a polynomial in X, Z over $\mathbb{F}_q[Y]$. Then, for each $z \in \mathcal{D}$, we require that Q has a root of multiplicity at least m at $(z, f(z) + Yg(z))$. Specifically, for every $r, s \ge 0$ with $r + s < m$, impose

$$\sum_{\substack{i \ge r,\ t \ge 0,\ j \ge s \\ t+j \le D,\ i+dj < mA-D}} Q_{i,t,j} \binom{i}{r} \binom{j}{s} z^{i-r} Y^t \big(f(z) + Yg(z)\big)^{j-s} = 0 \tag{4.3}$$

where the equation is over $\mathbb{F}_q[Y]$.

Each summand has Y-degree at most $t + j - s \le D - s$. Thus, in (4.3), we impose at most $D - s + 1$ homogeneous linear equations over $\mathbb{F}_q$, one for each coefficient of Y. For a fixed $s < m$, there are $m - s$ choices of r and n choices of z; if $s > M$, the sum is empty. The total number of scalar equations is therefore at most

$$n \sum_{s=0}^{\min(m-1, M)} (m - s)(D - s + 1). \tag{4.4}$$

Now we show that any Q satisfying these constraints satisfies the vanishing condition in the third item above. Suppose $z \in \mathcal{D}$ is a point where $P(z) = w_{\alpha,\beta}(z)$ for some $\alpha, \beta \in \mathbb{F}_q$. The multiplicity constraints imply that Q is a sum of polynomial multiples of

$$(X - z)^r \big(Z - f(z) - Yg(z)\big)^s, \qquad r + s = m.$$

After substituting $Y = \alpha + \beta X$ and $Z = P(X)$, the expression $Z - f(z) - Yg(z)$ becomes $P(X) - f(z) - (\alpha + \beta X)g(z)$ which is divisible by $X - z$. Thus each summand in (4.3) is divisible by $(X - z)^m$, and z is a root of multiplicity at least m of $Q(X, \alpha + \beta X, P(X))$.

If P agrees with $w_{\alpha,\beta}$ on at least A points, we obtain at least mA roots counted with multiplicity. On the other hand, $\deg P \le d$ and the weighted-degree bound on Q give

$$\deg Q(X, \alpha + \beta X, P(X)) \le \deg_{1,1,d} Q < mA.$$

Therefore $Q(X, \alpha + \beta X, P(X))$ vanishes identically as required by the third item above.

It remains to show that the constraints have a nonzero solution. For each $0 \le j \le M$, there are $mA - D - dj$ choices of i and $D - j + 1$ choices of t. The number of unknowns is therefore

$$\sum_{j=0}^{M} (mA - D - dj)(D - j + 1).$$

Set $U := \left\lceil (m + 1/2)\sqrt{nd} \right\rceil$. A direct calculation, as in the proof of [BCH+25, Lemma 3.1, equations (11)–(12)] with their $k = d$, gives

$$\sum_{j=0}^{\lfloor (U-1)/d \rfloor} (U - dj)(D - j + 1) > n \sum_{s=0}^{m-1} (m - s)(D - s + 1),$$

where $m \le \lfloor (U - 1)/d \rfloor \le D$. By our choice of m, we have $m \ge 3$ and $2 \le m\eta < 3$, so by (3.1),

$$D + 1 \le \frac{(m + 1/2)^2}{3\rho} + 1 \le \frac{2m^2}{3\rho} \le \frac{m\eta n}{2}.$$

Consequently,

$$mA - U \geq (m\eta - \tfrac{1}{2}\sqrt{\rho})n - 1 \geq \frac{m\eta n}{2} - 1 \geq D.$$

Hence $U \leq mA - D$ and $m \leq \lfloor (U-1)/d \rfloor \leq M$, and therefore

$$\sum_{j=0}^{M} (mA - D - dj)(D - j + 1) > n \sum_{s=0}^{\min(m-1,M)} (m - s)(D - s + 1). \tag{4.5}$$

There are therefore more unknowns than equations, whose number is bounded by (4.4), so a nonzero solution Q exists.

Fix such a Q. The last step is to show that $\deg_Z Q \geq 1$. To see this, note that if $\deg_Z Q = 0$, then $(X - z)^m$ divides Q for every $z \in \mathcal{D}$ by the multiplicity conditions, and this contradicts $\deg_X Q < mA \leq mn$. Thus $\deg_Z Q \geq 1$ and the polynomial Q satisfies all of the conditions of the lemma. □

We fix a polynomial Q provided by Lemma 4.2 for the rest of this section. As mentioned, our goal henceforth reduces to a root counting problem on Q and we now describe in more detail how weo do so using a partitioning of the. Specifically, we will want to bound the number of coefficient pairs (α, β) such that there exists a low degree P satisfying $Q(X, \alpha + \beta X, P(X)) = 0$ and (α, β, P) is exceptional.

We factor Q as

$$Q = C(X, Y) \prod_{i=1}^{r} Q_i(X, Y, Z)^{a_i},$$

where the Q_i are distinct irreducible factors of positive Z-degree. For each i, let e_i be the largest power of the field characteristic dividing every positive Z-exponent of Q_i. Note that we can find R_i such that

$$Q_i(X^{e_i}, Y^{e_i}, Z) = R_i(X, Y, Z)^{e_i}.$$

Henceforth, we fix these Q_i, R_i, and e_i. Define $h_i \coloneqq \deg_{Y,Z} R_i$. Note that

$$h_i \geq 1, \qquad e_i \leq M, \qquad \deg_Y C + \sum_i h_i \leq D. \tag{4.6}$$

The first bound is trivial, the second bound is because e_i divides every positive Z-exponent of Q_i and Q has Z-degree at most M. The third bound is because the Y, Z degree of Q is at most D.

Denote $\partial R_i / \partial Z$ by $R_{i,Z}$. If $Q_i(X, \alpha + \beta X, P(X)) = 0$, then, after substituting X^{e_i} for X we can obtain

$$R_i(X, \alpha^{1/e_i} + \beta^{1/e_i} X, P(X^{e_i})) = 0. \tag{4.7}$$

Now, define the excluded coefficient pairs by

$$\mathcal{E} \coloneqq \left\{ (\alpha, \beta) \in \mathbb{F}_q^2 : \begin{array}{l} C(X, \alpha + \beta X) \equiv 0, \text{ or} \\ R_i(X, \alpha^{1/e_i} + \beta^{1/e_i} X, Z) \text{ has a repeated} \\ \text{root in } Z \text{ for some } i \end{array} \right\}, \tag{4.8}$$

where, above, roots are taken over $\overline{\mathbb{F}_q(X)}$.

Set the intermediate agreement threshold

$$B \coloneqq \lceil (\sqrt{\rho} + \eta/2)n \rceil - 1.$$

We call a triple $(\alpha, \beta, P) \in \mathbb{F}_q^2 \times \mathbb{F}_q[X]$, with $\deg P \le d$, *persistent* if its agreement on a set of at least B coordinates has a "rational explanation" in the following sense. There exist a set $I \subseteq S_{\alpha,\beta}(P)$ with $|I| \ge B$, rational functions $F, G \in \mathbb{F}_q(X)$ such that $P = F + (\alpha + \beta X)G$ and

$$F(z) = f(z) \text{ and } G(z) = g(z) \quad \text{for all } z \in I.$$

where $G = H/L$, $\deg H \le d$, $L \neq 0$, and $\deg L \le 1$. We say that F, G *explains* the persistence of (α, β, P).

We call the triple *nonpersistent* otherwise. Denote the persistent and nonpersistent exceptional coefficient pairs outside $\mathcal{E}$ by

$$\mathcal{P} \coloneqq \left\{ (\alpha, \beta) \in \mathbb{F}_q^2 \setminus \mathcal{E} : \begin{array}{l} \exists P \text{ such that } (\alpha, \beta, P) \text{ is exceptional} \\ \text{and persistent} \end{array} \right\}, \tag{4.9}$$

$$\mathcal{N} \coloneqq \left\{ (\alpha, \beta) \in \mathbb{F}_q^2 \setminus \mathcal{E} : \begin{array}{l} \exists P \text{ such that } (\alpha, \beta, P) \text{ is exceptional} \\ \text{and nonpersistent} \end{array} \right\}. \tag{4.10}$$

Every coefficient pair occurring in an exceptional triple lies in $\mathcal{E} \cup \mathcal{P} \cup \mathcal{N}$ hence it suffices for us to bound the sizes of each of these three sets.

4.2 Bounding excluded coefficient pairs

We first bound $|\mathcal{E}|$ and also show a fact about roots outside of $\mathcal{E}$ which will be useful later.

Lemma 4.3. *We have*

$$|\mathcal{E}| \le D^2.$$

Furthermore, for every $(\alpha, \beta) \notin \mathcal{E}$ and every $P \in \mathbb{F}_q[X]$ of degree at most d satisfying

$$Q(X, \alpha + \beta X, P(X)) = 0,$$

there is some $i \in [r]$ such that

$$R_i(X, \alpha^{1/e_i} + \beta^{1/e_i} X, P(X^{e_i})) = 0,$$
$$R_{i,Z}(X, \alpha^{1/e_i} + \beta^{1/e_i} X, P(X^{e_i})) \not\equiv 0.$$

Proof. For any nonzero $G(X, Y)$,

$$G(X, \alpha + \beta X) = 0 \quad \iff \quad Y - \alpha - \beta X \text{ divides } G(X, Y).$$

Distinct pairs correspond to distinct factors of Y-degree one, so the restriction vanishes identically for at most $\deg_Y G$ pairs. Applying this to C, at most $\deg_Y C$ pairs satisfy the first condition defining $\mathcal{E}$.

For the second condition, fix i. The specialization $R_i(X, \alpha^{1/e_i} + \beta^{1/e_i} X, Z)$ is nonzero as otherwise $Y - \alpha^{1/e_i} - \beta^{1/e_i} X$ would divide the irreducible polynomial R_i, which has positive Z-degree. If the specialization has a repeated root, then its Z-derivative has the same root. By the polynomial Bézout identity,

$$\operatorname{Res}_Z(R_i, R_{i,Z})|_{Y = \alpha^{1/e_i} + \beta^{1/e_i} X} \equiv 0.$$

This resultant is nonzero because R_i is irreducible and $R_{i,Z} \neq 0$. Since R_i and $R_{i,Z}$ have total (Y, Z)-degrees at most h_i and $h_i - 1$ respectively, the Sylvester determinant has Y-degree at most

$h_i(h_i-1)$. Hence, there are at most $h_i(h_i-1)$ pairs satisfying the second condition for this i and, summing over all i, we obtain

$$\begin{aligned}|\mathcal{E}| &\le \deg_Y C + \sum_i h_i(h_i-1)\\ &\le \left(\deg_Y C + \sum_i h_i\right)^2 \le D^2,\end{aligned}$$

where the last inequality uses (4.6). This completes the bound on $|\mathcal{E}|$.

We move on to the second part of the lemma. Fix $(\alpha,\beta)\notin\mathcal{E}$ and a polynomial P satisfying (4.2). Substituting into the factorization of Q, we obtain

$$0 = Q(X,\alpha+\beta X,P(X)) = C(X,\alpha+\beta X)\prod_i Q_i(X,\alpha+\beta X,P(X))^{a_i}.$$

Since $(\alpha,\beta)\notin\mathcal{E}$, $C(X,\alpha+\beta X)$ is nonzero and hence for at least one i we have

$$R_i(X,\alpha^{1/e_i}+\beta^{1/e_i}X,P(X^{e_i}))\equiv 0,$$

and moreover $R_i(X,\alpha^{1/e_i}+\beta^{1/e_i}X,Z)$ has no repeated roots in Z. This means its root $Z=P(X^{e_i})$ is simple and cannot also be a root of its Z-derivative, so

$$R_{i,Z}(X,\alpha^{1/e_i}+\beta^{1/e_i}X,P(X^{e_i}))\not\equiv 0.$$

□

4.3 Bounding persistent coefficient pairs

Towards bounding $|\mathcal{P}|$, we first need the following fact. It bounds the number of exceptional coefficient pairs compatible with one fixed rational pair F,G.

