



Fast Reed-Solomon Interactive Oracle Proofs of Proximity

Eli Ben-Sasson* Iddo Bentov† Yinon Horesh* Michael Riabzev*

January 12, 2018

Abstract

The family of Reed-Solomon (RS) codes plays a prominent role in the construction of quasi-linear probabilistically checkable proofs (PCPs) and interactive oracle proofs (IOPs) with perfect zero knowledge and polylogarithmic verifiers. The large concrete computational complexity required to prove membership in RS codes is one of the biggest obstacles to deploying such PCP/IOP systems in practice.

To advance on this problem we present a new interactive oracle proof of proximity (IOPP) for RS codes; we call it the *Fast RS IOPP* (FRI) because (i) it resembles the ubiquitous Fast Fourier Transform (FFT) and (ii) the arithmetic complexity of its prover is strictly linear and that of the verifier is strictly logarithmic (in comparison, FFT arithmetic complexity is quasi-linear but not strictly linear). Prior RS IOPPs and PCPs of proximity (PCPPs) required super-linear proving time even for polynomially large query complexity.

For codes of block-length N , the arithmetic complexity of the (interactive) FRI prover is less than $6 \cdot N$, while the (interactive) FRI verifier has arithmetic complexity $\leq 21 \cdot \log N$, query complexity $2 \cdot \log N$ and constant soundness — words that are δ -far from the code are rejected with probability $\min\{\delta \cdot (1 - o(1)), \delta_0\}$ where δ_0 is a positive constant that depends mainly on the code rate. The particular combination of query complexity and soundness obtained by FRI is better than that of the quasilinear PCPP of [Ben-Sasson and Sudan, SICOMP 2008], even with the tighter soundness analysis of [Ben-Sasson et al., STOC 2013; ECCC 2016]; consequently, FRI is likely to facilitate better concretely efficient zero knowledge proof and argument systems.

Previous concretely efficient PCPPs and IOPPs suffered a constant *multiplicative* factor loss in soundness with each round of “proof composition” and thus used at most $O(\log \log N)$ rounds. We show that when δ is smaller than the unique decoding radius of the code, FRI suffers only a negligible *additive* loss in soundness. This observation allows us to increase the number of “proof composition” rounds to $\Theta(\log N)$ and thereby reduce prover and verifier running time for fixed soundness.

*Technion — Israel Institute of Technology, Haifa, Israel; supported by the Israel Science Foundation (grant # 1501/14) and the US–Israel Binational Science Foundation

†Cornell University, Ithaca, NY, USA

1 Introduction

The family of Reed-Solomon (RS) codes is a fundamental object of study in algebraic coding theory and theoretical computer science [RS60]. For an evaluation set S of N elements in a finite field \mathbb{F} and a rate parameter $\rho \in (0, 1]$, the code $\text{RS}[\mathbb{F}, S, \rho]$ is the space of functions $f : S \rightarrow \mathbb{F}$ that are evaluations of polynomials of degree $d < \rho N$ [RS60]. The *RS proximity problem* assumes a verifier has oracle access to $f : S \rightarrow \mathbb{F}$, and asks that verifier to distinguish, with “large” confidence and “small” query complexity, between the case that f is a codeword of $\text{RS}[\mathbb{F}, S, \rho]$ and the case that f is δ -far in relative Hamming distance from all codewords. This problem has been addressed in several different computational models (surveyed next and summarized in Table 1), and is also the focus of this paper.

RS proximity testing: When no additional data is provided to the verifier, the RS proximity problem is commonly called a *testing* problem, and has been first defined and addressed by Rubinfeld and Sudan in [RS92] (cf. [FS95]). In this case, one can see that $d + 1$ queries are necessary and sufficient to solve the problem: codewords are accepted by their tester with probability 1 whereas functions that are δ -far from the code are rejected with probability $\geq \delta$. Since no additional information is provided to the verifier in this model, we may say that a prover attempting to convince the verifier that $f \in \text{RS}[\mathbb{F}, S, \rho]$ spends zero computational effort, zero rounds of interaction and produces a proof of length zero.

