Khatam: Reducing the Communication Complexity of Code-Based SNARKs

Hadas Zeilberger

Yale University

February 20, 2025

Abstract

Every linear code satisfies the property of "correlated agreement", meaning that if π_L, π_R are two vectors in $\mathbb{F}^n$ and if $\pi_L + r\pi_R$ is close in Hamming distance to some codeword in C, then π_L and π_R each agree with a codeword in C in positions indexed by elements of $S \subset [n]$. In this work, we prove something stronger – that if $\pi_L + r\pi_R$ is close to C, then π_L, π_R *and* $(\pi_L + r\pi_R)$ all agree with codewords at positions indexed by elements of S, except with negligible probability over $r \leftarrow \mathbb{F}$. Our result holds as long as $|S| > (1 - \Delta_C + \epsilon)^{1/3}$, and with failure probability smaller than $\frac{2}{\epsilon^2|\mathbb{F}|}$. Furthermore, our results extend to any finite field and any linear code.

We use this result to prove that BaseFold(Crypto 2024), an efficient multilinear polynomial commitment scheme, is secure in the *list decoding regime*, which significantly reduces its communication complexity. Our result is agnostic with respect to both the field and code, and therefore can be used to reduce the communication complexity of a wide class of coding-based multilinear polynomial commitment schemes.

1 Introduction

In recent years, error-correcting codes have emerged as a key ingredient in the construction of efficient SNARKs. A prover of a code-based SNARK commits to its witness by encoding it with a linear error-correcting code, which uses relatively cheap operations such as finite-field addition and multiplication. The verifier tests the proximity of the prover's codeword to the error-correcting code by engaging with the prover in an interactive oracle proof of proximity[1](IOPP) [5, 22]. Then, using a collision-resistant hash function, we obtain a polynomial commitment scheme (PCS) [18], where a prover commits to a polynomial $P \in \mathbb{F}[X]$ so that it can later prove evaluation claims of the form $P(\alpha) = \beta$. Finally, a PCS compiles a polynomial interactive oracle proof (PIOP) into a SNARK. We refer to [7, 11, 24, 25] for more details on this transformation.

Despite their impressive *prover* efficiency, verifier costs remain a major bottleneck in code-based SNARKs, due mainly to the query complexity of the underlying IOPP. IOPPs have multiple rounds; in each round the prover sends an oracle in response to verifier randomness and the verifier queries and tests elements from these oracles. It would be too expensive for the verifier to query each element from the oracle. Instead the verifier obtains a *probabilistic*

[1]It may be useful to think of an IOPP as a PCPP [4, 13] but with multiple interactive rounds;

guarantee that a *large fraction*, $\beta \in [0,1]$, of elements from each prover oracle pass its tests. We choose the number of verifier queries to be l such that $\beta^l < 2^{-\lambda}$, where λ is a security parameter of our choosing. In the Fast Reed-Solomon IOPP(FRI) [6], $\beta > \sqrt{1-\Delta_C}$, where Δ_C is the minimum distance of the code. This setting is the best proven result and is commonly referred to as the "list-decoding" regime, since by the famous Johnson Bound, there is only a small *list* of codewords that agree with any vector in more than $\sqrt{1-\Delta_C}$ fraction of locations. Alternatively, we refer to the (inferior) case when β is greater than $(1-\Delta_C/2)$ as the "unique-decoding regime", as there is only one unique codeword that agrees with a vector in that many locations (due to the distance properties of the error-correcting code).

Since the FRI IOPP is proven secure in the *list-decoding regime*, its verifier is both asymptotically and concretely efficient. However, the FRI IOPP can only be directly used[2] as a univariate PCS, and SNARKs based on univariates have a higher overhead than those based on multilinear PCS, which is a generalization of PCS to multilinear polynomials, first introduced by [21][3]. BaseFold [25], introduced a technique for using FRI directly as a multilinear PCS by interleaving the sum-check protocol [20] for multilinear polynomial evaluation with (a generalization of) FRI. It avoids the overhead of univariate SNARKs while maintaining polylogarithmic communication. Several works ([2, 12, 17, 23]) have already adapted BaseFold to different settings. However, BaseFold's soundness proof requires a stronger result than the correlated agreement result that underpins FRI's security. As such, it was only proven secure in the *unique decoding regime* and as a result, it's proofs are concretely larger than FRI's. For that matter, no[4] multilinear polynomial commitment scheme has been proven secure in the list-decoding regime. In particular, Brakedown [15], Ligero [1], two other state-of-the-art multilinear PCS, are only proven secure in the *unique decoding regime*, and the same is true for two recent multilinear Polynomial Commitment Schemes, WHIR [2] and Blaze [10]. The results in this paper improve the verifier costs of all of these protocols.

1.1 Our Contributions

In this work, we prove a new and stronger notion of correlated agreement [3, 8, 9]. Correlated agreement states the following. For $\pi_L, \pi_R \in \mathbb{F}^n$, if $\Delta_C(\pi_L + r\pi_R, C) < 1-\beta$, then except with negligible probability over choice of $r \leftarrow \mathbb{F}$, there exists $c_L, c_R \in C$ such that

$$|\{i \in [n] : \pi_L[i] = c_L[i] \wedge \pi_R[i] = c_R[i]\}| > (\beta - \epsilon)n,$$

where $\beta, \epsilon \in [0,1]$. In this paper, we prove something stronger - that if $(\pi_L + r\pi_R)[S] \in C[S]$, then except with negligible probability over choice of $r \leftarrow \mathbb{F}$, there exists $c_L, c_R \in C$ such that

$$\pi_L[S] = c_L[S] \text{ and } \pi_R[S] = c_R[S].$$

Actually, we will prove a slight relaxation of the above, where we only guarantee that $\pi_L[S'] = c_L[S']$ and $\pi_R[S'] = c_R[S']$ for some large subset $S' \subset S$ such that $|S'| \geq \beta - \epsilon$. We show that when $\beta > (1-\Delta_C+\epsilon)^{1/3} + \epsilon$, the failure probability is less than $\frac{2}{\epsilon^2|\mathbb{F}|}$. In comparison, the result from [14] only bounds the probability by $O(\frac{1}{\epsilon^3})$, also for $\beta > (1-\Delta_C+\epsilon)^{1/3} + \epsilon$ and the result from [17] only manages to bound this probability by $O(\frac{n^2}{|\mathbb{F}|\epsilon^7})$, albeit with $\beta > \sqrt{1-\Delta_C+\epsilon}$.

Using this result, we prove that BaseFold [25] is secure in the list-decoding regime, meaning that a prover can use BaseFold to commit to a polynomial-size list of polynomials, which requires fewer queries from the IOPP verifier. Ultimately, this reduces the communication complexity of any SNARK that uses BaseFold as its PCS.

[2]There exists a generic transformation for univariate to multilinear PCS, but this incurs further overhead (e.g. [19])

[3]See [11, 25] for more details on this comparison

[4]Actually, recent concurrent work [17] also proves Basefold for RS codes secure in the list decoding regime, and we discuss this in detail in Section 1.2

Technical Overview We now give a high level overview of the proof. We would like to bound the size of a set $A \subset \mathbb{F}$, such that for all $r \in A$, there exists large $S \subset [n]$, where

$$\pi_L[S] + r\pi_R[S] \in C[S], \tag{1}$$

but

$$\pi_L[S], \pi_R[S] \notin C[S]. \tag{2}$$

where $\pi_L[S] = \{\pi_L[i] : i \in S\}$ (resp. for π_R) and $C[S]$ is the puncturing of C on the set S. First, suppose that Equation 2 holds for some set $S \subset [n]$. By linearity of C, this implies that there is only one $r \in \mathbb{F}$ such that $\pi_L[S] + r\pi_R[S] \in C[S]$. This already gives a naive upper bound on the size A, as it is *at most* the number of subsets of $[n]$ that are larger than βn. However, this is clearly insufficient as we would like $|A|/|\mathbb{F}|$ to be negligible.

Our strategy is as follows. First, define a "bad" codeword pair as one where there exists $S \subset [n], r \in \mathbb{F}$ satisfying Equations 1 and 2. We count the number of "bad" codeword pairs and then for each one, we count the number of $r \in \mathbb{F}$ that are witness to it. By a simple counting argument, this will yield the exact size of A. Counting the number of witnesses for a fixed codeword pair is easy, because by linearity of C, a fixed codeword pair can have *at most* $\frac{1}{\epsilon}$ such witnesses. Thus, the key challenge in computing $|A|$ will be in counting the total number of "bad" codeword pairs.

We proceed by labeling all the sets by a codeword pair (c_L, c_R) such that $\pi_L[S] + r\pi_R[S] \in C[S]$ for some $r \in \mathbb{F}$ but $\pi_L[S], \pi_R[S] \notin C[S]$. There are many possible such labelings, and without loss of generality we pick the smallest one. Next, we use an observation from Deep-FRI [8], that if three such sets have large three-way intersection, then there exists a single codeword pair that is a valid label for all three of them. We proceed by systematically iterating through sets in $[n]$, and assigning *the same codeword pair* to sets that share the same *large three-way intersection*. This implies that the number of distinct labels is *less than the number of large subsets of* $[n]$ *with small three-way intersection*. We show that his quantity is less than $\frac{1}{\epsilon}$. The full soundness proof is in Section 3.

1.2 Related Work

In concurrent work entitled "BaseFold In The List Decoding Regime" [17], Haböck proves the same stronger notion of correlated agreement that we do, but with $\beta > \sqrt{1 - \Delta_C + \epsilon}$ and where $|A| \in O(n^2)$. Furthermore, their result only applies to Reed-Solomon codes. In other concurrent work, entitled "Linear Proximity Gap for Linear Codes within the 1.5 Johnson Bound" [14], they also prove stronger correlated agreement, but where $|A| \in O(\frac{1}{\epsilon^3})$ when $\beta > (1 - \Delta_C + \epsilon)^{1/3}$. In Deep-Fold [16], Guo et al adapts Deep-FRI [8] to Basefold, and in that setting also proves the bound of $\beta > \sqrt{1 - \Delta_C}$, but this only applies to Reed-Solomon codes and only to the DeepFRI [8] Protocol, which is a variant of FRI with slightly more overhead.