Lemma 4.4. *For every pair of rational functions $F,G\in\mathbb{F}_q(X)$, the number of coefficient pairs $(\alpha,\beta)\in\mathbb{F}_q^2$ for which $P\coloneqq F+(\alpha+\beta X)G$ is a polynomial and (α,β,P) is exceptional is at most*

$$\begin{cases}\binom{n}{2}, & \text{if } F,G\in\mathbb{F}_q[X],\ \deg F\le d,\ \deg G\le d-1,\\ q, & \text{otherwise.}\end{cases}$$

Proof. If F,G are not both polynomials with $\deg F\le d$ and $\deg G\le d-1$, the pairs (α,β) for which $F+(\alpha+\beta X)G$ is a polynomial of degree at most d form an affine subspace of $\mathbb{F}_q^2$, if nonempty. This subspace cannot contain $\mathbb{F}_q^2$ as otherwise the pairs $(0,0),(1,0),(0,1)$ would force F,G,XG to all be polynomials of degree at most d, which is the opposite of what we assumed. There are thus at most q pairs in this case, taking care of the second case above.

Now suppose $F,G\in\mathbb{F}_q[X]$ with $\deg F\le d$ and $\deg G\le d-1$. Fix

$$I'\coloneqq\{z\in\mathcal{D}: F(z)=f(z) \text{ and } G(z)=g(z)\},$$

to be the set of points where F,G have correlated agreement with f,g respectively.

Every exceptional coefficient pair has at least two agreement coordinates outside I'; otherwise F,G would be admissible witnesses for P with agreement loss at most 1.

Suppose two coefficient pairs (α, β) and (α', β'), with candidates $P = F + (\alpha + \beta X)G$ and $P' = F + (\alpha' + \beta' X)G$, share two distinct agreement coordinates $z, z' \notin I'$. At either coordinate $x \in \{z, z'\}$, we have

$$(f(x) - F(x)) + (\alpha + \beta x)(g(x) - G(x)) = 0. \tag{4.11}$$

Here $g(x) \neq G(x)$: otherwise this equation would also give $f(x) = F(x)$, contradicting $x \notin I'$. Subtracting the corresponding equation for (α', β') and dividing by $g(x) - G(x)$ gives

$$(\alpha - \alpha') + (\beta - \beta')z = 0, \qquad (\alpha - \alpha') + (\beta - \beta')z' = 0.$$

Since $z \neq z'$, these equations imply $(\alpha, \beta) = (\alpha', \beta')$. Thus each exceptional coefficient pair has a pair of agreement coordinates outside I' that no other coefficient pair shares. There are at most $\binom{n}{2}$ choices for this unordered pair of coordinates, proving the first case of the lemma. □

Define

$$\ell \coloneqq \frac{n(B - d - 1)}{B^2 - n(d + 1)},$$

so that ℓ is the upper bound from Lemma 3.2, on the size of any family of subsets of $\mathcal{D}$, each of size at least B and with pairwise intersections of size at most $d + 1$.

Lemma 4.5. *We have*

$$|\mathcal{P}| \leq \ell \max\{q, \tbinom{n}{2}\}.$$

Proof. We first bound the number of possible agreement sets. For each rational pair F, G explaining the persistence of an exceptional triple, consider, as in the proof of Lemma 4.4, set

$$I \coloneqq \{z \in \mathcal{D} : F(z) = f(z) \text{ and } G(z) = g(z)\},$$

to be set of points where F, G have correlated agreement with f, g respectively.

Let (F_0, G_0) and (F_1, G_1) be two distinct such pairs, with correlated agreement sets I_0, I_1. For each $j = 0, 1$, choose a corresponding exceptional triple (α_j, β_j, P_j) and a representation $G_j = H_j / L_j$. At every $z \in I_0 \cap I_1$, we have

$$F_0(z) = F_1(z) = f(z), \qquad G_0(z) = G_1(z) = g(z).$$

If $G_0 \neq G_1$, then the polynomial

$$H_0 L_1 - H_1 L_0$$

is nonzero, has degree at most $d + 1$, and vanishes on every such z. If $G_0 = G_1$, then $F_0 \neq F_1$, so the polynomial

$$(P_0 - P_1)L_0 - (\alpha_0 - \alpha_1 + (\beta_0 - \beta_1)X)H_0$$

is nonzero, has degree at most $d+1$, and again vanishes on every such z. In either case, $|I_0 \cap I_1| \leq d+1$. Since each correlated agreement set has size at least $B > d + 1$, it determines a unique rational pair among those considered above. By Lemma 3.2, with $a = B/n$ and $b = (d + 1)/n$, there are at most ℓ distinct correlated agreement sets.

Finally, by Lemma 4.4, the rational pair determined by each correlated agreement set accounts for at most $\max\{q, \binom{n}{2}\}$ exceptional coefficient pairs. Every coefficient pair in $\mathcal{P}$ has an exceptional triple whose persistence is explained by one of these rational pairs, so

$$|\mathcal{P}| \leq \ell \max\{q, \tbinom{n}{2}\}.$$

□

4.4 Bounding nonpersistent coefficient pairs

We now bound $|\mathcal{N}|$. It will be helpful to define $v \coloneqq \frac{n}{A-B}$. Let T_0 be an indeterminate and let $\mathbb{K} \coloneqq \overline{\mathbb{F}_q(T_0)}$ be the algebraic closure of $\mathbb{F}_q(T_0)$. Since T_0 is transcendental, a nonzero polynomial in $\mathbb{F}_q[X]$ must be nonzero when evaluated at T_0.

Definition 4.6. We identify triples (α, β, P) with coefficient vectors in

$$\mathcal{A} \coloneqq \mathbb{K}^{d+3}, \qquad (\alpha, \beta, P) \longleftrightarrow (\alpha, \beta, p_0, \ldots, p_d),$$

where $P(X) = \sum_{j=0}^{d} p_j X^j$. The triples we count lie in $\mathbb{F}_q^{d+3} \subseteq \mathcal{A}$; the algebraic sets containing them are defined over $\mathbb{K}$.

Lemma 4.7. *We have*

$$|\mathcal{N}| \le 4d^2M^2D^3v^2.$$

For each $z \in \mathcal{D}$, define the agreement hyperplane

$$\mathcal{H}_z \coloneqq \left\{ (\alpha, \beta, p_0, \ldots, p_d) \in \mathcal{A} : \sum_{j=0}^{d} p_j z^j = f(z) + (\alpha + \beta z) g(z) \right\}.$$

Our first step is to adapt the induction in [DKT26, proof of Lemma 4.1] to the affine case we consider. Given an algebraic set of bounded dimension, we bound the number of nonpersistent triples (α, β, P) it contains.

Claim 4.8. *For every integer $r \ge 0$, an algebraic set $V \subseteq \mathcal{A}$ of dimension at most r contains at most $(\deg V)v^r$ nonpersistent triples $(\alpha, \beta, P) \in \mathbb{F}_q^{d+3}$ with $|S_{\alpha,\beta}(P)| \ge A$.*

Proof. We first show the result for an irreducible component V_0. We proceed by induction on $t \coloneqq \dim V_0$. If $t = 0$, then V_0 is a point and $\deg V_0 = 1$, so the bound holds. Suppose $t > 0$, and define

$$c \coloneqq \big|\{z \in \mathcal{D} : V_0 \subseteq \mathcal{H}_z\}\big|,$$

the number of agreement hyperplanes $\mathcal{H}_z$ that contain V_0.

First consider the case where $c > B$. We show that V_0 contains no nonpersistent $\mathbb{F}_q$-triples. Take any $\mathbb{F}_q$-triple $(\alpha, \beta, P) \in V_0$, and define

$$W \coloneqq \bigcap_{z \in \mathcal{D}:\, V_0 \subseteq \mathcal{H}_z} \mathcal{H}_z.$$

Since $V_0 \subseteq W$, we have $\dim W \ge t > 0$. The set W is defined by affine linear equations over $\mathbb{F}_q$ and contains (α, β, P), so it must contain a second $\mathbb{F}_q$-triple $(\alpha', \beta', P') \ne (\alpha, \beta, P)$. Both triples satisfy the same c agreement equations, so

$$P'(z) - P(z) = \big((\alpha' - \alpha) + (\beta' - \beta)z\big) g(z) \qquad \text{whenever } V_0 \subseteq \mathcal{H}_z.$$

Their coefficient pairs, (α, β) and (α', β'), must differ as otherwise $P' - P$ would have degree at most d and vanish at $c > d$ coordinates, forcing the triples to be the same. We can then define

$$G(X) \coloneqq \frac{P'(X) - P(X)}{(\alpha' - \alpha) + (\beta' - \beta)X}, \qquad F(X) \coloneqq P(X) - (\alpha + \beta X) G(X).$$

The polynomial $L \coloneqq (\alpha' - \alpha) + (\beta' - \beta)X$ vanishes on at most one of these c coordinates z where $\mathcal{H}_z \supseteq V_0$. At each remaining coordinate, $G(z) = g(z)$ and, since the triple lies in $\mathcal{H}_z$, also

$F(z) = f(z)$. Thus F, G have correlated agreement with f, g on at least $c - 1 \geq B$ coordinates which makes (α, β, P) persistent. Hence V_0 contains no nonpersistent $\mathbb{F}_q$-triples in this case.

Now suppose $c \leq B$. Let $\mathcal{T}$ be the set of nonpersistent $\mathbb{F}_q$-triples in V_0 with at least A agreements. We bound $|\mathcal{T}|$ by counting the edges of a bipartite graph. Its left vertices are the triples in $\mathcal{T}$, and its right vertices are the $n - c$ coordinates $z \in \mathcal{D}$ for which $V_0 \not\subseteq \mathcal{H}_z$. Add an edge between (α, β, P) and z if

$$P(z) = f(z) + (\alpha + \beta z)g(z),$$

or equivalently, if $(\alpha, \beta, P) \in V_0 \cap \mathcal{H}_z$.

Every left vertex has at least A agreement coordinates. Exactly c of these correspond to hyperplanes containing all of V_0; these coordinates are excluded from the right vertices. Thus every left vertex has at least $A - c$ neighbors, and the graph has at least $(A - c)|\mathcal{T}|$ edges.

For a right vertex z, its neighbors are precisely the triples in $\mathcal{T} \cap \mathcal{H}_z$. By Lemma 3.9, $V_0 \cap \mathcal{H}_z$ has dimension at most $t-1$ and degree at most $\deg V_0$. Applying induction to its irreducible components and summing their bounds, we get that it contains at most $(\deg V_0)v^{t-1}$ nonpersistent $\mathbb{F}_q$-triples with at least A agreements. Thus each right vertex has at most this many neighbors. Since there are $n - c$ right vertices, the graph has at most $(n - c)(\deg V_0)v^{t-1}$ edges.

Comparing these lower and upper bounds on the number of edges gives

$$(A - c)|\mathcal{T}| \leq (n - c)(\deg V_0)v^{t-1}.$$

Since $A - c \geq A - B > 0$ and $v = n/(A - B)$, we obtain

$$|\mathcal{T}| \leq \frac{n - c}{A - c}(\deg V_0)v^{t-1} \leq (\deg V_0)v^t.$$

This establishes the result for irreducible components. For the full lemma, we sum over irreducible components. □

To apply the bound above, we conclude by covering the polynomial roots by algebraic sets of dimension at most two.