RS proximity verification — PCPP model: Probabilistically checkable proofs of proximity (PCPP) [BGH⁺06, DR04] relax the testing problem to a setting in which the verifier is given oracle access also to an auxiliary proof, called a PCPP and denoted π . This PCPP is produced by the prover, which is given $f \in \text{RS}[\mathbb{F}, S, \rho]$ as input. The time required to produce π is the *prover complexity* and $|\pi|$ is called the *proof length*¹; similarly, *verifier complexity* is the total time required to generate queries and check query-answers. The techniques used to prove the celebrated PCP Theorem [ALM⁺98, AS98] also show that the proximity problem can be solved with constant query complexity and proof length and prover complexity $N^{O(1)}$, or with proof length $N^{1+\epsilon}$ and query complexity $(\log N)^{O(1/\epsilon)}$ [BFLS91]. The current state of the art in the PCPP model gives proofs of length $\tilde{O}(N) \triangleq N \cdot \log^{O(1)} N$ with constant query complexity [BS08, Din07] and prover complexity $\tilde{O}(N)$ [BCGT13]; verifier complexity is $\text{poly log } N$ [BGH⁺05, Mie09].

RS proximity verification — IOPP model: Interactive oracle proofs of proximity (IOPP), formally introduced in [BCG⁺16] and, independently, in [RRR16] (under the name “probabilistically checkable interactive proofs of proximity”), generalize IPs, PCPs and interactive PCPs (IPCP) [KR08]. As in an IP and IPCP, several rounds of interaction are used in which the prover sends messages $\pi_1, \pi_2, \dots, \pi_r$ in response to successive verifier messages. As in a PCP and IPCP, the verifier is not required to read prover messages in entirety but rather may query them at random locations (in an IPCP, verifier must read the full messages π_2, \dots but may query π_1 randomly); the query complexity is the total number of entries read from f and $\pi_1, \pi_2, \dots, \pi_r$. The prover is provided with $f \in \text{RS}[\mathbb{F}, S, \rho]$ as input and *prover complexity* is the total time required to produce all (prover) messages², while *proof length* is generalized from the PCPP setting to the IOPP setting and defined as $|\pi_1| + \dots + |\pi_r|$. IOPPs can be used to “replace” PCPP proof composition with more rounds of interaction, and thereby reduce proof length and prover complexity without compromising soundness (see Section 1.3). In particular, the IOPP version of the aforementioned

¹Typically π is a sequence of elements in \mathbb{F} . Therefore, proof length is measured over the alphabet \mathbb{F} .

²Notice that prover complexity does not include the time needed to produce f .

PCPP constructions reduces proof length to $O(N)$ with no change to soundness and/or query complexity [BBGR16a, BCG⁺16]. In spite of the shorter proof length, prover complexity in prior works was $\Theta(N \text{poly log } N)$ due to a limitation on the number of proof-composition rounds, explained in Section 2.1.

	prover comp.	proof length	verifier comp.	query comp.	round comp.
1. Testing [RS92]	0	0	$\tilde{O}(\rho N)$	ρN	0
2. PCP [ALM ⁺ 98, AS98]	$N^{O(1)}$	$N^{O(1)}$	$N^{O(1)}$	$O(\frac{1}{\delta})$	1
3. PCP [BFL90, BFLS91]	$N^{1+\epsilon}$	$N^{1+\epsilon}$	$\frac{1}{\delta} \log^{O(1/\epsilon)} N$	$\frac{1}{\delta} \log^{O(1/\epsilon)} N$	1
4. PCPP [BS08, BGH ⁺ 06, BCGT13]	$\geq N \log^{1.5} N$	$\geq N \log^{1.5} N$	$\geq \frac{1}{\delta} \log^{5.8} N$	$\frac{1}{\delta} \log^{5.8} N$	1
5. PCPP [Din07, Mie09]	$N \log^c N$	$N \log^c N$	$\frac{1}{\delta} \log^c N$	$O(\frac{1}{\delta})$	1
6. IOPP [BCF ⁺ 16, BBGR16b]	$N \log^c N$	$> 4 \cdot N$	$\frac{1}{\delta} \log^c N$	$O(\frac{1}{\delta})$	$\log \log N$
7. This work	$< 6 \cdot N$	$< \frac{N}{3}$	$\leq 21 \cdot \log N$	$2 \log N$	$\frac{\log N}{2}$

Table 1: Comparison of RS proximity protocols. For concreteness, all results stated for binary additive RS codes with rate $\rho = 1/8$ evaluated over a sufficiently large set $S, |S| = N$ satisfying $N/|\mathbb{F}| < 0.001$ with proximity parameter $\delta < \delta_0$ (cf. Theorem 1.3) and soundness at least 0.99δ ; i.e., the rejection probability of δ -far words is at least 0.99δ for $\delta < \delta_0$ (in particular, smaller δ leads to smaller soundness). Exponents for the 4th row taken from [BCGT13]; the various exponents c in the 5th and 6th row have not been estimated in prior works but are greater than the respective exponents in the 4th row.