Additionally, there are several papers [3, 8, 9] that analyze the communication complexity of FRI [6] using correlated agreement. In "Worst Case To Average Case Reduction For Distance to a Linear Code" [3], the authors improve upon the original FRI paper by proving that $|A|$ is small when $\beta \geq (1 - \Delta_C)^{1/4}$ and Deep-FRI [8] improves this to $\beta \geq (1 - \Delta_C + \epsilon)^{1/3}$. Deep-FRI additionally introduces a modification of the FRI protocol which further reduces the number of verifier queries, at the cost of some (slight) prover overhead. Finally, in "Proximity Gaps of Reed-Solomon Codes" [9], the authors use a list-decoding algorithm for Reed-Solomon codes to prove correlated agreement for $\beta \geq \sqrt{1 - \Delta_C + \epsilon}$, but (as mentioned earlier with regards to [17]), this result is only meaningful with a field that is at least quadratic in the instance size and with a suitably small rate, both of which impact prover time.

2 Preliminaries

2.1 Notation

Sets Let $n \in \mathbb{Z}$. Denote by $[n]$ the set $[0, n-1]$. For a set S, $\mathsf{even}(S)$ is the set of even integers in S, and $\mathsf{odd}(S)$ is the set of odd integers in S. Let $q \in \mathbb{N}$. Then $S/q = \{s/q : s \in S\}, S + q = \{s + q : s \in S\}$, etc. 2^S is the power set of S.

Strings and Functions Let $x \in \mathbb{F}, r \in \mathbb{N}$, then $x^{||r}$ is the string obtained by concatenating x to itself r times. Let $f : S \to S$ be a function and $n \in \mathbb{N}$. Then $f^{\circ n}$ denotes function composition of f with itself n times.

Error-Correcting Codes We will use C to denote a linear $[n, k, d]$ code, which is a subspace C of $\mathbb{F}^n$ with an encoding algorithm $\mathsf{Enc}_C : \mathbb{F}^k \to C$ (Definition 1). Δ_C is the minimum relative distance of the code C. Let $n \in \mathbb{N}$ and $S \subset [n]$. For a vector $\mathbf{x} \in \mathbb{F}^n$, $\mathbf{x}[S] = \{x[i] : i \in S\}$. Let $\mathbf{v} \in \mathbb{F}^n$, let C be a linear code, and let $S \subset [n]$. Then we say that $\mathbf{v}[S] \in C[S]$ if there exists a codeword $c \in C$ such that $\mathbf{v}[S] = c[S]$.

2.2 Definitions

We present a standard definition of a *linear error-correcting code.*

Definition 1 (Linear Code)**.** *A linear error-correcting code with message length k and codeword length n is an injective mapping from $\mathbb{F}^k$ to a linear subspace $C \subseteq \mathbb{F}^n$. C is associated with a generator matrix, $G \in \mathbb{F}^{k \times n}$ such that the rows of G are a basis of C and the encoding of a vector $\mathbf{v} \in \mathbb{F}^k$ is $\mathbf{v} \cdot G$. The minimum Hamming distance of a code is the minimum on the number of different entries between any two different codewords $c_1, c_2 \in C$. If C has a minimum distance $d \in [n]$, we say that C is an $[n, k, d]$ code and use Δ_C to denote d/n—the relative minimum distance.*

3 A Stronger Notion of Correlated Agreement

In this section, we state and prove our main result. We compute concrete bounds in subsection 3.2. Our main result is a stronger version of the correlated agreement theorem (Theorem 1.4) from "Proximity Gaps of Reed Solomon Codes" [9].

3.1 Strong Correlated Agreement Within The One-And-A-Half Johnson Bound

Theorem 1 (Strong Correlated Agreement A)**.** *Let C be a linear error-correcting code with $n \in \mathbb{N}$, $\epsilon \in [0, 1]$, and $\pi_L, \pi_R \in \mathbb{F}^n$. Let $\beta \geq (1 - \Delta_C + \epsilon)^{1/3} + \epsilon$. If*

$$\Pr_{r \in \mathbb{F}}[\Delta(\pi_L + r\pi_R, C) \leq (1 - \beta)n] \geq \frac{2}{\epsilon^2 |\mathbb{F}|},$$

then there exists $S' \subset S \subset [n]$ and $c_L, c_R \in C$ satisfying

- ***Density**: $|S| \geq \beta n, |S'| \geq (\beta - \epsilon)n$*
- ***Agreement**: $\pi_L[S'] = c_L[S'], \pi_R[S'] = c_R[S']$ and $\forall r \in \mathbb{F}, \pi_L[S] + r\pi_R[S] = c_L[S] + rc_R[S]$*

Actually, we will find it more useful to prove the following stronger statement, from which Theorem 1 easily follows.

Theorem 2 (Strong Correlated Agreement B)**.** *Let $\pi_L, \pi_R \in \mathbb{F}^n$ with $\pi = (\pi_L, \pi_R)$ and $\epsilon \in [0,1]$. Define $\beta_\epsilon \geq (1 - \Delta_C + \epsilon)^{1/3} + \epsilon$. Let $\pi = (\pi_L, \pi_R)$ and define the set $A_\pi(\epsilon)$[5] as follows.*

$$\left| \left\{ r \in \mathbb{F} : \exists S \subset [n], c \in C \text{ st} \quad \begin{array}{l} |S| > \beta_\epsilon n, \\ (\pi_L[S] + r\pi_R[S]) = c[S], \\ \text{ but } \forall (c_L, c_R) \in C \times C, \\ |\{i \in S : \pi_L[i] = c_L[i] \wedge \pi_R[i] = c_R[i]\}| < (\beta_\epsilon - \epsilon)n \end{array} \right\} \right|.$$

Then $\forall \pi \in \mathbb{F}^{2n}$,

$$|A_\pi(\epsilon)| \leq \frac{2}{\epsilon^2} \tag{3}$$

Remark 1. *There are a few other constraints on β and ϵ that we need for the statement to hold. For ease of exposition, we do not state them here and defer their details to the proof of Lemma 1 in Appendix A. We emphasize that these constraints hold for most values of β, ϵ.*

Proof. First, we define a subset $A_{\pi,\geq\beta_\epsilon} \subseteq A_\pi(\epsilon)$ as follows, where modifications from $A_\pi(\epsilon)$ are in bold.

$$\left| \left\{ r \in \mathbb{F} : \exists S \subset [n], \mathbf{c_L^*}, \mathbf{c_R^*}, c \in C \text{ st} \quad \begin{array}{l} |S| > \beta_\epsilon n, \\ (\pi_L[S] + r\pi_R[S]) = c[S], \\ \mathbf{c_L^* + r c_R^* = c} \\ \mathbf{|\{i \in [n] : c_L^*[i] = \pi_L[i] \wedge c_R^*[i] = \pi_R[i]\}| \geq \beta_\epsilon n} \\ \text{ but } \forall (c_L, c_R) \in C \times C, \\ |\{i \in S : \pi_L[i] = c_L[i] \wedge \pi_R[i] = c_R[i]\}| < (\beta_\epsilon - \epsilon)n \end{array} \right\} \right|,$$

and define $A_{\pi,<\beta_\epsilon} := A_{\pi,\epsilon} \setminus A_{\pi,\geq\beta_\epsilon}$. By definition of set compliment,

$$|A_{\pi,\epsilon}| = |A_{\pi,\geq\beta_\epsilon}| + |A_{\pi,<\beta_\epsilon}|. \tag{4}$$

Thus our task reduces to bounding the size of each of these individual sets. To do this, we will break each set down further by considering the list of pairs of codewords that "explain" the elements of $A_{\pi,<\beta_\epsilon}, A_{\geq\beta_\epsilon}$. To that end, we introduce two new sets.

Definition 2 ($\mathcal{F}_{<\beta_\epsilon}, \mathcal{F}_{\geq\beta_\epsilon}$)**.** *Let $\mathcal{L} \in \mathcal{F}_{<\beta_\epsilon}$. Then $\mathcal{L} \subset C \times C$ such that $\forall r \in A_{\pi,<\beta_\epsilon}$, there exists $(c_L, c_R) \in \mathcal{L}$ and $S \subset [n]$, such that*

$$c_L[S] + rc_R[S] = \pi_L[S] + r\pi_R[S]. \tag{5}$$

We define $\mathcal{F}_{\geq\beta_\epsilon}$ analogously, where if $\mathcal{L} \in \mathcal{F}_{\geq\beta_\epsilon}$ then for all $r \in A_{\pi,\geq\beta_\epsilon}$, there exists $(c_L, c_R) \in \mathcal{L}$ and $S \subset [n]$ satisfying Equation 5. We will denote by $\mathcal{L}_{<\beta_\epsilon}, \mathcal{L}_{\geq\beta_\epsilon}$ the smallest sets in $\mathcal{F}_{<\beta_\epsilon}, \mathcal{F}_{\geq\beta_\epsilon}$, respectively.