Claim 4.9. *There are algebraic sets $V_i \subseteq \mathcal{A}$, for $i \in [r]$, such that every $\mathbb{F}_q$-triple satisfying* (4.2) *with $(\alpha, \beta) \notin \mathcal{E}$ and $\deg P \leq d$ lies in $\bigcup_i V_i$. Furthermore, for each i, we have the following dimension and degree bounds*

$$\dim V_i \leq 2, \qquad \deg V_i \leq 4d^2 e_i^2 h_i^3.$$

Proof. We fix an $i \in [r]$ and write $R \coloneqq R_i$, $e \coloneqq e_i$, $h \coloneqq h_i$, and

$$\mathcal{R}_i \coloneqq \left\{ (\alpha, \beta, P) \in \mathbb{F}_q^2 \times \mathbb{F}_q[X] : \begin{array}{l} (\alpha, \beta) \notin \mathcal{E}, \quad \deg P \leq d, \\ Q_i(X, \alpha + \beta X, P(X)) \equiv 0 \end{array} \right\}. \tag{4.12}$$

Our goal is to cover $\mathcal{R}_i$ by an algebraic set of dimension at most two and bounded degree. Our strategy will be to construct a rational map of bounded degree whose image contains $\mathcal{R}_i$ and then use Lemma 3.10 to bound the dimension and degree of the closure of its image.

To this end, observe that for every triple $(\alpha, \beta, P) \in \mathcal{R}_i$, we can actually reconstruct P from α, β and a single initial value, $P(T_0^e)$, by Hensel lifting.

Indeed, letting $Y = \alpha^{1/e}$, $V = \beta^{1/e}$, and $Z = P(T_0^e)$. By (4.7),

$$R(X, Y + VX, P(X^e)) = 0.$$

At $X = T_0$ this gives $R(T_0, Y + VT_0, Z) = 0$ and let us write $S(Y, V, Z) \coloneqq R_Z(T_0, Y + VT_0, Z)$ for simplicity. Then since $(\alpha, \beta) \notin \mathcal{E}$ we have that $R_Z(X, Y + VX, P(X^e))$ is nonzero over $\mathbb{F}_q$, so $S(Y, V, Z) \neq 0$ when plugging in T_0.

We can thus apply Lemma 3.11 to construct a formal power series E such that

$$P(X^e) - Z = E(X - T_0) = \sum_{j \geq 1} c_j(Y, V, Z)(X - T_0)^j.$$

Since e is a power of the characteristic, $X^e - T_0^e = (X - T_0)^e$. Thus $c_{ej}(Y, V, Z)$ is the coefficient of $(X - T_0^e)^j$ in $P(X)$ for $1 \leq j \leq d$.

We now want to turn the equation above from Hensel lifting into a reconstruction map, which recovers P from α, β, and the initial value $P(T_0^e)$. Formally, we define

$$\Phi(Y, V, Z) \coloneqq \left(Y^e,\ V^e,\ Z + \sum_{j=1}^{d} c_{ej}(Y, V, Z)(X - T_0^e)^j \right),$$

where we think of the domain of the map as the set of triples (Y, V, Z) for which Hensel lifting applies, or

$$\mathcal{U} \coloneqq \{(Y, V, Z) \in \mathbb{K}^3 : R(T_0, Y + VT_0, Z) = 0, \quad S(Y, V, Z) \neq 0\}.$$

Then Φ gives a map from $\mathcal{U}$ to $\mathcal{A}$. When identifying the polynomial in the third coordinate with $\mathcal{A}$, we express it using its $d + 1$ coefficients in the basis $1, X - T_0^e, \ldots, (X - T_0^e)^d$, so

$$\Phi(Y, V, Z) = \Big(Y^e, V^e, Z, \underbrace{c_e(Y, V, Z), \ldots, c_{ed}(Y, V, Z)}_{j=1,\ldots,d} \Big) \in \mathbb{K}^2 \times \mathbb{K}^{d+1}.$$

Summarizing, we have now shown that each triple of interest $(\alpha, \beta, P) \in \mathcal{R}_i$, can be obtained through

$$\Phi(\alpha^{1/e}, \beta^{1/e}, P(T_0^e)) = (\alpha, \beta, P),$$

with $(\alpha^{1/e}, \beta^{1/e}, P(T_0^e)) \in \mathcal{U}$. Hence, $\mathcal{R}_i \subseteq \Phi(\mathcal{U})$.

To conclude, we bound the dimension and degree $\overline{\mathcal{U}}$, show that the map Φ is a rational map with numerators and common denominators of bounded degree, and then apply Lemma 3.10.

First, by Lemma 3.8, $\dim \overline{\mathcal{U}} \leq 2$ and $\deg \overline{\mathcal{U}} \leq h$. Second, for each $1 \leq j \leq d$, Lemma 3.11 gives

$$c_{ej}(Y, V, Z) = \frac{N_{ej}(Y, V, Z)}{S(Y, V, Z)^{2ej-1}}, \qquad \deg N_{ej} \leq 1 + (2ej - 1)(h - 1).$$

We write every coordinate of Φ with common denominator S^{2ed-1}. This yields,

$$\frac{1}{S^{2ed-1}} \left(S^{2ed-1} Y^e,\ S^{2ed-1} V^e,\ S^{2ed-1} Z,\ (S^{2e(d-j)} N_{ej})_{j=1}^{d} \right),$$

for the last d coordinates of Φ. Since $\deg S \leq h - 1$, the denominator has degree at most $(2ed-1)(h-1)$. The numerator degrees are at most $e + (2ed-1)(h-1)$ for the first two coordinates and $1 + (2ed - 1)(h - 1)$ for the polynomial coefficients.

Finally, we let

$$V_i \coloneqq \overline{\Phi(\mathcal{U})}.$$

Applying Lemma 3.10 to $\overline{\mathcal{U}}$ with numerator and denominator degrees at most $2edh$ gives

$$\dim V_i \leq 2, \qquad \deg V_i \leq h(2edh)^2 = 4d^2e^2h^3,$$

and $\mathcal{R}_i \subseteq V_i$ for each $i \in [r]$. Every triple in the claim belongs to some $\mathcal{R}_i$ by Lemma 4.3, so the V_i's give a covering of all triples in the claim. $\square$

Proof of Lemma 4.7. We prove the same bound for all nonpersistent triples (α, β, P) with $(\alpha, \beta) \notin \mathcal{E}$, $\deg P \le d$, and $|S_{\alpha,\beta}(P)| \ge A$. The same bound holds for $|\mathcal{N}|$, since each of its coefficient pairs occurs in at least one such triple.

By Lemma 4.2, every triple being counted satisfies (4.2) and, by Claim 4.9, belongs to $\bigcup_i V_i$. Apply Claim 4.8 to each V_i with the hyperplanes $\mathcal{H}_z$. Summing over factors, using $e_i \le M$ and $\sum_i h_i^3 \le D^3$ from (4.6), we bound the number of nonpersistent triples by

$$4d^2v^2 \sum_i e_i^2 h_i^3 \le 4d^2M^2D^3v^2.$$

□

Having bounded $|\mathcal{E}|$, $|\mathcal{P}|$, and $|\mathcal{N}|$, we can now prove Lemma 4.1.

Proof of Lemma 4.1. The parameter definitions give

$$D \le \frac{9}{2\rho\eta^2}, \qquad dM \le \frac{3n}{\eta}, \qquad v \le \frac{2}{\eta}.$$

Also, by (3.1),

$$\frac{B}{n} - \sqrt{\frac{d+1}{n}} \ge \frac{\eta}{2} - \frac{1}{n} - \frac{1}{2n\sqrt{\rho}} \ge \frac{\eta}{4},$$

so

$$\ell \le \frac{1}{B/n - \sqrt{(d+1)/n}} \le \frac{4}{\eta}.$$

Every coefficient pair counted in the lemma belongs to $\mathcal{E} \cup \mathcal{P} \cup \mathcal{N}$. Substituting these estimates into the bounds from Lemmas 4.3, 4.5 and 4.7, and using $\max\{q, \binom{n}{2}\} \le q + n^2/2$, bounds their total number by

$$\begin{aligned} |\mathcal{E}| + |\mathcal{P}| + |\mathcal{N}| &\le D^2 + \ell(q + n^2/2) + 4(dM)^2D^3v^2 \\ &\le \frac{4q}{\eta} + \frac{2^{14}n^2}{\rho^3\eta^{10}}. \end{aligned}$$

□

Theorem 1.1 gives a simplified bound when $n\eta^2 \ge 4/\rho$. This condition may fail at low rates. In that case, one can choose m, D directly and use the finite counts in the proof of Lemma 4.1.

Theorem 4.10. *Set $0 < \eta < 1 - \sqrt{\rho}$ and for positive integers m, D, set M by* (4.1)*, and suppose that $M \ge 1$, the interpolation inequality* (4.5) *holds, and $B^2 > n(d+1)$. Then the MCA error with one-coordinate agreement loss satisfies*

$$\begin{aligned} \varepsilon_{\text{aff}} \le{}& \frac{\left\lfloor \dfrac{n(B-d-1)}{B^2 - n(d+1)} \right\rfloor \max\{q, \binom{n}{2}\}}{q^2} \\ &+ \frac{D^2 + 4d^2M^2D^3 \left(\frac{n}{A-B}\right)^2}{q^2}. \end{aligned} \tag{4.13}$$

Proof. The interpolation inequality ensures that the construction in Lemma 4.2 produces a nonzero explainer for these choices of m, D. The condition $B^2 > n(d+1)$ allows the Johnson bound in Lemma 4.5 to be applied. The remaining root-counting arguments use the degree bounds on the explainer. Summing the bounds from Lemmas 4.3, 4.5 and 4.7 gives (4.13). □

When $q \geq \binom{n}{2}$, the first term is simply the rounded Johnson bound divided by q. For example, with $n = 2^{15}$, $\rho = 1/2048$, $\eta = 0.05$, and $q = 2^{128}$, we have $d = 16$, $A = 2363$, and $B = 1543$. Choosing $m = 1$ and $D = M = 28$ satisfies the interpolation inequality and gives

$$\varepsilon_{\text{aff}} \leq 27 \cdot 2^{-128} + 2^{-211} \approx 2^{-123.25},$$

even though the condition of Theorem 1.1 fails.

5 Adjustments for Equal-degree, linear constraints, and interleaving

For applications to proof systems, we need to modify the affine combiner to handle a few additional requirements, namely equal degree bounds, linear constraints, and higher-arity interleavings. We give adaptations that attain these properties without substantially compromising the MCA error from Theorem 1.1.

Equal degree bounds. In most IOPPs, it is convenient for the two received words in a fold to have the same degree bound. To handle this, we can modify the combiner so that both words have affine coefficients. Sample independent uniform $\theta, \alpha, \beta \in \mathbb{F}_q$ and define

$$\begin{aligned} w_{\theta,\alpha,\beta}(z) &\coloneqq g(z) + (\alpha + \beta z)(f(z) + \theta g(z)) \\ &= (\alpha + \beta z) f(z) + (1 + \theta\alpha + \theta\beta z) g(z). \end{aligned} \tag{5.1}$$

For received words $f, g : \mathcal{D} \to \mathbb{F}_q$, define

$$S_{\theta,\alpha,\beta}(P) \coloneqq \{x \in \mathcal{D} : P(x) = w_{\theta,\alpha,\beta}(x)\}.$$

For a fixed choice of (θ, α, β), say polynomials $F, G \in \mathbb{F}_q[X]$ are admissible witnesses for P with agreement loss at most s if

$$\begin{aligned} &\deg F, \deg G \le d - 1, \qquad P = G + (\alpha + \beta X)(F + \theta G), \\ &\left|\{x \in S_{\theta,\alpha,\beta}(P) : F(x) \ne f(x) \text{ or } G(x) \ne g(x)\}\right| \le s. \end{aligned} \tag{5.2}$$

We say that the equal-degree combiner has MCA error at most ε with agreement loss at most s on $\mathcal{D}$ if, for every pair of received words $f, g : \mathcal{D} \to \mathbb{F}_q$,

$$\Pr_{\theta,\alpha,\beta \leftarrow \mathbb{F}_q} \left[\begin{array}{l} \exists P \in \mathbb{F}_q[X],\ \deg P \le d : \\ |S_{\theta,\alpha,\beta}(P)| \ge A \text{ and } P \text{ has no} \\ \text{admissible witnesses with agreement loss at most } s \end{array} \right] \le \varepsilon.$$

Similar to our notation for MCA error for affine combiners, we denote the smallest such ε by $\varepsilon_{\mathrm{eq}}(n, \rho, A, s, q)$.