1.1 Main results

We present a new IOPP for RS codes, called the Fast RS IOPP (FRI) because of its resemblance to the Fast Fourier Transform (FFT) [CT65]; its analysis relies on the quasi-linear RS-PCPP [BS08] (see Section 2.1). FRI is the first RS-IOPP to have (i) *strictly linear* arithmetic complexity for the prover with (ii) *strictly logarithmic* arithmetic complexity for the verifier and (iii) *constant* soundness. We start by recalling IOPP systems as described in [BCF⁺16, Section 3.2], after informally summarizing the main complexity parameters of IOPs (introduced and discussed thoroughly in [BCS16]).

IOP An *Interactive Oracle Proof (IOP)* system S is defined by a pair of interactive randomized algorithms $S = (P, V)$, where P denotes the prover and V the verifier. On input x of length N , the number of rounds of interaction is denoted by $r(N)$ and called the *round complexity* of the system. During a single round the prover sends a message to which the verifier is given oracle access, and the verifier responds with a message to the prover. The *proof length*, denoted $\ell(N)$, is the sum of lengths of all messages sent by the prover. The *query complexity* of the protocol, denoted $q(N)$, is the number of entries read by V from the various prover messages; since the verifier has oracle access to those messages, typically $q(N) \ll \ell(N)$. (For the FRI system $q(N) = O(\log \ell(N))$). We denote by $\langle P \leftrightarrow V \rangle(x)$ the output of V after interacting with P on input x ; this output is either accept or reject. An IOP is said to be *transparent* (or have *public randomness*) if all messages sent from the verifier are public random coins and all queries are determined by public coins, which are broadcast to the prover (such protocols are also known as Arthur-Merlin protocols [Bab85]).

IOPP As its name suggests, an *IOP of proximity (IOPP)* is the natural generalization of a PCP of Proximity (PCPP) to the IOP model. An IOPP for a family of codes³ \mathcal{C} is a pair (P, V) of randomized algorithms, called *prover* and *verifier*, respectively. Both parties receive as common input a specification of a code $C \in \mathcal{C}$ which we view as a set of functions $C = \{f : S \rightarrow \Sigma\}$ for a finite set S and alphabet Σ . We also assume that the verifier has oracle access to a function $f^{(0)} : S \rightarrow \Sigma$ and that the prover receives the same function as explicit input. The number of rounds of interaction, or *round complexity*, is denoted by r , *query complexity* is denoted by q .

Definition 1.1 (Interactive Oracle Proof of Proximity (IOPP)). *An r -round Interactive Oracle Proof of Proximity (IOPP) $S = (P, V)$ is a $(r + 1)$ -round IOP. We say S is an $(r$ -round) IOPP for the error correcting code $C = \{f : S \rightarrow \Sigma\}$ with soundness $\mathbf{s}^- : (0, 1] \rightarrow [0, 1]$ with respect to distance measure Δ , if the following conditions hold:*

- **First message format:** *the first prover message, denoted $f^{(0)}$, is a purported codeword of C , i.e., $f^{(0)} : S \rightarrow \Sigma$*
- **Completeness:** $\Pr [\langle P \leftrightarrow V \rangle = \text{accept} | \Delta(f^{(0)}, C) = 0] = 1$
- **Soundness:** *For any P^* , $\Pr [\langle P^* \leftrightarrow V \rangle = \text{reject} | \Delta(f^{(0)}, C) = \delta] \geq \mathbf{s}^-(\delta)$*

The sum of lengths of all prover messages, except for $f^{(0)}$, is the IOPP proof length; the time required to generate all messages except for $f^{(0)}$ is the prover complexity. The IOPP query complexity is the total number of queries to all messages, including $f^{(0)}$ and the decision complexity is the computational complexity (see following remark) required by the verifier to reach its verdict, once the queries and query answers are provided as inputs.

