CONJURA · OPEN PROBLEM

Proximity Gaps for Reed–Solomon Codes Up to Capacity

From the Johnson bound $1-\sqrt{\rho}$ *to capacity* $1-\rho$*, with polynomial error*

Status. Statement: AI-written from Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty and Shubhangi Saraf, *Proximity Gaps for Reed–Solomon Codes*, IACR ePrint 2020/654. Not yet checked by a human. Numbered Conjecture 8.4 of the source. Proved up to the Johnson/Guruswami–Sudan bound $\delta < 1-\sqrt{\rho}$; conjectured up to capacity $\delta \le 1-\rho-\eta$. The source records that nothing known contradicts $c_1 = c_2 = 2$, and that for fields of characteristic greater than the degree nothing known contradicts $c_1 = c_2 = 1$ — while those smaller exponents are ruled out in characteristic two.

Category. *Proof Systems.*

Conjecture (informal). *Take a Reed–Solomon code and any affine space of vectors. The source's theorem says the space cannot be split: either all of its elements are close to the code, or almost none of them are. That dichotomy is what makes the FRI protocol sound, and hence what underlies the concrete security of the STARK proof systems built on it. The theorem is proved for closeness thresholds below the Johnson bound* $1-\sqrt{\rho}$*, because the proof runs through list-decoding algorithms that stop there. The conjecture is that the same dichotomy, with a polynomially small error, holds all the way up to capacity* $1-\rho$*. Settling it would close the factor-two gap between the provable and conjectured number of FRI query repetitions needed for a target security level.*

1 The setting

Write $\mathrm{RS}[\mathbb{F}_q, D, k]$ for the Reed–Solomon code of blocklength $n = |D|$ whose codewords are the polynomials of degree at most k evaluated on $D \subseteq \mathbb{F}_q$, and $\rho = (k+1)/n$ for its rate. A collection C of subsets of $\mathbb{F}_q^D$ displays a (δ,ε)-*proximity gap* with respect to a code V if every $A \in C$ satisfies one of two alternatives: either all

of its elements are within relative Hamming distance δ of V, or at most an ε fraction are.

The source proves this for C_{Affine}, the collection of affine spaces (its Theorem 1.2), in two regimes. Below the unique decoding radius, $\delta \in (0, (1-\rho)/2]$, with the essentially tight error $\varepsilon_U = n/q$. Above it, up to the Johnson/Guruswami–Sudan list-decoding bound $\delta < 1-\sqrt{\rho}$, with the larger error $\varepsilon_J = O\big((\eta\rho)^{-O(1)} \cdot n^2/q\big)$ where $\eta = 1-\sqrt{\rho}-\delta$. It also proves the *correlated agreement* form the applications actually use — its Theorem 1.4 for lines, Theorem 1.5 for low-degree parameterized curves and Theorem 1.6 for affine spaces: if a random element of the space is δ-close to V with probability exceeding ε, then there is a single set $D' \subseteq D$ of density at least $1-\delta$ on which *every* generator agrees with a codeword.

Two things bound what is proved. The proximity parameter stops at the Johnson bound because the proof leans on list-decoding algorithms that reach exactly that far; the source is explicit that this is an artefact of the technique, remarking that it conjectures Theorem 1.2 "holds even for larger proximity parameters, up to capacity $(1-\rho)$". And the error bound above the unique decoding radius is nontrivial only for fields quadratically larger than the blocklength.

The stake is concrete. Ben-Sasson et al. showed that FRI needs at least $s \geq \lambda/\log(1/\rho)$ invocations of its query phase to reach security parameter λ, and conjectured that this lower bound is also sufficient for large enough fields; the source's results give $s \approx 2\lambda/\log(1/\rho)$ for quadratically large fields. As the source puts it, "Closing the gap between the provable upper and lower bounds on s is left as an interesting open problem", and it identifies Conjecture 1 below as the improvement to its main correlated-agreement theorem that would imply that conjecture.

2 Notation and parameters

- $\mathbb{F}_q$ is the field of size q; $D \subseteq \mathbb{F}_q$ the evaluation domain, $n = |D|$ the blocklength.
- $V = \mathrm{RS}[\mathbb{F}_q, D, k]$ is the RS code of dimension $k+1$ and rate $\rho = (k+1)/n$.
- $\Delta(u, V)$ is the relative Hamming distance from $u \in \mathbb{F}_q^D$ to the code; δ is the proximity parameter and ε the error parameter.

- $\eta > 0$ is the slack below the relevant bound: the source writes $\eta = 1 - \sqrt{\rho} - \delta$ in Theorem 1.2 and $\delta \le 1 - \rho - \eta$ in the conjecture.
- ℓ is the degree of the parameterized curves in Theorem 1.5.
- c_1, c_2 are the two constants the conjecture asserts exist.

3 The conjecture

Definition 1 (Proximity gap, [BCIKS20, Section 1]). Let $V \subseteq \mathbb{F}_q^D$ be a code and C a collection of subsets of $\mathbb{F}_q^D$. Then C displays a (δ,ε)-*proximity gap* with respect to V if for every $A \in C$, either

$$\Pr_{u \in A}\left[\Delta(u,V) \le \delta\right] = 1 \qquad \text{or} \qquad \Pr_{u \in A}\left[\Delta(u,V) \le \delta\right] \le \varepsilon.$$

C_{Affine} denotes the collection of affine spaces in $\mathbb{F}_q^D$.