Next, for each $(c_L, c_R) \in C \times C$, define the following set, which deals with individual codeword pairs.

$$A_{\pi,\mathbf{c_L},\mathbf{c_R}} = \left| \left\{ r \in \mathbb{F} : \exists S \subset [n_i], c \in C \text{ st} \quad \begin{array}{l} |S| > \beta_\epsilon n, \\ (\pi_L[S] + r\pi_R[S]) = c[S], \\ \mathbf{c_L + r c_R = c} \\ \text{ but } \{\pi_R[S], \pi_L[S]\} \not\subset C[S] \end{array} \right\} \right|.$$

Then,

$$A_{\pi,\geq\beta_\epsilon} = \bigcup_{(c_L,c_R)\in\mathcal{L}_{\geq\beta_\epsilon}} A_{\pi,(c_L,c_R)}, \tag{6}$$

[5]The set $A_\pi(\epsilon)$ is a more formal version of A from the technical overview (Section 1.1).

and

$$A_{\pi,<\beta_\epsilon} = \bigcup_{(c_L,c_R)\in\mathcal{L}_{<\beta_\epsilon}} A_{\pi,(c_L,c_R)}. \tag{7}$$

Therefore,

$$|A_\pi| = |A_{\pi,<\beta_\epsilon} \cup A_{\pi,\geq\beta_\epsilon}| = \sum_{(c_L,c_R)\in\mathcal{L}_{<\beta_\epsilon}\cup\mathcal{L}_{\geq\beta_\epsilon}} |A_{\pi,(c_L,c_R)}| \tag{8}$$

Thus, our task reduces to bounding the size of the following three quantities:

1. $|\mathcal{L}_{\geq\beta_\epsilon}|$
2. $|\mathcal{L}_{<\beta_\epsilon}|$
3. $|A_{\pi,(c_L,c_R)}|$ for each $(c_L, c_R) \in \mathcal{L}_{\geq\beta_\epsilon} \cup \mathcal{L}_{<\beta_\epsilon}$

We will use the following Lemma to bound the size of items (1) and (2).

Lemma 1. *Let $\beta \in [0,1], n \in \mathbb{Z}$. For each $x \in \{2,3\}$, define $\mathcal{S}(x,\beta) \subset 2^{[n]}$ as follows.*

- *If $S \in \mathcal{S}(x,\beta)$, then $|S| > \beta n$*
- *For any x sets $S_1, .., S_x \in \mathcal{S}$, there exists $\epsilon \in [0,1]$ such that*

$$|\bigcap_{i\in[x]} S_i| < (\beta^x - \epsilon)n$$

.

Then,

$$|\mathcal{S}(x,\beta)| \leq \frac{1}{\epsilon}.$$

Proof. We defer the proof to Appendix A as it follows the proof of the Johnson Bound[6] closely. □

Lemma 2.

$$|\mathcal{L}_{\geq\beta_\epsilon}| \leq \frac{1}{\epsilon}.$$

Proof. For each $(c_L, c_R) \in \mathcal{L}_{\geq\beta_\epsilon}$, let $S_{c_L,c_R} = \{i \in [n] : \pi_L[i] = c_L[i] \wedge \pi_R[i] = c_R[i]\}$. Then we define,

$$\mathcal{S} = \{S_{c_L,c_R} : (c_L, c_R) \in \mathcal{L}_{\geq\beta_\epsilon}\}.$$

By definition of $\mathcal{L}_{\geq\beta_\epsilon}$, each set in $\mathcal{S}$ is larger than β_ϵ. Next, we show that any two sets in $\mathcal{S}$ must have pairwise intersection *smaller* than $1 - \Delta_C$. Let (c'_L, c'_R) be an element of $\mathcal{L}_{\geq\beta_\epsilon}$ *distinct* from (c_L, c_R). Then,

$$\begin{aligned}
&|S_{c_L,c_R} \cap S_{c'_L,c'_R}| \\
&= |\{i \in [n] : i \in S_{c_L,c_R} \wedge i \in S_{c'_L,c'_R}\}| \\
&= |\{i \in [n] : \pi_L[i] = c_L[i] \wedge \pi_R[i] = c_R[i] \wedge \pi_L[i] = c'_L[i] \wedge \pi_R[i] = c'_R[i]\}| \\
&\leq |\{i \in [n] : \pi_L[i] = c_L[i] \wedge \pi_L[i] = c'_L[i]\}| \\
&\leq (1 - \Delta_C)n \text{ (by definition of minimum distance)}
\end{aligned}$$

Since $(1 - \Delta_C + \epsilon)^{1/2} < (1 - \Delta_C + \epsilon)^{1/3} < \beta_\epsilon$, it follows that $(1 - \Delta_C) \leq (\beta_\epsilon^2 - \epsilon)$ and so we conclude that for any two distinct sets $S_{c_L,c_R}, S_{c'_L,c'_R} \in \mathcal{S}$, the following holds.

$$|S_{c_L,c_R} \cap S_{c'_L,c'_R}| \leq (1 - \Delta_C)n \leq (\beta_\epsilon^2 - \epsilon)n.$$

[6] `https://www.cs.cmu.edu/~venkatg/teaching/au18-coding-theory/lec-scribes/list-decoding.pdf`

Therefore, by Lemma 1, $|\mathcal{S}| < \frac{1}{\epsilon}$. Finally, it is clear by the definition of $\mathcal{S}$ that $|\mathcal{L}_{\geq\beta_\epsilon}| = |\mathcal{S}|$, which completes the proof. □

Lemma 3.

$$|\mathcal{L}_{<\beta_\epsilon}| \leq \frac{1}{\epsilon}.$$

Proof. Define $\mathcal{S} \subset 2^{[n]}$ to be the set such that $\forall r \in A_{\pi,<\beta_\epsilon}, \exists S \in \mathcal{S}$, and codeword pair $(c_L, c_R) \in C \times C$, such that

$$c_L[S] + rc_R[S] = \pi_L[S] + r\pi_R[S] \tag{9}$$

We show that there exists $\mathcal{L} \in \mathcal{F}_{<\beta_\epsilon}$ (Definition 2) that is *smaller* than $\frac{1}{\epsilon}$. Since $\mathcal{L}_{<\beta_\epsilon}$ is defined as the smallest set in $\mathcal{F}_{<\beta_\epsilon}$, this is enough to prove the Lemma. To define $\mathcal{L}$, we define a graph $G = (V, E)$ as follows. Every vertex in V is labeled by a set in $\mathcal{S}$ and we denote that set by $l(v)$. For any $v_1, v_2 \in V$, the edge (v_1, v_2) is in E if $|l(v_1) \cap l(v_2)| \geq (1 - \Delta_C)n$. A useful observation is that any three sets with large *three-way intersection* will form a 3-cycle in G. Next, let $G' = (V', E')$ be the graph where every $v' \in V'$ is labeled by a 3-cycle (v_1, v_2, v_3) in G, denoted by $l'(v')$. (v'_1, v'_2) is in E' if the two 3-cycles, $l'(v'_1)$ and $l'(v'_2)$, share an edge in G. To construct $\mathcal{L}$, we present the following protocol. This protocol also constructs an auxiliary set $\mathcal{S}'$, which we will use later on in the proof.

Protocol 1 `Construct` L

1. Initialize $\mathcal{S}' := \emptyset, \mathcal{L} := \emptyset$
2. For each unvisited $v \in V'$, do `label`(`NULL`, v).
3. For each $w \in V$ that is *not* contained in a 3-cycle in G do the following:
   - Solve for some $(c_L, c_R) \in C \times C$ and $r \in \mathbb{F}$ such that

     $$c_L[l(w)] + rc_R[l(w)] = \pi_L[l(w)] + r\pi_R[l(w)]$$

     .
   - Add (c_L, c_R) to $\mathcal{L}$ and $l(w)$ to $\mathcal{S}'$.

Next, we prove that $\mathcal{L} \in \mathcal{F}_{<\beta}$. To prove this, we need to show that for all $r \in A_{\pi,<\beta}$, there exists $S \subset [n]$ and $(c_L, c_R) \in \mathcal{L}$ such that Equation 9 holds, i.e. such that

$$c_L[S] + rc_R[S] = \pi_L[S] + r\pi_R[S].$$

We will make use of the following claim, which we prove at the end of this lemma's proof.

Claim 1. *$\forall S \in \mathcal{S}$, there exists $(c_L, c_R) \in \mathcal{L}$ such that Equation 9 holds for some $r \in A_{\pi,<\beta}$.*

By definition of $\mathcal{S}$, $\forall r \in A_{\pi,<\beta}$, $\exists S \in \mathcal{S}$ such that $|S| > \beta_\epsilon n$ and Equation 9 holds for some $(c_L, c_R) \in C \times C$. By Claim 1, there exists $(c'_L, c'_R) \in \mathcal{L}$ such that Equation 9 holds for some $r' \in A_{\pi,\beta}$. We now need to prove that $r' = r$. To do that, we state another small claim, whose proof we also defer to the end of this lemma's proof.

Claim 2. *Let $\pi_L, \pi_R \in \mathbb{F}^n$. Let $S \subset [n]$. Either $\pi_L[S], \pi_R[S] \in C[S]$ or there is exactly one $r \in \mathbb{F}$ such that*

$$\pi_L[S] + r\pi_R[S] \in C[S]$$

Protocol 2 `label`

1. Inputs: `codeword_pair` $\in C \times C$, vertex $v \in V'$.
2. If `codeword_pair` is NULL, then
   - Set `codeword_pair` to be the pair (c_L, c_R), satisfying
   $$\pi_L[S_x] + r_x \pi_R[S_x] = c_L[S_x] + r_x c_R[S_x] \tag{10}$$
   for each $x \in \{1, 2, 3\}$, where $S_x = l(l'(v)[x])$ and $r_x \in \mathbb{F}$.
   - Add S_1 to $\mathcal{S}'$ and `codeword_pair` to $\mathcal{L}$.
3. For each $w \in V'$ such that $(v, w) \in E$, invoke `label`(`codeword_pair`, w).