Remark 1.2 (Computational model for decision complexity). *The computational model in which decision complexity is computed is left undefined. A natural default is to use boolean circuit complexity. However, later we study families of linear codes in which each IOPP query is answered by a field element. The natural computational model in this case is that of arithmetic complexity, i.e., for a linear code C over a finite field \mathbb{F} , it is the number of arithmetic operations in \mathbb{F} made by the verifier to reach its decision.*

Main Theorem The finite field of size q is denoted here by \mathbb{F}_q ; when q is clear from context we omit it. A field is called *binary* if $q = 2^m, m \in \mathbb{N}$. A subset S of a binary field is an *additive coset* if it is a coset of a subgroup of the additive group \mathbb{F}^+ , i.e., if S is an additive shift of an \mathbb{F}_2 -linear space contained in \mathbb{F}_q . The *binary additive RS code family* is the collection of codes $\text{RS}[\mathbb{F}, S, \rho]$ where \mathbb{F} is a binary field and S an additive coset. This family of codes is one for which quasilinear PCPP were defined in [BS08], and our main theorem is stated for it (see Table 1).

Theorem 1.3 (Main — FRI properties). *The binary additive RS code family of rate $\rho = 2^{-\mathcal{R}}, \mathcal{R} \geq 2, \mathcal{R} \in \mathbb{N}$ has an IOPP (FRI) with the following properties, where $N = |S|$ denotes block-length (which equals the prover-side input length for a fixed $\text{RS}[\mathbb{F}, S, \rho]$ code) and $\rho N > 16$:*

- **Prover Complexity** *is less than $6N$ arithmetic operations in \mathbb{F} ; proof length is less than $N/3$ field elements and round complexity is at most $\frac{\log N}{2}$;*

³The definition of an IOPP can be generalized to arbitrary languages; we study an IOPP for a specific family of codes so prefer to limit the scope of our definition accordingly.

- **Verifier complexity** Query complexity is $2 \log N$; the verifier decision is computed using at most $21 \log N$ arithmetic operations over \mathbb{F}
- **Soundness:** There exists $\delta_0 \geq \frac{1}{4}(1 - 3\rho) - \frac{1}{\sqrt{N}}$ such that every f that is δ -far in relative Hamming distance from the code, is rejected with probability at least $\min\{\delta, \delta_0\} - \frac{3N}{|\mathbb{F}|}$
- **Parallelization:** Each prover-message can be computed in $O(1)$ time on a Parallel Random Access Machine (PRAM) with common read and exclusive write (CREW), assuming a single \mathbb{F} arithmetic operation takes unit time.

Generalizing Theorem 1.3 to arbitrary rate $\rho \in (0, 1]$ can be done as described in [BS08, Proposition 6.13] (cf. remark 6.2 there); this leads to slightly larger constants in the prover and verifier complexity. For practical applications like ZK-IOPs [BCGV16, BCF⁺16], rates of the form stated in the theorem above suffice.

Remark 1.4 (FRI for “smooth codes”). We call a multiplicative group $H \subset \mathbb{F}_q$ smooth if its order $(|H|)$ is 2^k for $k \in \mathbb{N}$. The family of smooth RS codes of rate ρ is the set of $\text{RS}[\mathbb{F}_q, H, \rho]$ codes in which H is a smooth multiplicative group. Theorem 1.3 holds also with respect to the family of smooth RS codes, with somewhat smaller constants than 6 and 21 for the prover and verifier arithmetic complexity (as explained in Remark 4.8); see Section 2.1 for a high-level overview of the smooth case and Remark 3.1 for more details on modifying the protocol to this case. The protocol can be further generalized to groups of order c^k for constant c (perhaps with different arithmetic complexity constants), details omitted.

The soundness bound of Theorem 1.3 is nearly tight for $\delta \leq \delta_0$. We conjecture that a similar bound holds for all δ . See Conjecture 3.4 for a more detailed version of the conjecture that implies it, and a discussion of Equation (1)

Conjecture 1.5. The soundness limit δ_0 of Theorem 1.3 approaches $1 - \rho$. Specifically, for all $\delta \leq 1 - \rho$, the rejection probability of any f that is δ -far from the RS-code of rate ρ and block-length N over \mathbb{F} , is at least

$$\delta - \frac{2 \log N}{\sqrt{|\mathbb{F}|}}. \tag{1}$$

1.2 Applications to transparent zero knowledge implementations

Prover-efficient IOPPs of the kind presented here are crucially needed to facilitate *practical* ZK argument systems that are (i) *transparent* (public randomness), (ii) *universal* — apply to any computation — and (iii) *succinctly* verifiable, meaning that verification time is negligible compared to naïve execution time (as will be reported in a follow up paper [BBHR17]). In this section we explain how our system could be incorporated in a larger practical ZK system, and in Section 1.3 briefly discuss the range of block-lengths that might be relevant in applications and the resulting communication complexity arising from their use.