Theorem 1 (What is proved, [BCIKS20, Theorems 1.2 and 1.4]). *C_{Affine} displays a (δ,ε)-proximity gap with respect to $V = \mathrm{RS}[\mathbb{F}_q, D, k]$ for every $\delta \in (0, 1 - \sqrt{\rho})$, with $\varepsilon = n/q$ when $\delta \le (1-\rho)/2$ and $\varepsilon = O\big((\eta\rho)^{-O(1)} n^2/q\big)$, $\eta = 1 - \sqrt{\rho} - \delta$, above that. Moreover (correlated agreement over lines) for $u_0, u_1 \in \mathbb{F}_q^D$ and δ in the same range, if $\Pr_{z \in \mathbb{F}_q}[\Delta(u_0 + z u_1, V) \le \delta] > \varepsilon$ then there are $D' \subseteq D$ with $|D'|/|D| \ge 1 - \delta$ and $v_0, v_1 \in V$ such that v_0 agrees with u_0 and v_1 with u_1 everywhere on D'.*

Conjecture 1 ([BCIKS20, Conjecture 8.4]). *There exist constants c_1, c_2 such that the following hold for all $\eta > 0$.*

- *Theorems 1.2, 1.4 and 1.6 of the source hold for proximity parameter $\delta \le 1 - \rho - \eta$ with error*

$$\varepsilon \le \frac{1}{(\eta\rho)^{c_1}} \cdot \frac{n^{c_2}}{q}.$$

- *Theorem 1.5 of the source holds for proximity parameter $\delta \le 1 - \rho - \eta$ and parameterized curves of degree ℓ with error*

$$\varepsilon \le \frac{1}{(\eta\rho)^{c_1}} \cdot \frac{(\ell \cdot n)^{c_2}}{q}.$$

Remark 1 (Why capacity is the right target, and why the proof stops short). The distance $1-\sqrt{\rho}$ is the Johnson/Guruswami–Sudan bound, the largest for which efficient list decoding is known, and the source's proof above the unique decoding radius uses the Guruswami–Sudan algorithm in a non-algorithmic but essential way. Capacity $1-\rho$ is the information-theoretic limit beyond which the list of nearby codewords need not be small at all. So the conjecture asks for the dichotomy to survive in the range where the only tool the proof has is gone — which is why the source calls the current ceiling an artefact of the technique rather than of the statement.

Remark 2 (The constants, and what is known not to hold). The source records two calibrations, both worth keeping in view. To its knowledge nothing contradicts $c_1 = c_2 = 2$. Restricted to fields of characteristic greater than k, nothing known contradicts $c_1 = c_2 = 1$ — but those smaller exponents provably cannot hold in characteristic two, by [BGKS20, Appendix B]. So a proof of Conjecture 1 should report the exponents it achieves and the characteristic it needs, and a refutation is most likely to come from characteristic two.

Remark 3 (What settling it buys). By the source's own accounting, Conjecture 1 implies the conjecture of [BBHR18b] that $s \geq \lambda/\log(1/\rho)$ query repetitions suffice for FRI to reach security parameter λ over sufficiently large fields, against the $s \approx 2\lambda/\log(1/\rho)$ the source can prove. That is a factor of two in the query complexity of every FRI-based proof system, which is why the statement is of practical and not only combinatorial interest.

Remark 4 (Field size is a separate axis). Conjecture 1 fixes the error's shape as $(\eta\rho)^{-c_1} n^{c_2}/q$ and so still needs q polynomially larger than n to be nontrivial. The source separately notes that its error above the unique decoding radius is nontrivial only for q quadratically larger than n, while below that radius linear suffices. Reducing the required field size at a fixed proximity parameter is a different question and is not what this statement asks.

4 Bibliography

[BCIKS20] Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty and Shubhangi Saraf. *Proximity gaps for Reed–Solomon codes.* IACR ePrint 2020/654. The source. The proximity-gap definition and Theo-

rem 1.2 are on page 2, Theorem 1.4 on page 4, Theorems 1.5 and 1.6 on page 5, and Conjecture 8.4 with its discussion on page 43 (PDF page 43).

[BBHR18b] Eli Ben-Sasson, Iddo Bentov, Ynon Horesh and Michael Riabzev. *Fast Reed–Solomon interactive oracle proofs of proximity.* 45th International Colloquium on Automata, Languages, and Programming (ICALP), 2018. Introduces FRI, proves the $s \geq \lambda/\log(1/\rho)$ lower bound on the number of query repetitions, and conjectures it is also sufficient for large enough fields.

[BGKS20] Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty and Shubhangi Saraf. *DEEP-FRI: sampling outside the box improves soundness.* 11th Innovations in Theoretical Computer Science Conference (ITCS), 2020, LIPIcs 151, pages 5:1–5:32. The prior state of the art the source improves on, and the work whose Appendix B rules out $c_1 = c_2 = 1$ in characteristic two.

[BKS18] Eli Ben-Sasson, Swastik Kopparty and Shubhangi Saraf. *Worst-case to average case reductions for the distance to a code.* 33rd Computational Complexity Conference (CCC), 2018, pages 24:1–24:23. Part of the prior state of the art on proximity testing for RS codes that the source improves.