Linear constraints. In a typical proof system, the prover sends, along with the received words f, g, linear claims of the form

$$\langle m_f, v \rangle = \gamma_0, \qquad \langle m_g, v \rangle = \gamma_1,$$

where m_f, m_g are the messages that underlie the received words f, g, and $v \in \mathbb{F}_q^{d+1}$ is a fixed public weight vector. In the case of Reed-Solomon codes, m_f, m_g might be the d coefficients of the degree $< d$ polynomial equal to f, g respectively. That is, suppose that for honest degree $< d$ polynomials F, G, we have $m_f = \mathrm{coef}(F)$ and $m_g = \mathrm{coef}(G)$, where $\mathrm{coef}(U)$ denotes the coefficient vector of U in the monomial basis $(1, X, \ldots, X^d)$.

We want to carry these claims to a linear constraint on the coefficients of the

$$P = G + \alpha(F + \theta G) + \beta X(F + \theta G).$$

The first two terms contribute $\gamma_1 + \alpha(\gamma_0 + \theta\gamma_1)$ to its inner product with v. The remaining term requires a claim on the message of $X(F + \theta G)$, which the two original claims do not determine. After receiving θ, the prover therefore sends one field element c, intended to equal

$$\langle \mathrm{coef}(X(F + \theta G)), v \rangle.$$

Only then are α, β sampled. Writing $m_P = \mathrm{coef}(P)$, the new claim on m_P is

$$\langle m_P, v \rangle = \gamma_1 + \alpha(\gamma_0 + \theta\gamma_1) + \beta c. \tag{5.3}$$

For honest messages, the claim holds by linearity.

We show that our equal degree combiners can handle linear claims in the sense above with small additional proximity loss.

Theorem 5.1. *For the combiner* (5.1)*, the MCA error with agreement loss $s = 1$ satisfies*

$$\begin{aligned} \varepsilon_{\mathrm{eq}}\Big(\underbrace{n}_{\substack{\text{block}\\ \text{length}}}, \underbrace{\rho}_{\substack{\text{degree}\\ \text{fraction}}}, \underbrace{A}_{\substack{\text{agreement}\\ \text{threshold}}}, \underbrace{1}_{\substack{\text{agreement}\\ \text{loss}}}, \underbrace{q}_{\substack{\text{field}\\ \text{size}}}\Big) &\le \varepsilon_{\mathrm{aff}} + \frac{\ell_0}{q} \\ &\le \frac{4/\eta + \ell_0}{q} + \frac{2^{14} n^2}{\rho^3 \eta^{10} q^2}, \end{aligned} \tag{5.4}$$

where $\ell_0 = \lfloor n(A - 1 - d)/((A - 1)^2 - nd) \rfloor$.

Moreover, for the combination with linear claims as described above, the MCA guarantee implies that, except with probability at most

$$\varepsilon_{\mathrm{aff}} + \frac{2\ell_0}{q} \le \frac{4/\eta + 2\ell_0}{q} + \frac{2^{14} n^2}{\rho^3 \eta^{10} q^2}, \tag{5.5}$$

every candidate P of degree at most d with $|S_{\theta,\alpha,\beta}(P)| \ge A$ satisfying (5.3) *has admissible witnesses F, G with agreement loss at most one satisfying*

$$\begin{gathered} \langle \mathrm{coef}(F), v \rangle = \gamma_0, \qquad \langle \mathrm{coef}(G), v \rangle = \gamma_1, \\ \langle \mathrm{coef}(X(F + \theta G)), v \rangle = c. \end{gathered} \tag{5.6}$$

This holds for every prover strategy.

Proof. Consider the list

$$\mathcal{L} \coloneqq \left\{ (F, G) \in \mathbb{F}_q[X]^2_{\le d} : |\{z \in \mathcal{D} : F(z) = f(z),\ G(z) = g(z)\}| \ge A - 1 \right\}.$$

By (3.1), we have $n\eta > 1$, so $A - 1 > \sqrt{nd}$ and the denominator defining ℓ_0 is positive. Two distinct pairs have at most d common agreement coordinates, since at least one of their polynomial differences is nonzero and has degree at most d. Lemma 3.2, with $a = (A - 1)/n$ and $b = d/n$, gives $|\mathcal{L}| \le \ell_0$. This list is independent of all challenges.

For each fixed θ, apply Theorem 1.1 to the received words $(g, f + \theta g)$. Except with probability $\varepsilon_{\mathrm{aff}}$ over (α, β), every candidate under consideration has witnesses U, V with $\deg U \le d$, $\deg V \le d - 1$, and $P = U + (\alpha + \beta X)V$, explaining all but one of its agreements. Set $G = U$ and $F = V - \theta U$. Then $(F, G) \in \mathcal{L}$, the combination identity and agreement guarantee in (5.2) hold, and

$$[X^d]F + \theta [X^d]G = 0.$$

For each pair in $\mathcal{L}$ with $\deg G = d$, this equation determines exactly one θ. With probability at most ℓ_0/q over θ, any such pair can occur. Outside these choices, $\deg G \le d-1$ and $F = V - \theta G$ also has degree at most $d-1$. This proves the first assertion.

For the claim guarantee, condition on θ and the prover's message c. Fix a pair $(F, G) \in \mathcal{L}$ with both degrees at most $d-1$, and define

$$\Delta_F = \langle \mathrm{coef}(F), v\rangle - \gamma_0, \qquad \Delta_G = \langle \mathrm{coef}(G), v\rangle - \gamma_1,$$
$$\Delta_X = \langle \mathrm{coef}(X(F+\theta G)), v\rangle - c.$$

For its combined polynomial, the residual in (5.3) is

$$\Delta_G + \alpha(\Delta_F + \theta\Delta_G) + \beta\Delta_X.$$

If any equality in (5.6) fails, this is a nonzero affine polynomial in α, β and vanishes with probability at most $1/q$. A union bound over the at most ℓ_0 pairs, followed by the two exceptional events above, proves (5.5). □

Since $\ell_0 = O_{\rho,\eta}(1)$, both guarantees retain error $O_{\rho,\eta}(1/q + n^2/q^2)$.

Round complexity. The schedules above introduce one extra interaction round per fold. This is done mainly for simplicity. An optimization removes this round at the cost of one additional field element per functional, while preserving the soundness bounds for the original claims.

The extra round arises because the prover's auxiliary claims are sent after θ and before the remaining challenges. For a single functional, the prover can instead announce claims $c_F = \langle \mathrm{coef}(XF), v\rangle$ and $c_G = \langle \mathrm{coef}(XG), v\rangle$ alongside γ_0, γ_1, before any challenges are sampled. The verifier then samples θ, α, β together and uses $c = c_F + \theta c_G$ in (5.3). The bound in (5.5) still protects the original claims $\langle \mathrm{coef}(F), v\rangle = \gamma_0$ and $\langle \mathrm{coef}(G), v\rangle = \gamma_1$: conditioned on θ, the derived c is fixed before α, β. This guarantee need not certify c_F, c_G individually.

Higher-arity interleavings. We next fold many received words using shared challenges, requiring agreement on a domain shared by all words. Start with received words $f_1, \dots, f_k : \mathcal{D} \to \mathbb{F}_q$, where $k \ge 2$ is a power of two, and set $r := \log_2 k$. Sample independent uniform triples $(\theta_j, \alpha_j, \beta_j) \in \mathbb{F}_q^3$ for $1 \le j \le r$. Starting with $f_i^{(0)} = f_i$, define round j by

$$f_i^{(j)}(z) = f_{2i}^{(j-1)}(z) + (\alpha_j + \beta_j z)\big(f_{2i-1}^{(j-1)}(z) + \theta_j f_{2i}^{(j-1)}(z)\big), \qquad 1 \le i \le k/2^j.$$

Thus each round halves the number of rows, with the same challenges for every pair. Denote the final word by $w := f_1^{(r)}$. Assume Theorem 5.1 applies in round j with degree bound $d + j$ and agreement threshold A, and let ε be the maximum of its MCA bounds (5.4) over the r rounds.

Theorem 5.2. *Except with probability at most $r\varepsilon$, every polynomial P of degree at most $d + r$ with an agreement domain S of size at least $A + r$ with w has polynomial witnesses $F_1, \dots, F_k$ of degree at most d whose folding with the same challenges gives P and for which*

$$\big|\{z \in S : \exists i,\ F_i(z) \ne f_i(z)\}\big| \le r.$$

Proof. We first prove the one-round guarantee for t simultaneous combinations. Fix received words $f_i, g_i : \mathcal{D} \to \mathbb{F}_q$ for $1 \le i \le t$, and use the same uniform $\alpha, \beta \in \mathbb{F}_q$ in all combinations $f_i(z) + (\alpha + \beta z) g_i(z)$.

For a tuple $P = (P_1, \dots, P_t)$, define its shared agreement domain by

$$S_{\alpha,\beta}(P) \coloneqq \{z \in \mathcal{D} : P_i(z) = f_i(z) + (\alpha + \beta z) g_i(z) \text{ for every } i\}.$$

Tuples $F = (F_1, \dots, F_t)$ and $G = (G_1, \dots, G_t)$ of polynomials over $\mathbb{F}_q$ are admissible witnesses for P if

$$\deg F_i \le d, \quad \deg G_i \le d-1, \quad P_i = F_i + (\alpha + \beta X) G_i \qquad (1 \le i \le t),$$
$$\left|\{z \in S_{\alpha,\beta}(P) : \exists i,\ F_i(z) \ne f_i(z) \text{ or } G_i(z) \ne g_i(z)\}\right| \le 1.$$

Thus the agreement loss of one coordinate is shared by all rows.

We establish the one-round bound

$$\Pr_{\alpha,\beta \leftarrow \mathbb{F}_q} \left[\begin{array}{l} \exists P \in (\mathbb{F}_q[X]_{\le d})^t : \\ |S_{\alpha,\beta}(P)| \ge A \text{ and } P \text{ has no} \\ \text{admissible witness tuples} \end{array} \right] \le \frac{4}{\eta q} + \frac{2^{14} n^2}{\rho^3 \eta^{10} q^2}. \tag{5.7}$$

Fix an $\mathbb{F}_q$-basis $\zeta_1, \dots, \zeta_t$ of $\mathbb{F}_{q^t}$. Pack each tuple into one $\mathbb{F}_{q^t}$-valued word or polynomial by taking its linear combination with coefficients ζ_i. Since $\mathcal{D} \subseteq \mathbb{F}_q$ and the challenges lie in $\mathbb{F}_q$, packing commutes with evaluation and with the combiner. Equality of packed symbols is equivalent to equality in every row.

We show that the proof of Theorem 1.1 gives the same bound for the packed words, with polynomial coefficients in $\mathbb{F}_{q^t}$ and challenges in $\mathbb{F}_q$. Interpolation, factorization, simple-root lifting, and the algebraic degree bounds hold over $\mathbb{F}_{q^t}$ with the same parameters. The excluded coefficient pairs are still counted by line factors; restricting the challenges to $\mathbb{F}_q^2$ cannot increase their number. Taking e_i-th roots also preserves $\mathbb{F}_q$.