It follows directly from Claim 2 that $r' = r$. Thus, we can conclude that $\mathcal{L} \in \mathcal{F}_{<\beta}$. Next, we argue that $|\mathcal{L}| \leq \frac{1}{\epsilon}$. We show that $|\mathcal{S}'| \leq \frac{1}{\epsilon}$ and that $|\mathcal{L}| = |\mathcal{S}'|$. By construction of Protocol 1, if $l(v_1), l(v_2), l(v_3) \in \mathcal{S}'$, then v_1, v_2, v_3 must not share a 3-cycle in G, and therefore $|l(v_1) \cap l(v_2) \cap l(v_3)| < (1-\Delta_C)n$. Since $\beta_\epsilon \geq (1-\Delta_C+\epsilon)^{1/3}+\epsilon$, it follows that $(1-\Delta_C) \leq \beta_\epsilon^3 - \epsilon$. Therefore, by Lemma 1, $|\mathcal{S}'| < \frac{1}{\epsilon}$. Next, by construction of Protocol 1, a set is added to $\mathcal{S}'$ if and only if it is added to $\mathcal{L}$, and therefore $|\mathcal{S}'| = |\mathcal{L}|$. Since $\mathcal{L}_{\pi,<\beta}$ is the smallest set in $\mathcal{F}_{<\beta_\epsilon}$, it is smaller than L, and therefore,

$$|\mathcal{L}_{\pi,<\beta}| \leq |L| < \frac{1}{\epsilon}.$$

Finally, to complete the proof, we prove the remaining claims.

Proof of Claim 1. To prove this claim, we need to prove the following two statements.

1. A codeword pair $(c_L, c_R) \in C \times C$ is added to $\mathcal{L}$ for every set in $\mathcal{S}$
2. Every codeword pair added to $\mathcal{L}$ satisfies Equation 9.

To prove item (1), note that in Protocol 1, we process every vertex that does not exist in a 3-cycle in G and in Protocol 2, we process every vertex that *does* exist in a 3-cycle in G. Therefore, a codeword pair is added to $\mathcal{L}$ for each vertex $v \in V$, (and therefore for each $S \in \mathcal{S}$ (by definition of G)). To prove 2) we consider the two steps where a codeword pair is added to $\mathcal{L}$; Step 3 in Protocol 1 and Step 3 in Protocol 2. In the first case, the pair clearly satisfies Equation 9. In the second case, it satisfies Equation 9 as long as (c_L, c_R) in Step 2 of Protocol 2 exists and satisfies Equation 10. The proof that this exists follows directly from the following claim, whose proof we defer to the end.

Claim 3 (Intersecting Foldings are Co-linear)**.** *Let $r_1, r_2, r_3 \in \mathbb{F}$, $S_1, S_2, S_3 \subset [n]$ and $\beta \in [0, 1]$ such that $\beta > (1 - \Delta_C)$. Let $S_I = \bigcap_{x \in \{1,2,3\}} S_x$ and suppose that $|S_I| > \beta n$. Suppose that for all $x \in \{1, 2, 3\}$,*

$$\pi_L[S_x] + r_x \pi_R[S_x] \in C[S_x].$$

Then there exists a unique pair of codewords (c_L, c_R) such that $\pi_L[S_I] = c_L[S_I]$, $\pi_R[S_I] = c_R[S_I]$ and for all $x \in \{1, 2, 3\}$,

$$\pi_L[S_x] + r_x \pi_R[S_x] = c_L[S_x] + r_x c_R[S_x].$$

□

Next, we prove Claim 2.

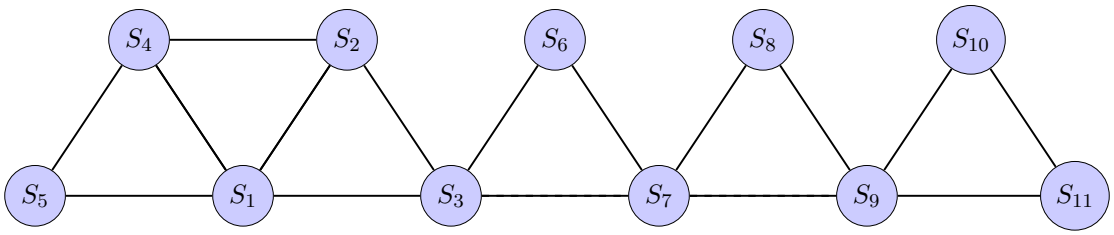

Figure 1: Representation of the graph G, where an edge connects two vertices if their labeled sets have large intersection.

Proof of Claim 2. Suppose otherwise. Then there exists $r_1, r_2 \in \mathbb{F}$ and $c_1, c_2 \in C$ such that

$$\pi_L[S] + r\pi_R[S] = c_1[S],$$

$$\pi_L[S] + r'\pi_R[S] = c_2[S],$$

and such that $\pi_L[S], \pi_R[S] \notin C[S]$. By linearity of the code, we subtract the bottom equation from the top equation and obtain

$$(r - r') \cdot (\pi_R[S]) = (c_1 - c_2)[S].$$

Again by linearity of the code, we divide both sides by $(r_1 - r_2)$, which proves that

$$\pi_R[S] = \frac{(c_1 - c_2)}{r - r'}[S].$$

Therefore, $\pi_R[S] \in C[S]$. Furthermore, plugging this back into Equation 1, we find that $\pi_L[S] \in C[S]$, which leads to a contradiction and completes the proof. □

Finally, we prove Claim 3.

Proof Of Claim 3. By assumption, there exists $c_1, c_2 \in C$ such that

$$\pi_L[S_1] + r_1\pi_R[S_1] = c_1[S_1],$$

$$\pi_L[S_2] + r_2\pi_R[S_2] = c_2[S_2].$$

Let $S'_I = S_1 \cap S_2$. Then since $S'_I \subset S_1, S_2$,

$$\pi_L[S'_I] + r_1\pi_R[S'_I] = c_1[S'_I],$$

and

$$\pi_L[S'_I] + r_2\pi_R[S'_I] = c_2[S'_I].$$

By Claim 2, this implies that $\pi_L[S'_I], \pi_R[S'_I]$ are in $C[S'_I]$. Since $S_I \subset S_{I'}$, it follows that $\pi_L[S_I], \pi_R[S_I] \in C[S_I]$ and therefore there exists $v, v^* \in C$ such that $\pi_L[S_I] = v[S_I]$, $\pi_R[S_I] = v^*[S_I]$. Finally, for each $x \in \{1, 2, 3\}$, $\pi_L[S_I] + r_x\pi_R[S_I] = v[S_I] + r_x v^*[S_I]$. Since $|S'_I| \geq |S_I| \geq \beta n \geq (1 - \Delta_C)^{1/3}n > (1 - \Delta_C)n$, it follows from definition of minimum distance of the code that for each $x \in \{1, 2, 3\}$,

$$\pi_L[S_x] + r_x\pi_R[S_x] = v[S_x] + r_x v^*[S_x],$$

which completes the proof. □

□

Finally, bound the size of each $A_{\pi,(c_L,c_R)}$.

Lemma 4. *For all* $(c_L, c_R) \in C \times C$,

$$|A_{\pi,(c_L,c_R)}| \leq \frac{1}{\epsilon}$$

Proof. Define $\neg S_{c_L,c_R} \subset [n]$ as

$$\neg S_{c_L,c_R} = \{i \in [n] : c_L[i] \neq \pi_L[i] \wedge c_R[i] \neq \pi_R[i]\}.$$

For each $r \in A_{\pi,(c_L,c_R)}$, let S_r be the maximal subset of $\neg S_{c_L,c_R}$ such that $\pi_L[S_r] + r\pi_R[S_r] = c_L[S_r] + rc_R[S_r]$. Define $\mathcal{S} \subset 2^{[n]}$ as,

$$\mathcal{S} = \{S_r : r \in A_{\pi,(c_L,c_R)}\}.$$

We use the following claim, whose proof we defer.

Claim 4. *Let* $\mathbf{c}_L, \mathbf{c}_R \in C$, *and* $\pi_L, \pi_R \in \mathbb{F}^n$. *Let* $i \in [n]$ *where* $\pi_L[i] \neq \mathbf{c}_L[i], \pi_R[i] \neq c_R[i]$. *Then there is exactly one* $r \in \mathbb{F}$ *satisfying*

$$\pi_L[i] + r\pi_R[i] = \mathbf{c}_L[i] + rc_R[i]$$

By Claim 4, every two sets in $\mathcal{S}$ are *disjoint*. Therefore

$$|\bigcup_{S \in \mathcal{S}} \mathcal{S}| = \sum_{S \in \mathcal{S}} |S| \leq n.$$

For each $r \in \mathbb{F}$, $|S_r| > \epsilon n$, because otherwise π_L, π_R both agree with codewords at $\geq \beta_\epsilon - \epsilon$ locations and so $r \notin A_\pi$. Therefore,

$$|\mathcal{S}| \cdot \epsilon n \leq \sum_{S \in \mathcal{S}} |S|,$$

and so, combining the previous two equations,

$$|\mathcal{S}| \leq \frac{\sum_{S \in \mathcal{S}} |S|}{\epsilon n} \leq \frac{n}{n\epsilon} = \frac{1}{\epsilon}$$

Since there is a one-to-one relationship between $A_{\pi,(c_L,c_R)}$ and $\mathcal{S}$, we have

$$|A_{\pi,(c_L,c_R)}(\epsilon)| = |\mathcal{S}| \leq \frac{1}{\epsilon}.$$

Finally, we prove Claim 4.