For persistent triples, the Johnson bound counts shared agreement sets and is unchanged by extending the coefficient field. A rational pair over $\mathbb{F}_{q^t}$ that fails the polynomial degree bounds imposes a nonzero affine equation on α, β. Even with coefficients in $\mathbb{F}_{q^t}$, such an equation has at most q solutions in $\mathbb{F}_q^2$: fixing one coordinate determines the other whenever its coefficient is nonzero. For a pair satisfying the degree bounds, two accidental agreement coordinates determine at most one challenge, giving the same $\binom{n}{2}$ bound.

The nonpersistent bound in Claim 4.8 counts all $\mathbb{F}_{q^t}$-rational triples in the algebraic cover, so it also bounds the subset whose coefficient pairs lie in $\mathbb{F}_q^2$. In its case $c > B$, the affine linear system defining W is now over $\mathbb{F}_{q^t}$. If it contains a triple and has positive dimension, it contains a second $\mathbb{F}_{q^t}$-rational triple; the same subtraction gives a persistent rational explanation over $\mathbb{F}_{q^t}$. The second triple need not have its coefficient pair in $\mathbb{F}_q^2$. The hyperplane-section induction and the reconstruction cover therefore retain their bounds.

Consequently, the number of exceptional base-field challenges is at most

$$D^2 + \ell \max\{q, \tbinom{n}{2}\} + 4d^2 M^2 D^3 v^2.$$

The estimates in the proof of Lemma 4.1, followed by division by q^2, give (5.7). Unpacking the witness polynomials gives the required tuples, with at most one discarded coordinate shared by all rows.

For the equal-degree and linear-claim conclusions, apply the same packing to the witness lists, claimed values, and auxiliary messages. Distinct packed polynomial pairs have at most d common agreement coordinates, so their list has size at most ℓ_0. The degree test in θ has at most one base-field solution for each pair that fails the input degree bounds. The affine zero counts in the claim-binding arguments are unchanged: each nonzero equation over $\mathbb{F}_{q^t}$ has at most q base-field pairs. Common weight vectors make the inner products commute with packing. Thus the proof of Theorem 5.1 gives its stated bounds without a factor of t.

Now apply the interleaved equal-degree bound in each of the r rounds, starting with $t = k/2$ pairs. Conditioning on the preceding rounds fixes the current received words before the fresh challenges are sampled, so each round fails with probability at most ε. A union bound gives total failure probability at most $r\varepsilon$. Outside these failures, work backward from P. Each round recovers polynomials for all parent rows and discards at most one shared agreement coordinate. Since $|S| \geq A + r$, every intermediate tuple retains at least A agreements, so the one-round guarantee continues to apply. After r steps, the recovered polynomials have degree at most d and agree with all original words on all but at most r coordinates of S. Their combination identities compose to give P. □

Remark 5.3. Flock uses Ligerito [NA25] to encode each message row with the same Reed–Solomon code and fold the encoded rows by multilinear partial evaluation [BRW26, Appendices C.1–C.2]. Each scalar folding round replaces every pair u, v by $(1 - \lambda)u + \lambda v$, using the same fresh uniform λ for all pairs. Theorem 5.2 gives the MCA guarantee for an affine variant of this row folding, with common-weight linear constraints handled as above.

We now spell out the MCA error and the grinding needed to reach a 128-bit target for this contribution. When $q \geq \binom{n}{2}$, (4.13) gives

$$\varepsilon_{\text{aff}} \leq \frac{\ell}{q} + \delta, \qquad \delta := \frac{D^2 + 4d^2M^2D^3\big(n/(A-B)\big)^2}{q^2}.$$

For $n = 2^{20}$, $d = 2^{18}$, $\eta = 0.05$, $A = 576717$, and $q = 2^{128}$, the choices $m = 3$, $D = 49$, $M = 6$, and $B = 567980$ satisfy the interpolation dimension inequality and give $\ell \leq 6$, $\ell_0 \leq 5$, and $\delta < 2^{-182}$. Thus the original affine combiner has MCA error at most $6 \cdot 2^{-128} + \delta$, or approximately 125.42 bits of error ($-\log_2$ of the bound). For equal degree bounds, (5.4) adds at most $5 \cdot 2^{-128}$, giving error at most $11 \cdot 2^{-128} + \delta$, or about 124.54 bits. Carrying a common-weight linear claim adds another $5 \cdot 2^{-128}$ by (5.5), giving error at most $16 \cdot 2^{-128} + \delta$, or about 124 bits.

In the standard Fiat–Shamir grinding accounting, b bits of proof of work reduce the protected round's error contribution by a factor of 2^b [BT24, Section 3.2]. Consequently, three bits of grinding bring the original combiner's MCA contribution below 2^{-128}, and four bits suffice for the equal-degree version. With the linear claim, four bits cover the leading term $16 \cdot 2^{-128}$; five bits also absorb the positive correction and give a bound below 2^{-128}.

6 Separation from scalar combiners and Necessity of Agreement Loss

In this section we show a separation between affine and scalar combiners by adapting the counterexample of [BCI+20, Remark 1.1]. In particular, in the regime $n \ll q$, Theorem 1.1 shows that affine combiners can achieve error $O(1/q)$, while the separation result below shows that scalar combiners must have error $\Omega(n/q)$, even for larger agreement losses $s > 1$. For scalar combiners, write

$$S_\alpha(P) \coloneqq \{z \in \mathcal{D} : P(z) = f(z) + \alpha g(z)\}.$$

and let $\varepsilon_{\mathrm{sc}}(n, d, A, s, q)$ denote the MCA error at agreement threshold A, defined as for affine combiners but with a uniform scalar multiplier α instead and degree bounds $\deg F, \deg G \le d$.

Proposition 6.1. *Let $s \ge 0$ be an integer and suppose $A - s - 1 > d$. There are received words f, g for which no polynomial of degree at most d agrees with g on $A - s$ coordinates, while*

$$\varepsilon_{\mathrm{sc}}(n, d, A, s, q) \ge \Pr_{\alpha \leftarrow \mathbb{F}_q}\left[|S_\alpha(0)| = A\right] \ge \frac{1 + \lfloor (n - A)/(s + 1) \rfloor}{q}. \tag{6.1}$$

Proof. Choose $C \subseteq \mathcal{D}$ of size $A - s - 1$ and define

$$V_C(X) \coloneqq \prod_{z' \in C} (X - z'), \qquad g(z) \coloneqq V_C(z).$$

For every polynomial G of degree at most d, the difference $V_C - G$ has degree $A - s - 1$. Thus G agrees with g on at most $A - s - 1$ coordinates.

Partition $\mathcal{D} \setminus C$ into blocks of size $s + 1$ and, if needed, one smaller nonempty block. Assign a distinct $\gamma_T \in \mathbb{F}_q$ to each block T; this is possible since there are at most $n \le q$ blocks. Define

$$f(z) \coloneqq 0 \quad (z \in C), \qquad f(z) \coloneqq -\gamma_T g(z) \quad (z \in T).$$

Because g is nonzero outside C, each block T of size $s + 1$ satisfies $S_{\gamma_T}(0) = C \cup T$, of size A. With agreement loss s, at least $A - s$ agreements must be preserved, and no degree-d polynomial can agree with g on that many coordinates. Hence $P = 0$ has no admissible witnesses for each of these distinct challenges. There are $1 + \lfloor (n - A)/(s + 1) \rfloor$ blocks of size $s + 1$. Since α is uniform in $\mathbb{F}_q$, we obtain (6.1). □

For the comparison with affine combiners, set $s = 1$ and write $\varepsilon_{\mathrm{sc}} \coloneqq \varepsilon_{\mathrm{sc}}(n, d, A, 1, q)$. For fixed ρ, η and sufficiently large n, we have $A - 2 > d$ and $n - A = \Theta_{\rho,\eta}(n)$. By Proposition 6.1 and Theorem 1.1,

$$\varepsilon_{\mathrm{sc}} = \Omega_{\rho,\eta}(n/q), \qquad \varepsilon_{\mathrm{aff}} = O_{\rho,\eta}(1/q) \quad \text{when } q \ge n^2.$$

Thus the affine combiner improves the worst-case MCA error by a factor of $\Omega_{\rho,\eta}(n)$ at the same agreement threshold and with the same agreement loss $s = 1$.

We also show that without allowing the one point of agreement loss, affine combiners cannot achieve the same low error as with $s = 1$.

Proposition 6.2. *Suppose $A - 1 > d$. Then*

$$\varepsilon_{\mathrm{aff}}(n, \rho, A, 0, q) \ge \frac{(n - A + 1)(q - 1)}{q^2} \ge \frac{n - A + 1}{2q}. \tag{6.2}$$

Proof. Choose $C \subseteq \mathcal{D}$ of size $A-1$. Let $f \coloneqq 0$, and let g be zero on C and one elsewhere. For each $z \in \mathcal{D} \setminus C$, there are $q-1$ nonzero pairs (α, β) with $\alpha + \beta z = 0$. Every such pair has $\beta \neq 0$ and

$$S_{\alpha,\beta}(0) = C \cup \{z\}, \qquad |S_{\alpha,\beta}(0)| = A.$$

Any polynomial G of degree at most $d-1$ agreeing with g on C must be zero, since $|C| = A-1 > d$. It therefore disagrees with g at z, so $P = 0$ has no witnesses with agreement loss $s = 0$.

The nonzero pairs counted for distinct z are disjoint. Dividing $(n-A+1)(q-1)$ by q^2 proves the first inequality in (6.2); $q - 1 \geq q/2$ proves the second. □

Constant fractional agreement loss. Even when a constant fraction of agreement loss is allowed, we show that scalar combiners cannot achieve $O(1/q)$ MCA error even on additive evaluation domains. We adapt the DEEP-FRI subspace counterexample [BGKS20, Appendix B] to degree bound $d = n/4$ by multiplying both received words and their candidate polynomials by a common vanishing polynomial.

Proposition 6.3. *Suppose $\mathbb{F}_q$ has characteristic two, $n \geq 16$ is a power of two, $q \geq n$, $d = n/4$, and $A \leq 9n/16$. For every $\mathbb{F}_2$-linear evaluation domain $\mathcal{D} \subseteq \mathbb{F}_q$ of size n and every integer $0 \leq s < 3n/16$,*

$$\varepsilon_{\mathrm{sc}}(n, d, A, s, q) \geq \frac{n-8}{q}. \tag{6.3}$$

Proof. We construct $n-8$ combinations with $9n/16$ agreements, while every degree-at-most-d polynomial agrees with g on at most $3n/8$ coordinates. The gap of $3n/16$ makes these combinations exceptional for every allowed loss s in the statement.

Choose a codimension-three subspace $\mathcal{U} \subseteq \mathcal{D}$, so $|\mathcal{U}| = n/8$, and define

$$W(X) \coloneqq \prod_{\gamma \in \mathcal{U}} (X - \gamma).$$

For $z \in \mathcal{D}$, set

$$f(z) \coloneqq W(z) z^{n/2}, \qquad g(z) \coloneqq W(z) z^{n/4}.$$

For every polynomial G of degree at most d, the polynomial $W(X)X^{n/4} - G(X)$ is nonzero and has degree $3n/8$. Thus G agrees with g on at most $3n/8$ coordinates.

For each $\mathbb{F}_2$-linear hyperplane $\mathcal{H} \subseteq \mathcal{D}$, consider its subspace polynomial

$$L_{\mathcal{H}}(X) \coloneqq \prod_{\gamma \in \mathcal{H}} (X - \gamma).$$

Such polynomials have only power-of-two exponents. This follows inductively from $L_{\{0\}}(X) = X$: for an $\mathbb{F}_2$-subspace $\mathcal{W}$ and $\gamma \notin \mathcal{W}$, we have

$$L_{\mathcal{W} + \mathbb{F}_2 \gamma}(X) = L_{\mathcal{W}}(X)^2 - L_{\mathcal{W}}(\gamma) L_{\mathcal{W}}(X).$$

Consequently,

$$L_{\mathcal{H}}(X) = X^{n/2} + \alpha_{\mathcal{H}} X^{n/4} + R_{\mathcal{H}}(X), \qquad \deg R_{\mathcal{H}} \leq n/8.$$

Distinct hyperplanes give distinct coefficients $\alpha_{\mathcal{H}}$: otherwise their subspace polynomials would differ by a polynomial of degree at most $n/8$ vanishing on their intersection of size $n/4$. They would therefore be equal, contradicting the distinct root sets.