Proof. Suppose otherwise. Then the following two equations are true for distinct $r_1, r_2 \in \mathbb{F}$:

$$\pi_L[i] - c_L[i] + r_1(\pi_R[i] - \mathbf{c}_2[i]) = 0$$

$$\pi_L[i] - \mathbf{c}_1[i] + r_2(\pi_R[i] - \mathbf{c}_R[i]) = 0$$

Let $P(X) = (\pi_L[i] - \mathbf{c}_L[i]) + X(\pi_R[i] - \mathbf{c}_R[i])$. Since $\pi_L[i] \neq \mathbf{c}_L[i], \pi_R[i] \neq c_R[i]$, it follows that $P(X)$ is a non-zero, degree-one, polynomial. Therefore, by the Schwartz-Zippel Lemma, it has only one zero and so it is not possible that $P(r_1) = P(r_2) = 0$. This is a contradiction and completes the proof. □

□

Combining these three bounds with Equation 8, we have

$$|A_\pi| \leq \sum_{(c_L,c_R) \in \mathcal{L}_{<\beta_\epsilon} \cup \mathcal{L}_{\geq \beta_\epsilon}} \frac{1}{\epsilon} = (\frac{2}{\epsilon}) \cdot \frac{1}{\epsilon} = \frac{2}{\epsilon^2},$$

which completes the proof of Theorem 2. □

3.2 Concrete Bounds and Comparison To Other Work

In the BaseFold IOPP, the verifier needs to check that a) $\pi_L + r\pi_R$ is within the list-decoding radius of some $c \in C$, and that b) π_L, π_R are within the list-decoding radius of c_L, c_R such that $c_L + rc_R = c$. Theorem 2 implies that the verifier needs to check more than β_ϵ locations, where

$$\beta_\epsilon \geq (1 - \Delta_C + \epsilon)^{1/3} + \epsilon.$$

Actually, the verifier will only make a constant number of queries. If no $\beta_\epsilon n$-sized set exists satisfying the above conditions, then each query is a Bernulli trial with success rate $< \beta_\epsilon$, and so after q trials, the probability of verifier acceptance is $< (\beta_\epsilon)^q$. Moreover, if the verifier accepts all q queries, then the statement holds accept with probability

$$\leq \frac{2}{\epsilon^2}|\mathbb{F}|.$$

Next, we compare to the results from [17] and [14]. The bound from [17] over Reed-Solomon codes (translated to our notation and for a batch of only 2 polynomials) achieves

$$|A_\pi| \leq 2\frac{(m + 1/2)}{\sqrt{1 - \Delta_C}} \cdot \max\left(\frac{(m + 1/2)^6}{3(1 - \Delta_C)} \cdot n^2, 2 \cdot (B \cdot n + 1)\right),$$

where m is a parameter larger than 3 and for verifier query complexity

$$\beta > \sqrt{(1 - \Delta_C)}(\frac{1}{2m}),$$

i.e. the verifier only needs to query π_L, π_R in $\beta > \sqrt{1 - \Delta}(\frac{1}{2m})$ locations (which is better than our bound), but will be *incorrect* with the probability of $\frac{O(n^2)}{|\mathbb{F}|}$ (which is worse than our bound). The authors from [14], on the other hand, prove that for $\beta_\epsilon \geq (1 - \Delta_C + \epsilon)^{1/3} + \epsilon$, the failure probability is only in $O(\frac{2}{\epsilon^2|\mathbb{F}|})$. We show a comparison of concrete results in Figure 1.

Remark 2. *It has been proven in a blog post*[7] *that the bound of $\beta > 1 - \Delta_C/3$ is tight for general linear codes, which seems to contradict the results in this paper. However, upon closer inspection, they only consider codes where $1 - \Delta_C/3$ is always less than $\sqrt{1 - \Delta_C}$. Thus, there is no contradiction after all, as we do not expect the Correlated Agreement/Proximity Gaps statement to hold for $\beta < (1 - \Delta_C)^{1/3}$, let alone $\beta < \sqrt{1 - \Delta_C}$.*

4 Improved Soundness of the BaseFold Protocol

In this section, we re-prove the soundness theorem of BaseFold, and show that soundness holds even if the verifier only makes $l := \frac{\lambda}{\log_2(\beta_\epsilon)}$ queries for $\beta_\epsilon \geq (1 - \Delta + \epsilon)^{1/3} + \epsilon$. Previously, we only proved this for $\beta_\epsilon \geq 1 - \Delta/3 + \epsilon$. The BaseFold IOP remains unchanged from [25][8]. We restate the IOPP in Figure 2 for completeness. We also restate the definition of a *foldable code*[9], which was introduced in BaseFold [25]. A *foldable code* is a family of codes, which we will denote $(C_d, .., C_0)$, which are characterized by a set of vectors $\{\mathbf{t}_i \in \mathbb{F}^{n_i} : i \in [1, d]\}$. Each codeword $c_i \in C_i$ is composed of two codewords in C_{i-1}. Additionally, the structure of a *foldable codes* enables local consistency checks between codewords in adjacent codes. These consistency tests allow the BaseFold[10] IOPPs to maintain logarithmic query complexity.

[7] https://notes.0xparc.org/results/counterexample-proximity-gap/

[8] The syntax of this description is slightly different than that in [25], but the protocol itself is equivalent

[9] For ease of exposition, we define the codes according to a different ordering than the original.

[10] This structure also underlies certain types of Reed-Solomon codes and is the reason for FRI's efficiency

Distance	$(\lvert A_\pi\rvert/\lvert\mathbb{F}\rvert)$	β_ϵ
Hab24		
3/4	2^{-53}	$2^{-0.77}$
7/8	2^{-51}	$2^{-1.2}$
15/16	2^{-49}	$2^{-1.7}$
31/32	2^{-48}	$2^{-2.27}$
GKL24		
3/4	2^{-65}	$2^{-0.66}$
7/8	2^{-64}	2^{-1}
15/16	2^{-63}	$2^{-1.32}$
31/32	$2^{-62.5}$	$2^{-1.65}$
This result		
3/4	2^{-111}	$2^{-0.66}$
7/8	2^{-111}	2^{-1}
15/16	2^{-111}	$2^{-1.32}$
31/32	2^{-111}	$2^{-1.65}$

Table 1: We consider instance sizes of 2^{30} over finite field, $\mathbb{F}$ such that $\log_2(|\mathbb{F}|) = 128$. We report the number of Basefold verifier repetitions needed to achieve a soundness error below 2^{-100}. For our result and the result from [14], we set $\epsilon = 0.0005$. For [17], we set $m = 3$ to minimize $|A_\pi|/|\mathbb{F}|$. Smaller is better in all three categories.

Definition 3 (Foldable Code)**.** *Let $c, d \in \mathbb{N}$ and for each $i \in [d]$, define $k_i = 2^i, n_i = c \cdot k_i$. A $[n_d, k_d]$, foldable code is a family of codes $(C_0, .., C_d)$, where the base code, C_0, is equal to the repetition code, $\{m^{||c} : m \in \mathbb{F}\}$ and each $[n_i, k_i]$ code, C_i, and each $\mathbf{v} \in C_i$ satisfies the following:*

$$\begin{aligned}\mathbf{v} := Enc_{C_i}(\mathbf{m}) = Enc_{C_{i-1}}(\mathbf{m}_L) + \mathbf{t}_i \circ Enc_{C_{i-1}}(\mathbf{m}_R)\\ ||Enc_{C_{i-1}}(\mathbf{m}_L) - \mathbf{t}_i \circ Enc_{C_{i-1}}(\mathbf{m}_R)\end{aligned}$$

where $\{\mathbf{t}_i \in \mathbb{F}^{n_i} : i \in [1, d]\}$ is given in the description of the code, $\mathbf{m} = (\mathbf{m}_L||\mathbf{m}_R)$ is a vector in $\mathbb{F}^{k_i}$, and $\circ$ denotes the Hadamard product.

Remark 3. *In the above definition, elements $\mathbf{v}[i]$ and $\mathbf{v}[i + n_i/2]$ are two points on the same line for each $i \in [n_i/2]$. For the remainder of this section, we will assume that the codeword has been re-ordered, so that $\mathbf{v}[i], \mathbf{v}[i + 1]$ are on the same line for each $i \in \mathsf{even}([n_i])$. It is easy to prove that folding preserves this ordering.*

Foldable codes are attractive because a codeword in C_i can be transformed into a smaller codeword in C_{i-1} using only local operations. More specifically, we query the same random point on each of the $n_i/2$ lines defined by the pairs $\{(\mathbf{v}[j], \mathbf{v}[j + 1]) : j \in \mathsf{even}([n_i])$, and obtain a new codeword in C_{i-1}. We describe this formally with the following definition.

Definition 4 (Fold)**.** *Define $\mathsf{interp} : \mathbb{F}^2 \times \mathbb{F}^2 \to \mathbb{F}[X]$ to be Lagrange Interpolation of a degree-one univariate polynomial. Let $(C_d, .., C_0)$ be a family of foldable codes characterized by $\{\mathbf{t}_i \in$*

$\mathbb{F}^{n_i} : i \in [1,d]\}$. *For each* $\mathbf{v} \in C_i$, *and* $j \in \mathsf{even}([n_i])$, *define the pair* (p_j, p_{j+1}) *as follows.*

$$(p_j, p_{j+1}) = ((\mathbf{t}[j], \mathbf{v}[j]), (-\mathbf{t}[j], \mathbf{v}[j+1])).$$

Then, the fold *of* $\mathbf{v}$ *with respect to* $r \in \mathbb{F}$ *is the vector,* $\mathsf{fold}(\mathbf{v}, r)[j]$ *satisfying,*

$$\mathsf{fold}(\mathbf{v}, r)[j] = \mathsf{interp}(p_j, p_{j+1})(r).$$

At times, we will need to work with the univariate polynomials defined by $\mathsf{interp}(p_j, p_{j+1})$ directly. We call these polynomials the *unfolding* of $\mathbf{v}$. To ease exposition, we denote by $\mathbf{v}_L, \mathbf{v}_R$ the codewords that for all $j \in \mathsf{even}([n_i])$ satisfy,

$$\mathsf{interp}(p_j, p_{j+1}) = \mathbf{v}_L[j] + X\mathbf{v}_R[j]. \tag{11}$$

We remark that fold can also be defined over arbitrary vectors that are not codewords, and indeed the FRI and Basefold IOPPs rely on this fact. For a generic $\pi \in \mathbb{F}^{n_i}$, define the pair of points $(p_j, p_{j+1}) = ((\mathbf{t}[j], \pi[j]), (-\mathbf{t}[j], \pi[j+1]))$. Then, as before $\mathsf{fold}(\pi, r) = \mathsf{interp}(p_j, p_{j+1})$. Finally, we will sometimes fold over entire sets $S \subset [n_i]$, and this operation is well defined as long as S contains $j+1$ whenever it contains j. We introduce additional notation for this as follows.

$$\mathsf{fold}(\pi, r)[S] = \{\mathsf{interp}((\mathbf{t}_i[j], \pi[j]), (-\mathbf{t}_i, \pi[j+1]))(r) : j \in \mathsf{even}(S)\} \tag{12}$$

Theorem 3 (Soundness of Basefold IOP)**.** *Let* $\lambda \in \mathbb{N}$ *be a security parameter,* $\pi_d \in \mathbb{F}^{n_d}$, $l \in \mathbb{N}$, *and* $\epsilon \in [0,1]$, *with* $\beta_\epsilon > (1 - \Delta_C + \epsilon)^{1/3} + d\epsilon$, *and* $\beta_\epsilon^l \leq \mathsf{negl}(\lambda)$. *Then, with probability greater than*

$$1 - \frac{2d}{\epsilon^2 |\mathbb{F}|},$$

over verifier randomness $(r_d, .., r_1) \leftarrow \mathbb{F}$ *in the* `commit` *phase, and letting* $\{\pi_i \in \mathbb{F}^{n_i} : i \in [d+1]\}$ *be the corresponding oracles sent by the prover, either the verifier accepts with probability less than*

$$\beta_\epsilon^l,$$

or $\exists P \in \mathbb{F}[X_1, .., X_d]$ *such that* $Enc_{C_0}(P(r_d, .., r_1)) = \pi_0$ *and* $\Delta(Enc_{C_d}(P), \pi_d) < (1 - (\beta_\epsilon - d\epsilon))n_d$.

High Level Overview of Proof In the remainder of the paper, we assume all linear codes are punctured Reed-Muller codes, that are evaluations of multilinear polynomials over some subset of $\mathbb{F}^d$. Recall from Lemma 2, that if $\pi_L[S] + r\pi_R[S] = c[S]$ for some $c \in C$, then $\pi_L[S], \pi_R[S]$ differ from $c_L[S], c_R[S]$ in very few locations, where $c_L + rc_R = c$. We will show in Lemma 5, that in this case, c is the encoding of polynomial P, c_L is the encoding of polynomial P_L and c_R is the encoding of polynomial P_R where $P_L + rP_R = P$. Next, in Lemma 6, we prove soundness for just one round of the IOPP, and finally, we show how to extend this to the full, multi-round IOPP.

Lemma 5. *Let* $n, k \in \mathbb{N}$ *and let* C *be an* $[n,k]$ *linear error-correcting code. Let* $\beta, \tau_1, \tau_2 \in [0,1]$ *such that* $\beta - (\tau_1 + \tau_2) > 1 - \Delta_C$. *Let* $\pi_L, \pi_R \in \mathbb{F}^n$ *and* $d = \log_2(n)$. *Suppose that* $S \subset [n]$ *where* $|S| > \beta n$ *and there exists* $P_L, P_R \in \mathbb{F}[X_1, .., X_d]$ *such that*

$$|\{i \in S : \pi_L[i] = \mathsf{Enc}_C(P_L)[i] \wedge \pi_R[i] = \mathsf{Enc}_C(P_R)[i]\}| > (\beta - \tau_1)n.$$

Suppose further that there exists $c \in C$ *such that*

$$|\{i \in S : \pi_L[S] + r\pi_R[S] = c[S]\}| > (\beta - \tau_2)n. \tag{13}$$

Then

$$c = \mathsf{Enc}_C(P_L + rP_R). \tag{14}$$

Protocol 3 IOPP.commit

Input oracle: $\pi_d \in \mathbb{F}^{n_d}$
Output oracles: $(\pi_{d-1}, \ldots, \pi_0) \in \mathbb{F}^{n_{d-1}} \times \cdots \times \mathbb{F}^{n_0}$

- For i from $d-1$ downto 0:
  1. The verifier samples and sends $\alpha_i \leftarrow\!\!\$\, \mathbb{F}$ to the prover
  2. For each index $j \in \mathsf{even}[0, n_{i+1} - 1]$, the prover
     (a) sets $f(X) := \mathsf{interp}((\mathbf{diag}(T_i)[j], \pi_{i+1}[j]), (\mathbf{diag}(-T_i)[j], \pi_{i+1}[j+1]))$
     (b) sets $\pi_i[j] = f(\alpha_i)$
  3. The prover outputs oracle $\pi_i \in \mathbb{F}^{n_i}$.

Protocol 4 IOPP.query

Oracles: $(\pi_d, .., \pi_0)$
Repetition Parameter: $\lambda \in \mathbb{N}$

- For $j \in [0, \lambda - 1]$
  - The verifier samples an index $\mu_j \leftarrow\!\!\$\, \mathsf{even}[1, n_d - 1]$
  - For i from $d-1$ downto 0, the verifier
    1. queries oracle entries $\pi_{i+1}[\mu_j]$, $\pi_{i+1}[\mu_j + 1]$
    2. computes $p(X) := \mathsf{interpolate}((\mathbf{diag}(T_i)[\mu_j], \pi_{i+1}[\mu_j]), (\mathbf{diag}(-T_i)[\mu_j], \pi_{i+1}[\mu_j + 1]))$
    3. checks that $p(\alpha_i) = \pi_i[\mu_j/2]$
    4. if $i > 0$ and $\mu_j/2 \mod 2 = 0$, update $\mu_j \leftarrow \mu_j/2$, otherwise update $\mu_j \leftarrow \mu_j/2 - 1$.
  - If π_0 is a valid codeword w.r.t. generator matrix C_0, output accept, otherwise output reject

Figure 2: The IOPP protocol for foldable codes.

Proof. Let $c_L := \mathsf{Enc}_C(P_L), c_R = \mathsf{Enc}_C(P_R)$, $\pi^* := \pi_L + r\pi_R$, $P^*(X_1, .., X_d) = P_L + rP_R$ and $c^* = \mathsf{Enc}_C(P^*)$. Then,

$$
\begin{aligned}
&|\{i \in S : \pi^*[i] = c^*[i]\}| \\
&= |\{i \in S : \pi_L[i] + r\pi_R[i] = c_L[i] + rc_R[i]\}| \text{ (By Definition of } \pi^* \text{ and linearity of } C) \\
&\geq |\{i \in S : \pi_L[i] = c_L[i] \wedge \pi_R[i] = c_R[i]\}| \\
&\geq (\beta - \tau_1)n.
\end{aligned}
$$

Let $c \in C$ be the codeword satisfying Equation 13. Then by a simple counting argument,

$$|\{i \in S : c^*[i] = c[i]\}| > (\beta - (\tau_1 + \tau_2))n.$$

By assumption of the lemma, $\beta - (\tau_1 + \tau_2) > 1 - \Delta_C$. Therefore, by minimum distance properties of C, $c^* = c$, which completes the proof. □

Next, we combine Lemma 5 with Lemma 2 to prove soundness over one single round of the IOPP.

Lemma 6 (One Round Soundness)**.** *Let $d \in \mathbb{N}$ and let C_d, C_{d-1} be a pair of codes from an $[n_d, k_d]$ foldable code family (Definition 3). Let $\pi \in \mathbb{F}^{n_d}, \epsilon \in [0,1]$ and let $\beta_\epsilon \in [0,1]$ be defined as in Throrem 2. Suppose that $\beta_\epsilon - 2d\epsilon > 1 - \Delta_{C_d}$, and that $\exists S \subset [n_d]$ and $c \in C_{d-1}$ such that*

$$|\{i \in S : \mathsf{fold}(\pi, r)[i] = c[i]\}| > (\beta_\epsilon - (d-1)\epsilon)n_{d-1}. \tag{15}$$

Then with probability greater than $1 - \frac{2d}{\epsilon^2|\mathbb{F}|}$ (over verifier randomness r) there exists $P_L, P_R \in \mathbb{F}[X_1, .., X_{d-1}]$ such that $c = \mathsf{Enc}_{C_{d-1}}(P_L + rP_R)$ and

$$|\{i \in S : \pi[i] = \mathsf{Enc}_{C_i}(P_L + XP_R)[i]\}| > (\beta_\epsilon - d\epsilon)n_d. \tag{16}$$

Proof. By Equation 15 and by Definition of fold (Definition 3), it follows that

$$|\{i \in S : \pi_L[i] + r\pi_R[i] = c[i]\}| > (\beta_\epsilon - (d-1)\epsilon)n_{d-1},$$

where (π_L, π_R) is the *unfolding* (Definition 11) of π. Therefore, it follows from Theorem 2 that if $r \notin A_\pi(\epsilon)$ then there exists $c_L, c_R \in C_{d-1}$ such that $c_L + rc_R = c$ and,

$$|\{i \in S : \pi_L[i] = c_L[i] \wedge \pi_R[i] = c_R[i]\}| > (\beta_\epsilon - (d-1)\epsilon - \epsilon)n_{d-1} = (\beta_\epsilon - d\epsilon)n_{d-1}. \tag{17}$$

Let $P_L, P_R \in \mathbb{F}[X_1, .., X_d]$ satisfy $c_L = \mathsf{Enc}_{C_{d-1}}(P_L), c_R = \mathsf{Enc}_{C_{d-1}}(P_R)$. Then by definition of a foldable code (Definition 3), it follows that

$$|\{i \in S : \pi[i] = Enc_{C_d}(P_L + XP_R)[i]\}| \geq 2(\beta_\epsilon - d\epsilon)n_{d-1} = (\beta_\epsilon - d\epsilon)n_d.$$

Furthermore, by Lemma 5,

$$c = \mathsf{Enc}_{C_{d-1}}(P_L + rP_R).$$

By Theorem 2, $|A_\pi(\epsilon)| \leq \frac{2}{\epsilon^2}$, and therefore the probability that $r \notin A_\pi$ is greater than $1 - \frac{2}{\epsilon^2|\mathbb{F}|}$, which completes the proof. □

Next, we show that if the verifier accepts with probability greater than β^l, then this implies the existence of d large sets, one in each oracle, that are consistent with each other with respect to the fold operation.