Set $P_{\mathcal{H}} \coloneqq -WR_{\mathcal{H}}$, so $\deg P_{\mathcal{H}} \le n/4 = d$. For $z \in \mathcal{D}$,

$$f(z) + \alpha_{\mathcal{H}} g(z) - P_{\mathcal{H}}(z) = W(z) L_{\mathcal{H}}(z),$$

and hence

$$S_{\alpha_{\mathcal{H}}}(P_{\mathcal{H}}) = \mathcal{U} \cup \mathcal{H}.$$

If $\mathcal{H}$ does not contain $\mathcal{U}$, then $|\mathcal{H} \cap \mathcal{U}| = n/16$, giving

$$|S_{\alpha_{\mathcal{H}}}(P_{\mathcal{H}})| = \frac{n}{8} + \frac{n}{2} - \frac{n}{16} = \frac{9n}{16} \ge A.$$

Of the $n-1$ hyperplanes of $\mathcal{D}$, exactly $2^3 - 1 = 7$ contain $\mathcal{U}$, leaving $n - 8$ distinct challenges with this agreement. For each of them, witnesses with agreement loss at most s would have to agree with g on at least $9n/16 - s > 3n/8$ coordinates, which is impossible. Dividing by q proves (6.3). □

Taking $s = \lfloor \lambda n \rfloor$ for any fixed $0 < \lambda < 3/16$ gives scalar MCA error $\Omega(n/q)$, including for 1% and 10% agreement loss. At $\rho = 1/4$, the condition on A includes every fixed Johnson slack $0 < \eta \le 1/16$. For sufficiently large n and $q \ge n^2$, Theorem 1.1 gives affine MCA error $O_{\rho,\eta}(1/q)$ already with one-coordinate loss, and hence also with this larger allowed loss. Thus the separation persists when a constant fraction of agreement may be discarded.

7 Extension to Reed–Muller codes

Fix an integer $k \geq 1$ and a subfield $\mathbb{F}_{q_0} \subseteq \mathbb{F}_q$, and set $n \coloneqq q_0^k$, the size of the grid $\mathbb{F}_{q_0}^k \subseteq \mathbb{F}_q^k$. In this section, degrees of polynomials in $\mathbb{F}_q[\mathbf{X}] = \mathbb{F}_q[X_1, \ldots, X_k]$ are total degrees. Received words are functions on the grid, so their coordinates are the grid points. For $\rho \in (0,1)$ such that $d \coloneqq \rho q_0$ is an integer, the Reed–Muller code consists of the evaluations on $\mathbb{F}_{q_0}^k$ of polynomials of degree at most d, and it has relative distance $1-\rho$. We call ρ the degree fraction; unlike for Reed–Solomon codes, it is not the rate of the code when $k \geq 2$.

Grid points on a hyperplane. Let $H = \{\mathbf{x} \in \mathbb{F}_q^k : \mathbf{a} \cdot \mathbf{x} = b\}$ be a hyperplane of $\mathbb{F}_q^k$, where $\mathbf{a} \neq 0$. The grid points in H, if any, form a coset of the $\mathbb{F}_{q_0}$-linear subspace $\{\mathbf{z} \in \mathbb{F}_{q_0}^k : \mathbf{a} \cdot \mathbf{z} = 0\}$, which is proper since $\mathbf{a} \neq 0$. Therefore

$$|H \cap \mathbb{F}_{q_0}^k| \leq n/q_0, \tag{7.1}$$

with equality for the hyperplane $x_1 = 0$. Consequently, a set of more than n/q_0 grid points is not contained in any hyperplane, so it contains $k+1$ affinely independent points.

The affine combiner. For received words $f, g : \mathbb{F}_{q_0}^k \to \mathbb{F}_q$, we sample a uniform $\boldsymbol{\alpha} = (\alpha_0, \ldots, \alpha_k) \in \mathbb{F}_q^{k+1}$ and form

$$w_{\boldsymbol{\alpha}}(\mathbf{x}) \coloneqq f(\mathbf{x}) + L_{\boldsymbol{\alpha}}(\mathbf{x}) g(\mathbf{x}), \qquad L_{\boldsymbol{\alpha}}(\mathbf{X}) \coloneqq \alpha_0 + \alpha_1 X_1 + \cdots + \alpha_k X_k.$$

As in Section 1.1, the factor $L_{\boldsymbol{\alpha}}$ is the affine combiner, and for polynomials with $\deg F \leq d$ and $\deg G \leq d-1$, the polynomial $F + L_{\boldsymbol{\alpha}} G$ has degree at most d.

Definition of the MCA error. Let $1 \leq A \leq n$ and $s \geq 0$ be integers, and let $S_{\boldsymbol{\alpha}}(P) \coloneqq \{\mathbf{x} \in \mathbb{F}_{q_0}^k : P(\mathbf{x}) = w_{\boldsymbol{\alpha}}(\mathbf{x})\}$. For $\boldsymbol{\alpha}$, polynomials $F, G \in \mathbb{F}_q[\mathbf{X}]$ of degrees at most d and $d-1$ respectively are admissible witnesses for $P = F + L_{\boldsymbol{\alpha}} G$ with agreement loss at most s if

$$\big|\{\mathbf{x} \in S_{\boldsymbol{\alpha}}(P) : F(\mathbf{x}) \neq f(\mathbf{x}) \text{ or } G(\mathbf{x}) \neq g(\mathbf{x})\}\big| \leq s.$$

As in Section 1.1, the MCA error with agreement loss at most s is the smallest ε such that, for every pair f, g, the probability over $\boldsymbol{\alpha}$ that some polynomial P with $\deg P \leq d$ and $|S_{\boldsymbol{\alpha}}(P)| \geq A$ has no admissible witnesses is at most ε. We denote it by $\varepsilon_{\mathrm{RM}}(k, q_0, \rho, A, s, q)$.

Theorem 7.1. *Let $k \geq 1$ be an integer and $\mathbb{F}_{q_0}$ a finite field, and set $n \coloneqq q_0^k$. Fix $\rho \in (0,1)$ with ρq_0 an integer, let $0 < \eta < 1 - \sqrt{\rho}$, and set $A \coloneqq \lceil (\sqrt{\rho} + \eta) n \rceil$. Suppose $q_0 \eta^2 \geq 4/\rho$. Then, for every finite field $\mathbb{F}_q$ containing $\mathbb{F}_{q_0}$, the MCA error with agreement loss $s = n/q_0$ satisfies*

$$\varepsilon_{\mathrm{RM}}\Big(\underbrace{k}_{\substack{\text{number of}\\ \text{variables}}}, \underbrace{q_0}_{\substack{\text{subfield}\\ \text{size}}}, \underbrace{\rho}_{\substack{\text{degree}\\ \text{fraction}}}, \underbrace{A}_{\substack{\text{agreement}\\ \text{threshold}}}, \underbrace{n/q_0}_{\substack{\text{agreement}\\ \text{loss}}}, \underbrace{q}_{\substack{\text{field}\\ \text{size}}} \Big) \leq \frac{4}{\eta q^k} + \frac{2^{6k+8} n^{k+1}}{\rho^{k+2} \eta^{4k+6} q^{k+1}}. \tag{7.2}$$

For $k = 1$, Theorem 7.1 is Theorem 1.1 for the domain $\mathcal{D} = \mathbb{F}_{q_0}$. For fixed ρ, η, k and $q \geq n^{k+1}$, the error is $O_{\rho,\eta,k}(q^{-k})$.

The exponent k cannot be improved. The q challenges with $\alpha_1 = \cdots = \alpha_k = 0$, a q^{-k} fraction of all challenges, have the constant multiplier α_0, so a received word g of degree d leads to combined words of degree d; these have no admissible witnesses, which require $\deg G \leq d-1$. Indeed, take $f = 0$ and $g = X_1^d$. For each such challenge, the polynomial $P = \alpha_0 X_1^d$ agrees with $w_{\boldsymbol{\alpha}}$ on all n

coordinates. By the Schwartz–Zippel lemma, every polynomial of degree at most $d-1$ disagrees with X_1^d on at least $(1-\rho)n = (q_0-d)n/q_0$ coordinates. This is more than n/q_0, since the conditions of Theorem 7.1 imply $d \le q_0 - 2$: if $d = q_0 - 1$, then $\eta < 1 - \sqrt{\rho} < 1 - \rho = 1/q_0$, so $q_0\eta^2 < 1/q_0 < 4/\rho$. Thus P has no admissible witnesses, and $\varepsilon_{\mathrm{RM}}(k, q_0, \rho, A, n/q_0, q) \ge q/q^{k+1} = q^{-k}$.

7.1 Proof idea

We prove Theorem 7.1 in Appendix A by following the proof of Theorem 1.1. The proof changes in three places.

Agreement loss and the order of the error. Fix polynomials F, G with $\deg F \le d$ and $\deg G \le d-1$, and call an agreement of $P = F + L_{\boldsymbol{\alpha}} G$ with $w_{\boldsymbol{\alpha}}$ *accidental* if it occurs at a coordinate $\mathbf{z}$ where F, G do not both agree with f, g. For each such coordinate, the value $L_{\boldsymbol{\alpha}}(\mathbf{z})$ is fixed as in the proof of Lemma 4.4. In the univariate case, the challenge is determined by two such values at distinct coordinates. Here, values at $k+1$ affinely independent points are needed: values on a hyperplane $\{L_{\boldsymbol{\sigma}} = 0\}$ are the same for at least q challenges, since adding a multiple of $\boldsymbol{\sigma}$ to $\boldsymbol{\alpha}$ does not change them. A hyperplane contains at most n/q_0 grid points by (7.1), so with agreement loss n/q_0, an exceptional candidate has accidental agreements at $k+1$ affinely independent points. Hence at most $\binom{n}{k+1}$ of the q^{k+1} challenges are exceptional for the pair.

If F, G are rational functions that do not satisfy these degree bounds, we use, as in the proof of Lemma 4.4, that the challenges for which $F + L_{\boldsymbol{\alpha}} G$ is a polynomial of degree at most d form a proper affine subspace. For $k \ge 2$, a proper subspace can contain q^k challenges, and the resulting error bound would only be of order $1/q$. By Lemma A.1, this subspace has dimension at most one: for any two of its elements $\boldsymbol{\alpha}, \boldsymbol{\alpha}'$, the product $L_{\boldsymbol{\alpha}-\boldsymbol{\alpha}'} G$ is a polynomial of degree at most d, and $\boldsymbol{\alpha} - \boldsymbol{\alpha}'$ is therefore a multiple of one fixed vector. Summing over the $O(1/\eta)$ persistent pairs, with at most q challenges per pair, we obtain a leading term of order $1/(\eta q^k)$.

Reduction to univariate polynomials. The interpolation in Section 4.1 is univariate, so we first map the Reed–Muller code into a Reed–Solomon code of length n with the same degree fraction [PW04]. This map identifies the grid with the field $\mathbb{F}_n$ through an $\mathbb{F}_{q_0}$-linear bijection whose components, polynomials of degree n/q_0, are substituted for $X_1, \dots, X_k$. Polynomials of degree at most d become univariate polynomials of degree at most $dn/q_0 = \rho n$, so the agreement threshold is the same as in Theorem 1.1. The affine combiner, however, becomes a polynomial of degree n/q_0, so we interpolate with the $(1, n/q_0, \rho n)$-weighted degree. Because of this larger degree, Theorem 7.1 requires $q_0\eta^2 \ge 4/\rho$, while Theorem 1.1 for block length n would require only $n\eta^2 \ge 4/\rho$.

Nonpersistent candidates. The challenge now consists of $k+1$ field elements, so the nonpersistent candidates lie in algebraic sets of dimension at most $k+1$, and the agreement-incidence argument bounds their number by $O_{\rho,\eta,k}(n^{k+1})$, so their contribution to the error is of order n^{k+1}/q^{k+1}. The reconstructed univariate polynomials need not have the form $\widetilde{P}$ for a Reed–Muller polynomial P, so we intersect these algebraic sets with the linear subspace where they do, which does not increase their dimension or degree. In the agreement-incidence argument, the denominator $L_{\boldsymbol{\sigma}}$ of a persistent pair can vanish at up to n/q_0 coordinates by (7.1), which is why the persistence threshold is lowered by n/q_0. Simple-root lifting and the agreement-incidence argument are otherwise unchanged, since they only use that the multiplier is linear in $\boldsymbol{\alpha}$.

Acknowledgments

We thank Ron Rothblum for his comments on an earlier version of this paper, particularly on its connection to Flock. We also thank Ariel Gabizon for his comments.

Eylon Yogev is supported by the Israel Science Foundation (Grant No. 2302/22) and the European Union (ERC, CRYPTOPROOF, 101164375). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

AI usage. The authors were initially motivated to explore this direction by the connection between proximity gaps and robust versions of the line-versus-point test, as described in Section 2. In an earlier version of this paper, the authors were able to obtain proximity gaps for high-variate Reed-Muller codes over small fields up to their list-decoding radius using this agreement testing connection and the list-decoding results of [GKZ08, BL18]. However, this result used combinatorial techniques and we were unable to obtain interesting results for Reed-Solomon codes using them.

The authors then asked ChatGPT 6 Astra to prove an improved proximity-gap result for these combiners and immediately obtained Theorem 1.1. The analysis, including all proofs appearing in this paper, was carried out entirely with AI. The authors edited, verified, reorganized, and Lean-formalized the initial writeup. The Lean formalization was developed using ChatGPT 6 Astra, and the theorem statement was independently checked by Claude Code Opus 5.5. The equal-degree reductions in Section 5 and the Reed–Muller extension in Section 7 and Appendix A are not covered by the Lean formalization. The authors take full responsibility for the contents of this paper.

References

[ABF26] Gal Arnon, Dan Boneh, and Giacomo Fenzi. Open problems in list decoding and correlated agreement. Cryptology ePrint Archive, Paper 2026/680, 2026. URL: https://eprint.iacr.org/2026/680.

[ACFY24a] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. STIR: Reed–Solomon proximity testing with fewer queries. Cryptology ePrint Archive, Paper 2024/390, 2024. URL: https://eprint.iacr.org/2024/390.

[ACFY24b] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. WHIR: Reed–Solomon proximity testing with super-fast verification. Cryptology ePrint Archive, Paper 2024/1586, 2024. URL: https://eprint.iacr.org/2024/1586.

[AGG26] Omar Alrabiah, Rohan Goyal, and Venkatesan Guruswami. Combinatorial limitations of list-decoding Reed–Solomon codes. arXiv:2611.07607, November 2026. URL: https://arxiv.org/abs/2611.07607.

[AS03] Sanjeev Arora and Madhu Sudan. Improved low-degree testing and its applications. *Combinatorica*, 23(3):365–426, 2003. doi:10.1007/s00493-003-0025-0.

[BBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast Reed–Solomon interactive oracle proofs of proximity. In *ICALP 2018*, volume 107 of *LIPIcs*, pages 14:1–14:17, 2018. URL: https://doi.org/10.4230/LIPIcs.ICALP.2018.14.

[BCH+25] Eli Ben-Sasson, Dan Carmon, Ulrich Haböck, Swastik Kopparty, and Shubhangi Saraf. On proximity gaps for Reed–Solomon codes. Technical Report TR25-169, Electronic Colloquium on Computational Complexity, November 2025. Especially Theorem 4.6. URL: https://eccc.weizmann.ac.il/report/2025/169/.

[BCI+20] Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty, and Shubhangi Saraf. Proximity gaps for Reed–Solomon codes. Technical Report TR20-083, Electronic Colloquium on Computational Complexity, 2020. Revision 3, July 2021. URL: https://eccc.weizmann.ac.il/report/2020/083/revision/3/.

[BCP+26] Joshua Brakensiek, Yeyuan Chen, Aaron Putterman, Zihan Zhang, and Kai Zhe Zheng. Algorithmic list decoding of Reed–Solomon codes up to capacity. arXiv:2609.08005, September 2026. Version 1. URL: https://arxiv.org/abs/2609.08005.

[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)*, volume 151 of *LIPIcs*, pages 5:1–5:32, 2020. doi:10.4230/LIPIcs.ITCS.2020.5.

[BL18] Abhishek Bhowmick and Shachar Lovett. The list decoding radius for Reed–Muller codes over small fields. *IEEE Transactions on Information Theory*, 64(6):4382–4391, 2018. URL: https://doi.org/10.1109/TIT.2018.2822686.

[BRW26] Benedikt Bünz, Ron Rothblum, and William Wang. Flock: Fast proving for batch boolean computations. Cryptology ePrint Archive, Paper 2026/1329, 2026. URL: https://eprint.iacr.org/2026/1329.

[BT24] Alexander R. Block and Pratyush Ranjan Tiwari. On the concrete security of non-interactive FRI. Cryptology ePrint Archive, Paper 2024/1161, 2024. URL: https://eprint.iacr.org/2024/1161.

[DKT26] Quang Dao, Scott Duke Kominers, and Justin Thaler. Reed–Solomon codes beyond Johnson: Efficient decoding and smaller cryptographic proofs. Cryptology ePrint Archive, Paper 2026/2056, September 2026. Revised November 3, 2026. URL: https://eprint.iacr.org/2026/2056.

[DKTZ26] Quang Dao, Scott Duke Kominers, Justin Thaler, and Kai Zhe Zheng. Counterexamples to beyond-Johnson proximity gaps over binary fields. Technical Report TR26-237, Electronic Colloquium on Computational Complexity, October 2026. URL: https://eccc.weizmann.ac.il/report/2026/237/.

[Fs26] Matteo Frigo and abhi shelat. Polygero: Proximity testing from polynomial challenges. Cryptology ePrint Archive, Paper 2026/2481, October 2026. URL: https://eprint.iacr.org/2026/2481.

[Gab26] Ariel Gabizon. Linear list size bounds for Reed–Solomon beyond the Johnson radius. Cryptology ePrint Archive, Paper 2026/2048, 2026. URL: https://eprint.iacr.org/2026/2048.

[GKZ08] Parikshit Gopalan, Adam R. Klivans, and David Zuckerman. List-decoding Reed–Muller codes over small fields. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC 2008)*, pages 265–274, 2008. URL: https://www.cs.utexas.edu/~klivans/rm.pdf.

[GS99] Venkatesan Guruswami and Madhu Sudan. Improved decoding of Reed–Solomon and algebraic-geometry codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, 1999.

[Hei83] Joos Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoretical Computer Science*, 24(3):239–277, 1983. doi:10.1016/0304-3975(83)90002-6.

[HKS26] Prahladh Harsha, Mrinal Kumar, and Ramprasad Saptharishi. Reed–Solomon codes at capacity: Algorithmic list-decoding and proximity gaps. arXiv:2610.08610, October 2026. URL: https://arxiv.org/abs/2610.08610.

[HKSS24] Prahladh Harsha, Mrinal Kumar, Ramprasad Saptharishi, and Madhu Sudan. An improved line-point low-degree test. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1883–1892. IEEE, 2024. doi:10.1109/FOCS61266.2024.00113.

[Jer26] Fernando Granha Jeronimo. Algorithmic list decoding at capacity and optimal proximity gaps for Reed–Solomon codes. Technical Report TR26-169, Electronic Colloquium on Computational Complexity, September 2026. URL: https://eccc.weizmann.ac.il/report/2026/169/.

[KTZ26] Scott Duke Kominers, Justin Thaler, and Kai Zhe Zheng. Improved soundness for the line–versus–point test. Technical Report TR26-147, Electronic Colloquium on Computational Complexity, August 2026. Revision 1. URL: https://eccc.weizmann.ac.il/report/2026/147/revision/1/download.

[lea26] leanEthereum. leanVM: Minimal zkvm for ethereum. Software repository, 2026. Accessed October 31, 2026. URL: https://github.com/leanEthereum/leanVM.

[MS77] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error-Correcting Codes*, volume 16 of *North-Holland Mathematical Library*. North-Holland, 1977. URL: https://neilsloane.com/doc/ms77.html.

[MZ25] Dor Minzer and Kai Zhe Zheng. Improved round-by-round soundness IOPs via Reed–Muller codes. arXiv:2504.00346, 2025. Full version of FOCS 2025. URL: https://arxiv.org/abs/2504.00346.

[NA25] Andrija Novakovic and Guillermo Angeris. Ligerito: A small and concretely fast polynomial commitment scheme. Cryptology ePrint Archive, Paper 2025/1187, 2025. URL: https://eprint.iacr.org/2025/1187.

[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, pages 194–203. ACM, 1994. doi:10.1145/195058.195132.

[PW04] Ruud Pellikaan and Xin-Wen Wu. List decoding of q-ary Reed–Muller codes. *IEEE Transactions on Information Theory*, 50(4):679–682, 2004. URL: https://ruudp.win.tue.nl/paper/43.pdf.

[Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, 1980.

[Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Symbolic and Algebraic Computation (EUROSAM '79)*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226. Springer, 1979.

A Proof of the Reed–Muller extension

The proof of Theorem 7.1 follows Section 4, and only the steps that change are described here. A pair $(\boldsymbol{\alpha}, P)$ is *exceptional* if $\deg P \le d$, $|S_{\boldsymbol{\alpha}}(P)| \ge A$, and P has no admissible witnesses with agreement loss at most n/q_0. As in Section 4, we divide the coefficient vectors $\boldsymbol{\alpha} \in \mathbb{F}_q^{k+1}$ for which some $(\boldsymbol{\alpha}, P)$ is exceptional into the sets $\mathcal{E}$, $\mathcal{P}$, and $\mathcal{N}$ of excluded, persistent, and nonpersistent ones, and bound each set separately.

Reduction to univariate polynomials. Following [PW04], we map the Reed–Muller code into a Reed–Solomon code with the same degree fraction by identifying the grid with the field $\mathbb{F}_n$ through an $\mathbb{F}_{q_0}$-linear bijection $\psi = (\psi_1, \dots, \psi_k)$. Its components $\psi_i(X) = \mathrm{Tr}_{\mathbb{F}_n/\mathbb{F}_{q_0}}(\gamma_i X)$, for a basis $\gamma_1, \dots, \gamma_k$ of $\mathbb{F}_n$ over $\mathbb{F}_{q_0}$, are polynomials of degree n/q_0. Substituting $\psi_1, \dots, \psi_k$ for $X_1, \dots, X_k$ maps each polynomial P of degree at most d to a univariate polynomial $\widetilde{P}$ of degree at most ρn, and distinct P give distinct $\widetilde{P}$: if $\widetilde{P}_1 = \widetilde{P}_2$, then $P_1 - P_2$ vanishes at all n grid points, and since its total degree is less than q_0, $P_1 = P_2$ by the Schwartz–Zippel lemma [Sch80, Zip79]. With the received words $f \circ \psi$ and $g \circ \psi$ on $\mathbb{F}_n$, each candidate P corresponds to the Reed–Solomon candidate $\widetilde{P}$, with the same number of agreements, for the multiplier $\widetilde{L}_{\boldsymbol{\alpha}} = \alpha_0 + \sum_i \alpha_i \psi_i$, which has degree n/q_0. The interpolation points therefore lie in $\mathbb{F}_n$, while the received words take values in $\mathbb{F}_q$, so we work over the smallest field $\mathbb{F}'$ containing $\mathbb{F}_q$ and $\mathbb{F}_n$. Up to the agreement-incidence argument, Section 4 uses no property of its base field other than being finite, so the factorization of Q, the normalization by e_i-th powers, the bound on $\mathcal{E}$, and simple-root lifting hold over $\mathbb{F}'$, with $K \coloneqq \overline{\mathbb{F}'(T_0)}$. Claim 4.8 is the exception, treated below.