Lemma 7 (Verifier Queries)**.** *Let $\epsilon \in [0,1]$, let $\beta_\epsilon \in [0,1]$ be defined as in Theorem 2, and let $l \in \mathbb{N}$. If the verifier accepts the query phase with probability greater than β_ϵ^l then there exists d large sets $\{f_i(S) \subset [n_i] : i \in [d], |f_i(S)| > \beta_\epsilon n_i\}$ such that for all $i \in [d]$,*

$$\mathsf{fold}(\pi_{i+1}, \mathbf{r}[i+1])[f_{i+1}(S)] = \pi_i[\mathsf{even}(f_{i+1}(S))/2] \tag{18}$$

where fold is defined in Equation 12.

Proof. Define the function $Q : \mathsf{even}([n_d]) \to \{0,1\}$ as $Q(\mu) = 1$ if the unique verifier query beginning with $\mu \leftarrow \mathsf{even}(n_d)$ (defined in Protocol 4) passes the verifier tests and $Q(\mu) = 0$ otherwise. Let $S = Q^{-1}(1)$. Then, each verifier sample is a Bernoulli trial with success probability $\frac{|S|}{|\mathsf{even}([n_d])|}$. After l queries, the probability of acceptance is $(\frac{|S|}{|\mathsf{even}([n_d])|})^l$. Therefore, if the verifier accepts with probability greater than β^l, then $(\frac{|S|}{\mathsf{even}([n_d])})$ must be larger than β, and so $|S| > \beta|\mathsf{even}([n_d])| = \beta n_{d-1}$. Next, we define $f_i(S)$.

Definition 5 ($f_d(S)$)**.** *Let $d \in \mathbb{N}$ and $S \subset \mathsf{even}([n_d])$. Then,*

$$f_d(S) = S \cup (S + 1).$$

For $i \in [d]$, $f_i(S)$ satisfies the following:

$$\mathsf{even}(f_i(S)) = \{\mathsf{even}(\{j/2, j/2 - 1\}) : j \in \mathsf{even}(f_{i+1}(S)\}$$

$$\mathsf{odd}(f_i(S)) = \mathsf{even}(f_i(S)) + 1.$$

To complete the Lemma, we need to prove that for each $i \in [d + 1]$,

$$\mathsf{fold}(\pi_i, \mathbf{r}[i])[f_i(S)] = \pi_{i-1}[\mathsf{even}(f_i(S))/2].$$

For each $\mu \in S$, let $((\mu_d, \mu_d + 1), .., (\mu_1, \mu_1 + 1))$ be the unique set of queries associated with μ (defined in Protocol 4). Then, by definition of fold (Definition 12), for each $i \in [1, d]$

$$\mathsf{fold}(\pi_i, r_i)[f_i(S)] = \{\pi_{i,L}[\mu_i] + r_i \pi_{i,R}[\mu_i + 1]\} \tag{19}$$

Furthermore,

$$\pi_{i,L}[\mu_i] + r_i \pi_{i,R}[\mu_i + 1] = \pi_{i-1}[\mu_i/2]. \tag{20}$$

Combining Equations 19 and 20 gives

$$\mathsf{fold}(\pi_i, r_i)[f_i(S)] = \{\pi_{i-1}[\mu_i/2] : \mu_i \in \mathsf{even}(f_i(S))\} = \pi_{i-1}[\mathsf{even}(f_i(S)/2)],$$

which completes the proof. □

Finally, we are ready to prove the main statement of the Theorem.

Proof. Suppose by contradiction that the verifier accepts with probability $> \beta^l$ but there does not exist $P \in \mathbb{F}[X_1, .., X_d]$ such that $\mathsf{Enc}_{C_0}(P(r_d, .., r_1)) = \pi_0$ and $\Delta(\mathsf{Enc}_{C_d}(P), \pi_d) < (1 - (\beta_\epsilon - d\epsilon))$. By Lemma 7, there exists a set $S \subset \mathsf{even}(n_d)$ such that for each $i \in [1, d]$,

$$\mathsf{fold}(\pi_i, \mathbf{r}[i])[f_i(S)] = \pi_{i-1}[\mathsf{even}(f_i(S))/2].$$

Therefore, since $\pi_0 \in C_0$ there must exist a round where Equation 15 holds but Equation 16 does not. By Lemma 6, this only happens with probability $\frac{2}{\epsilon^2 |\mathbb{F}|}$. Taking the union bound over d rounds completes the proof. □

Acknowledgements

I would like to thank the following people: Binyi Chen for feedback on earlier drafts and ongoing helpful discussions; Ulrich Haböck for checking correctness of the proof and giving very useful feedback on how to present it; Ron Rothblum and Ben Fisch for general guidance and advice; Michael Rosenberg and Paul Grubbs for edits and feedback; and Giacomo Fenzi for early brainstorming sessions.

Additionally, I would like to thank AR for coming up with the name Khatam; Oded Goldreich for useful advice on how to write a proof; Joan Feigenbaum for gifting me stacks of old conference proceedings, from which I drew inspiration on writing; and KL for gifting me the fountain pen I used to figure out the solution.

References

[1] Scott Ames et al. "Ligero: Lightweight Sublinear Arguments Without a Trusted Setup". In: CCS '17. Dallas, Texas, USA: Association for Computing Machinery, 2017, 2087–2104. DOI: 10.1145/3133956.3134104. URL: https://doi.org/10.1145/3133956.3134104 (cited on page 2).

[2] Gal Arnon et al. *WHIR: Reed–Solomon Proximity Testing with Super-Fast Verification*. Cryptology ePrint Archive, Paper 2024/1586. 2024. URL: https://eprint.iacr.org/2024/1586 (cited on page 2).

[3] Eli Ben-Sasson, Swastik Kopparty, and Shubhangi Saraf. "Worst-case to average case reductions for the distance to a code". In: *Proceedings of the 33rd Computational Complexity Conference*. CCC '18. San Diego, California: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018 (cited on pages 2, 3).

[4] Eli Ben-Sasson et al. "Robust pcps of proximity, shorter pcps and applications to coding". In: *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*. STOC '04. Chicago, IL, USA: Association for Computing Machinery, 2004, 1–10. ISBN: 1581138520. DOI: 10.1145/1007352.1007361. URL: https://doi.org/10.1145/1007352.1007361 (cited on page 1).

[5] Eli Ben-Sasson et al. *Short Interactive Oracle Proofs with Constant Query Complexity, via Composition and Sumcheck*. Cryptology ePrint Archive, Report 2016/324. 2016. URL: https://eprint.iacr.org/2016/324 (cited on page 1).

[6] Eli Ben-Sasson et al. "Fast Reed-Solomon Interactive Oracle Proofs of Proximity". In: *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*. Edited by Ioannis Chatzigiannakis et al. Volume 107. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2018, 14:1–14:17. DOI: 10.4230/LIPIcs.ICALP.2018.14. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ICALP.2018.14 (cited on pages 2, 3).

[7] Eli Ben-Sasson et al. "Aurora: Transparent Succinct Arguments for R1CS". In: *Advances in Cryptology – EUROCRYPT 2019, Part I*. Edited by Yuval Ishai and Vincent Rijmen. Volume 11476. Lecture Notes in Computer Science. Darmstadt, Germany: Springer, Cham, Switzerland, 2019, pages 103–128. DOI: 10.1007/978-3-030-17653-2_4 (cited on page 1).

[8] Eli Ben-sasson et al. "DEEP-FRI: Sampling outside the box improves soundness". In: (Mar. 2019). DOI: 10.48550/arXiv.1903.12243 (cited on pages 2, 3).

[9] Eli Ben-Sasson et al. "Proximity Gaps for Reed–Solomon Codes". In: *J. ACM* 70.5 (Oct. 2023). DOI: 10.1145/3614423. URL: https://doi.org/10.1145/3614423 (cited on pages 2–4).

[10] Martijn Brehm et al. *Blaze: Fast SNARKs from Interleaved RAA Codes*. Cryptology ePrint Archive, Paper 2024/1609. 2024. URL: https://eprint.iacr.org/2024/1609 (cited on page 2).

[11] Binyi Chen et al. *HyperPlonk: Plonk with Linear-Time Prover and High-Degree Custom Gates*. Cryptology ePrint Archive, Report 2022/1355. 2022. URL: https://eprint.iacr.org/2022/1355 (cited on pages 1, 2).

[12] Benjamin E. Diamond and Jim Posen. *Polylogarithmic Proofs for Multilinears over Binary Towers*. Cryptology ePrint Archive, Paper 2024/504. https://eprint.iacr.org/2024/504. 2024. URL: https://eprint.iacr.org/2024/504 (cited on page 2).

[13] Irit Dinur and Omer Reingold. "Assignment testers: Towards a combinatorial proof of the PCP theorem". In: *SIAM Journal on Computing* 36.4 (2006), pages 975–1024 (cited on page 1).

[14] Yiwen Gao, Haibin Kan, and Yuan Li. *Linear Proximity Gap for Reed-Solomon Codes within the 1.5 Johnson Bound*. Cryptology ePrint Archive, Paper 2024/1810. 2024. URL: https://eprint.iacr.org/2024/1810 (cited on pages 2, 3, 11, 12).

[15] Alexander Golovnev et al. *Brakedown: Linear-time and field-agnostic SNARKs for R1CS*. Cryptology ePrint Archive, Paper 2021/1043. 2021. URL: https://eprint.iacr.org/2021/1043 (cited on page 2).

[16] Yanpei Guo et al. *DeepFold: Efficient Multilinear Polynomial Commitment from Reed-Solomon Code and Its Application to Zero-knowledge Proofs*. Cryptology ePrint Archive, Paper 2024/1595. 2024. URL: https://eprint.iacr.org/2024/1595 (cited on page 3).

[17] Ulrich Haböck. *Basefold in the List Decoding Regime*. Cryptology ePrint Archive, Paper 2024/1571. 2024. URL: https://eprint.iacr.org/2024/1571 (cited on pages 2, 3, 11, 12).