Interpolation. The interpolation follows Lemma 4.2 over $\mathbb{F}'$, with the $(1, n/q_0, \rho n)$-weighted degree. With m and D from Section 4.1, the Z-degree is bounded by

$$M \coloneqq \min\left\{D, \left\lfloor \frac{mA - (n/q_0)D - 1}{\rho n} \right\rfloor\right\}.$$

Since $\deg \widetilde{L}_{\boldsymbol{\alpha}} \le n/q_0$, the condition $D \le mA - U$ in the proof of Lemma 4.2, where $U \coloneqq \lceil (m + 1/2)\sqrt{\rho}\, n \rceil$, becomes $(n/q_0)D \le mA - U$. In that proof, $mA - U \ge m\eta n/2 - 1$; hence this condition holds when $q_0\eta^2 \ge 4/\rho$, because $D + 1 \le 2m^2/(3\rho)$ and $m < 3/\eta$. The result is a nonzero polynomial Q such that $Q(X, \widetilde{L}_{\boldsymbol{\alpha}}(X), \widetilde{P}(X)) = 0$ whenever $\deg P \le d$ and $|S_{\boldsymbol{\alpha}}(P)| \ge A$.

Excluded coefficient vectors. We factor Q as in Section 4 and define $\mathcal{E}$ by (4.8), with $\widetilde{L}_{\boldsymbol{\alpha}}$ instead of $\alpha + \beta X$ and with the e_i-th root of $\widetilde{L}_{\boldsymbol{\alpha}}(X^{e_i})$ instead of $\alpha^{1/e_i} + \beta^{1/e_i} X$. Then $|\mathcal{E}| \le D^2$ by the proof of Lemma 4.3, which only uses that distinct coefficient vectors give distinct polynomials $\widetilde{L}_{\boldsymbol{\alpha}}$; this follows from the argument above with $P = L_{\boldsymbol{\alpha}}$.

Persistent coefficient vectors. As in Section 4.3, we first bound the number of exceptional coefficient vectors for one fixed rational pair.

Lemma A.1. *Let $F, G \in \mathbb{F}_q(\mathbf{X})$. The number of coefficient vectors $\boldsymbol{\alpha}$ for which $P \coloneqq F + L_{\boldsymbol{\alpha}} G$ is a polynomial and $(\boldsymbol{\alpha}, P)$ is exceptional is at most $\binom{n}{k+1}$ if F, G are polynomials with $\deg F \le d$ and $\deg G \le d - 1$, and at most q otherwise.*

Proof. First suppose that F, G are polynomials with these degree bounds, and let I' be their full correlated agreement set with f, g. At an accidental agreement $\mathbf{z} \notin I'$, we have $g(\mathbf{z}) \ne G(\mathbf{z})$ as

in (4.11), so the agreement imposes the linear condition

$$L_{\boldsymbol{\alpha}}(\mathbf{z}) = \frac{F(\mathbf{z}) - f(\mathbf{z})}{g(\mathbf{z}) - G(\mathbf{z})}$$

on $\boldsymbol{\alpha}$. Since F and G are fixed, the right-hand side does not depend on $\boldsymbol{\alpha}$. If $(\boldsymbol{\alpha}, P)$ is exceptional, then F, G are not admissible witnesses, so P has more than n/q_0 accidental agreements. By (7.1), these include $k+1$ affinely independent points, and the conditions at these points determine $\boldsymbol{\alpha}$. Choosing such $k+1$ points for each exceptional coefficient vector therefore assigns distinct sets to distinct vectors, and there are at most $\binom{n}{k+1}$ sets.

Otherwise, the coefficient vectors for which $F + L_{\boldsymbol{\alpha}} G$ is a polynomial of degree at most d form an affine subspace of $\mathbb{F}_q^{k+1}$, if nonempty, and we show that it has dimension at most one. In Lemma 4.4, it sufficed that the subspace is proper, but a proper subspace of $\mathbb{F}_q^{k+1}$ can contain q^k coefficient vectors. Note that G is not a polynomial of degree at most $d-1$, since then $F = P - L_{\boldsymbol{\alpha}} G$ would be a polynomial of degree at most d, which is the first case. Now let $\boldsymbol{\alpha} \neq \boldsymbol{\alpha}'$ lie in the subspace, so that $L_{\boldsymbol{\alpha}-\boldsymbol{\alpha}'} G$ is a polynomial of degree at most d. If G is not a polynomial, its denominator in lowest terms divides the nonzero polynomial $L_{\boldsymbol{\alpha}-\boldsymbol{\alpha}'}$ of degree at most one, so the denominator is $L_{\boldsymbol{\sigma}}$ for some nonzero $\boldsymbol{\sigma}$ determined by G, and $\boldsymbol{\alpha} - \boldsymbol{\alpha}'$ is a multiple of $\boldsymbol{\sigma}$. If G is a polynomial of degree at least d, then $L_{\boldsymbol{\alpha}-\boldsymbol{\alpha}'}$ must be constant, so $\boldsymbol{\alpha} - \boldsymbol{\alpha}'$ is a multiple of $(1, 0, \ldots, 0)$. In both cases, all differences between elements of the subspace are multiples of one vector, so the subspace has dimension at most one. □

Persistence is defined as in Section 4, with rational pairs F, G such that $G = H/L_{\boldsymbol{\sigma}}$ for a polynomial H of degree at most d and a nonzero $\boldsymbol{\sigma} \in \mathbb{F}_q^{k+1}$. The pair defined in the proof of Claim 4.8 from two solutions $(\boldsymbol{\alpha}, P)$ and $(\boldsymbol{\alpha}', P')$ is of this form, with $H = P' - P$ and $\boldsymbol{\sigma} = \boldsymbol{\alpha}' - \boldsymbol{\alpha}$. There, the pair obtained from c agreement hyperplanes agrees with f, g on the c corresponding coordinates, except at the zeros of $L_{\boldsymbol{\sigma}}$. If $L_{\boldsymbol{\sigma}}$ is not constant, its zeros form a hyperplane of $\mathbb{F}_q^k$, so by (7.1), $L_{\boldsymbol{\sigma}}$ has at most n/q_0 zeros on the grid, while the denominator L in Section 4 has at most one zero. We therefore set $B \coloneqq \lceil (\sqrt{\rho} + \eta/2)n \rceil - n/q_0$. For two distinct pairs (F_j, G_j) of this form, with $G_j = H_j / L_{\boldsymbol{\sigma}_j}$, explaining exceptional $(\boldsymbol{\alpha}_j, P_j)$ for $j = 0, 1$, the intersection of their correlated agreement sets lies in the zero set of a nonzero polynomial of degree at most $d+1$, as in the proof of Lemma 4.5. By the Schwartz–Zippel lemma, this intersection therefore contains at most $(d+1)n/q_0$ coordinates. Since $B > (d+1)n/q_0$, as shown below, Lemma 3.2 with $a = B/n$ and $b = (d+1)/q_0$ bounds the number of such pairs by $\ell \coloneqq (a-b)/(a^2-b)$, and by Lemma A.1, $|\mathcal{P}| \leq \ell \max\{q, \binom{n}{k+1}\}$.

Nonpersistent coefficient vectors. Simple-root lifting applies to $\widetilde{L}_{\boldsymbol{\alpha}}$ as to the line $\alpha + \beta X$. For a solution $(\boldsymbol{\alpha}, P)$, set $Y \coloneqq \widetilde{L}_{\boldsymbol{\alpha}}(T_0^{e_i})^{1/e_i}$, $V_j \coloneqq \alpha_j^{1/e_i}$, and $Z \coloneqq \widetilde{P}(T_0^{e_i})$. At $X = T_0 + W$, the e_i-th root of $\widetilde{L}_{\boldsymbol{\alpha}}(X^{e_i})$ equals $Y + \sum_j V_j \theta_j(W)$ for fixed polynomials $\theta_1, \ldots, \theta_k \in K[W]$ with $\theta_j(0) = 0$. The proof of Lemma 3.11 also applies to this polynomial. In the proof, the coefficients of $Y + V(T_0 + W)$ in W are linear in (Y, V), and here the coefficients of $Y + \sum_j V_j \theta_j(W)$ are linear in $(Y, V_1, \ldots, V_k)$. The parameters $(Y, V_1, \ldots, V_k, Z)$ lie on the hypersurface $R_i(T_0, Y, Z) = 0$ in K^{k+2}, of dimension $k+1$. The reconstruction in Claim 4.9, with ρn instead of d, therefore produces algebraic sets of dimension at most $k+1$ and degree at most $h_i(2e_i \rho n h_i)^{k+1}$.

Claim 4.8 cannot be applied to these sets directly. Its proof needs agreement hyperplanes defined over $\mathbb{F}_q$, so that the rational pair it constructs is defined over $\mathbb{F}_q$, but in the coefficient space of $(\boldsymbol{\alpha}, \widetilde{P})$, the agreement condition at a point $x \in \mathbb{F}_n$ involves the powers $x^j \in \mathbb{F}_n$. The sets are therefore moved to the coefficient space $K^{k+1} \times K^{\binom{d+k}{k}}$ of $(\boldsymbol{\alpha}, P)$, where the agreement hyperplanes $P(\mathbf{z}) = f(\mathbf{z}) + L_{\boldsymbol{\alpha}}(\mathbf{z}) g(\mathbf{z})$ are defined over $\mathbb{F}_q$. To do so, we intersect them with the

image of the injective linear map $(\boldsymbol{\alpha}, P) \mapsto (\boldsymbol{\alpha}, \widetilde{P})$, which also discards reconstructed polynomials not of the form $\widetilde{P}$, and pull the intersection back along this map. By Lemma 3.9, applied to each hyperplane defining the image, the intersection does not increase the dimension or the degree, and the pullback preserves both. The claim then applies with two changes to its proof: $\boldsymbol{\alpha} \neq \boldsymbol{\alpha}'$ by the Schwartz–Zippel lemma, and the two cases are $c > B + n/q_0 - 1$ and $c \leq B + n/q_0 - 1$, so that in the first case the pair agrees with f, g on at least $c - n/q_0 \geq B$ coordinates outside the zeros of $L_{\boldsymbol{\sigma}}$. Summing over the factors as in the proof of Lemma 4.7, with $v \coloneqq n/(A - B - n/q_0 + 1)$, we obtain $|\mathcal{N}| \leq (2\rho n M v)^{k+1} D^{k+2}$.

The bound. As in the proof of Lemma 4.1, with ρn instead of d, the parameters satisfy $D \leq 9/(2\rho\eta^2)$, $\rho n M \leq 3n/\eta$, $v \leq 2/\eta$, and $\ell \leq 4/\eta$. For v, note that $A - B - n/q_0 + 1 = A - \lceil(\sqrt{\rho} + \eta/2)n\rceil + 1 > \eta n/2$. For ℓ, the condition $q_0\eta^2 \geq 4/\rho$, which also forces $q_0 > 64$, implies $a - \sqrt{b} \geq \eta/4$ by the same computation; in particular, $B > (d+1)n/q_0$. Every coefficient vector we count lies in $\mathcal{E} \cup \mathcal{P} \cup \mathcal{N}$, so their number is at most

$$D^2 + \ell \max\{q, \tbinom{n}{k+1}\} + (2\rho n M v)^{k+1} D^{k+2} \leq \frac{4q}{\eta} + \frac{2^{6k+8} n^{k+1}}{\rho^{k+2}\eta^{4k+6}}.$$

Dividing by q^{k+1} proves Theorem 7.1.