[18] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. "Constant-Size Commitments to Polynomials and Their Applications". In: *Advances in Cryptology – ASIACRYPT 2010*. Edited by Masayuki Abe. Volume 6477. Lecture Notes in Computer Science. Singapore: Springer Berlin Heidelberg, Germany, 2010, pages 177–194. DOI: 10.1007/978-3-642-17373-8_11 (cited on page 1).

[19] Tohru Kohrita and Patrick Towa. *Zeromorph: Zero-Knowledge Multilinear-Evaluation Proofs from Homomorphic Univariate Commitments*. Cryptology ePrint Archive, Paper 2023/917. 2023. URL: https://eprint.iacr.org/2023/917 (cited on page 2).

[20] Carsten Lund et al. "Algebraic methods for interactive proof systems". In: *J. ACM* 39.4 (1992), pages 859–868. DOI: 10.1145/146585.146605. URL: https://doi.org/10.1145/146585.146605 (cited on page 2).

[21] Charalampos Papamanthou, Elaine Shi, and Roberto Tamassia. "Signatures of Correct Computation". In: *TCC 2013: 10th Theory of Cryptography Conference*. Edited by Amit Sahai. Volume 7785. Lecture Notes in Computer Science. Tokyo, Japan: Springer Berlin Heidelberg, Germany, 2013, pages 222–242. DOI: 10.1007/978-3-642-36594-2_13 (cited on page 2).

[22] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. "Constant-round interactive proofs for delegating computation". In: *48th Annual ACM Symposium on Theory of Computing*. Edited by Daniel Wichs and Yishay Mansour. Cambridge, MA, USA: ACM Press, 2016, pages 49–62. DOI: 10.1145/2897518.2897652 (cited on page 1).

[23] Hang Su, Qi Yang, and Zhenfei Zhang. *Jolt-b: recursion friendly Jolt with basefold commitment*. Cryptology ePrint Archive, Paper 2024/1131. https://eprint.iacr.org/2024/1131. 2024. URL: https://eprint.iacr.org/2024/1131 (cited on page 2).

[24] Alexander Vlasov and Konstantin Panarin. *Transparent Polynomial Commitment Scheme with Polylogarithmic Communication Complexity*. Cryptology ePrint Archive, Report 2019/1020. 2019. URL: https://eprint.iacr.org/2019/1020 (cited on page 1).

[25] Hadas Zeilberger, Binyi Chen, and Ben Fisch. *BaseFold: Efficient Field-Agnostic Polynomial Commitment Schemes from Foldable Codes*. Cryptology ePrint Archive, Paper 2023/1705. 2023. URL: https://eprint.iacr.org/2023/1705 (cited on pages 1, 2, 11).

Appendix

A Proof of Lemma 1

Proof. Let $x \in \{2,3\}, \beta, \epsilon \in \{0,1\}$ such that $\epsilon < \beta/100, \epsilon \leq 1+\beta^2-\beta$, $(\beta^2-(\beta^3-\epsilon)) < 1/3$ and $\beta > 0.0930$ (it is easy to check that these constraints are practical). Suppose further that $\mathcal{S}(x,\beta) \subset 2^{[n]}$ defined as in the lemma statement. Recall that our goal is bound the size of $\mathcal{S}(x,\beta)$. We do this by constructing a bipartite graph, G, where right vertices are labeled by elements in $[n]$ and left vertices are labeled by subsets of $[n]$. We place an edge between a vertex $v \in L$ and a vertex $w \in R$ if the element associated with w is contained inside the set associated with v. By definition of $\mathcal{S}(x,\beta)$, this implies that that any vertex in L has more than βn neighbors but shares less than $(\beta^x - \epsilon)n$ neighbors with any other vertex in L. In other words, letting $N(v)$ denote the neighbor set of $v \in L$:

- $\forall v \in L, |N(v)| > \beta n$
- $\forall v_1, .., v_x \in L, |\bigcap_{i \in x} N(v)| < (\beta^x - \epsilon)n$

We prove now that in a bipartite graph with these two properties,

$$|L| \leq \frac{1}{\epsilon}.$$

This proof is a generalization of the proof of the Johnson bound [11]. For $v_1, .., v_x \in L$ and $w \in R$, define an "x-angle" as an (x+1)-tuple, $(v_1, .., v_x, w)$, such that $(v_i, w) \in E$ for all $i \in [x]$. Let $d(w)$ be the degree of node w. Then the number of x-angles in the graph is equal to

$$\sum_{v \in R} \binom{d(v)}{x}$$

[11] https://www.cs.cmu.edu/~venkatg/teaching/au18-coding-theory/lec-scribes/list-decoding.pdf

By Jensen's Inequality, we know that for each $x \in \{2, 3\}$,

$$\sum_{v \in R} \binom{d(v)}{x} \geq |R| \binom{(\sum_{v \in R} d(v))/|R|}{x}. \tag{21}$$

Since each vertex in L has more than βn neighbors, it follows that the sum of the degrees of all vertices in R is greater than $|L| \cdot \beta |R|$, i.e. $\sum_{v \in R} d(v) \geq |L| \cdot \beta |R|$. Thus, by Equation 21, it follows that for $x \in \{2, 3\}$,

$$\sum_{v \in R} \binom{d(v)}{x} \geq n \binom{\beta |L|}{x}.$$

On the other hand, for each $x \in \{2, 3\}$, any x vertices in L can share at most $(\beta^x - \epsilon)|R| = (\beta^x - \epsilon)n$ neighbors in R. Thus the total number of angles in G is at most

$$\binom{|L|}{x} (\beta^x - \epsilon) n$$

Combining the two inequalities, we have for each $x \in \{2, 3\}$,

$$\binom{\beta |L|}{x} \leq \binom{|L|}{x} (\beta^x - \epsilon).$$

Solving for L when $x = 2$, we have,

$$\frac{\beta |L| (\beta |L| - 1)}{2} \leq \frac{(\beta^2 - \epsilon)|L|(|L| - 1)}{2} \tag{22}$$

$$\beta |L| (\beta |L| - 1) \leq (\beta^2 - \epsilon)|L|(|L| - 1) \tag{23}$$

$$\beta (\beta |L| - 1) \leq (\beta^2 - \epsilon)(|L| - 1) \tag{24}$$

$$\beta^2 |L| - \beta \leq (\beta^2 - \epsilon)|L| - (\beta^2 - \epsilon) \tag{25}$$

$$\beta^2 |L| - \beta^2 |L| + \epsilon |L| \leq \beta - (\beta^2 - \epsilon) \tag{26}$$

$$|L| \epsilon \leq \beta - (\beta^2 - \epsilon) \tag{27}$$

$$\tag{28}$$

Therefore, since $\epsilon \leq 1 + \beta^2 - \beta$,

$$|L| \leq \frac{\beta - (\beta^2 - \epsilon)}{\epsilon} \leq \frac{1}{\epsilon}.$$

Next, we solve for L when $x = 3$. Let $\alpha := \beta^3 - \epsilon$. Then,

$$\frac{\beta |L| (\beta |L| - 1)(\beta |L| - 2)}{3} \leq \frac{\alpha |L| (|L| - 1)(|L| - 2)}{3} \tag{29}$$

$$\beta |L| (\beta |L| - 1)(\beta |L| - 2) \leq \alpha |L| (|L| - 1)(|L| - 2) \tag{30}$$

$$\beta (\beta |L| - 1)(\beta |L| - 2) \leq \alpha (|L| - 1)(|L| - 2) \tag{31}$$

$$(\beta^2 |L| - \beta)(\beta |L| - 2) \leq (\alpha |L| - \alpha)(|L| - 2) \tag{32}$$

$$(\beta^3 |L|^2 - 2\beta^2 |L| - \beta^2 |L| + 2\beta) \leq (\alpha |L|^2 - 2\alpha |L| - \alpha |L| + 2\alpha \tag{33}$$

$$(\beta^3 |L|^2 - 3\beta^2 |L| + 2\beta) \leq (\alpha |L|^2 - 3\alpha |L| + 2\alpha \tag{34}$$

$$(\beta^3 - \alpha)|L|^2 - 3|L|(\beta^2 - \alpha) + 2(\beta - \alpha) \leq 0 \tag{35}$$

Let $P(X) = (\beta^3 - \alpha)X^2 - 3(\beta^2 - \alpha)X + 2(\beta - \alpha)$. Our goal is to solve for L, such that $P(L) \leq 0$. Since P is concave, it is smaller than 0 for X between $[x_0, x_1]$, where x_0, x_1 are its two roots and $x_0 << x_1$. Next, we solve for x_1, which is an upper bound on L. By the quadratic equation,

$$x_1 = \frac{3(\beta^2 - \alpha)) + \sqrt{(-3(\beta^2 - \alpha))^2 - 8\epsilon(\beta - \alpha)}}{2\epsilon}$$

For now, lets assume that x_1 is a real root, i.e. that the discriminate is greater than 0. In that case, it follows that,

$$\begin{aligned} x &\leq \frac{3(\beta^2 - \alpha)) + \sqrt{(-3(\beta^2 - \alpha))^2}}{2\epsilon} \\ &= \frac{3(\beta^2 - \alpha)) + 3(\beta^2 - \alpha)}{2\epsilon} \\ &= \frac{2 \cdot (3(\beta^2 - \alpha))}{2\epsilon} = \frac{3(\beta^2 - \alpha)}{\epsilon} \\ &= \frac{3(\beta^2 - (\beta^3 - \epsilon))}{\epsilon} \end{aligned}$$

Since we assume that $(\beta^2 - (\beta^3 - \epsilon)) < 1/3$, it follows that

$$|L| \leq x_1 \leq \frac{1}{\epsilon}.$$

Next, we show that x_1 is a real root. To do so, we state the following claim. We do not provide a proof as it can be easily verified.

Claim 5. *Let $\beta, \epsilon \in [0, 1]$ such that $\epsilon < \frac{\beta}{100}$. Then, for all $\beta > 0.0930$,*

$$(-3(\beta^2 - \alpha))^2 - 8\epsilon(\beta - \alpha) \geq 0.$$

□