The seminal works of Babai et al. [BFL90, BFLS91] showed that verifying the correctness of an arbitrary nondeterministic computation running for $T(N)$ steps can be achieved by a verifier running in time $\text{poly}(N, \log T(N))$ in the PCP model. Kilian’s construction transforms such PCPs into a 4-round ZK argument in which the total communication complexity and verifier running time are bounded by $\text{poly} \log T(N)$ [Kil92] (cf. [KPT97, IMSX15, IW14]), assuming a family of

collision-resistant hash functions. Micali further compressed this system into a non-interactive computationally sound (CS) proof system, assuming both prover and verifier share access to the same random function [Mic00a]; this is typically realized in practice using a hash function like SHA2 and relying on the Fiat-Shamir heuristic [FS86]. No implementation of these marvelous techniques has appeared during the quarter century that has passed since they were first published. This is explained, in part, by concerns about the efficiency of these constructions for concrete programs and run-times. Among the numerous components involved in building these systems, a significant computational bottleneck is that of computing solutions to the Reed-Muller (RM) proximity problem, also known as “low degree testing” of multivariate polynomials.

Quasilinear PCPs based on RS codes have prover complexity that is asymptotically more efficient than RM codes, and a number of works have explored the concrete efficiency of these protocols [BCGT13, BBGR16a]. Recently, Ben-Sasson et al. suggested an IOP with perfect zero knowledge (PZK) for NP [BCGV16], later extended to NEXP [BCF⁺16], in which prover complexity is quasilinear and verifier complexity is $\text{poly}(N, \log T(N))$; this PZK-IOP can be compiled, using Kilian’s technique, into an interactive ZK argument with succinct⁴ communication complexity, or, using Micali’s technique (cf. [Val08]), into a *non-interactive random oracle proof (NIROP)* as defined in [BCS16]. In light of this, the practicality of Kilian- and Micali-type ZK argument systems with polylogarithmic verifiers should be reconsidered.

To add motivation, a number of interesting practical succinct argument systems (with and without zero-knowledge) have been reported recently (see [WB15] for an excellent updated survey of the subject and [BBC⁺16] for a comparison of PCP/IOP-based solutions to other approaches). A particular system based on the *quadratic span programs* (QSP) of Gennaro et al. [GGPR13] (cf. [BCG⁺13]) has been used by Ben-Sasson et al. to build a decentralized anonymous payment (DAP) system called “Zerocash” [BCG⁺14], later deployed as a practical commercial crypto-currency called “Zcash” [Pec16, HBHW17]. However, the QSP based ZK system used in Zerocash/Zcash, called a “preprocessing SNARK” [BCCT12], requires a setup phase that involves *private* randomness; additionally, it relies on rather strong cryptographic “knowledge of exponent” assumptions, and quantum computers can create pseudo-proofs of falsities in polynomial time for such systems [Sho94] (cf. [PZ03]). In contrast, the aforementioned succinct interactive and non-interactive (NIROP) systems based on quasilinear PZK-IOPs require only public randomness for their setup, and the only cryptographic assumption required to realize them⁵ is the existence of a family of collision resistant hash functions [Kil92], in particular, they are not known to be breakable by quantum computers in polynomial time. Therefore, there is great interest in understanding whether succinct (interactive and non-interactive) ZK argument systems which require only public randomness (and resistant to known polynomial time quantum algorithms) can be *practically* built and used, say, by Zcash. Ben-Sasson et al. [BBC⁺16] describe such an implemented system, called “succinct computational integrity (SCI)” which is not ZK and has comparatively large communication complexity⁶. As mentioned above, we hope to incorporate the RS proximity solution described in Theorem 1.3 within practical ZK systems [BBHR17].

⁴Here, as in past works, “succinct” is synonymous to “polylogarithmic”.

⁵To reach a (non-interactive) computationally sound (CS) proof [Mic00b], the “random oracle” is assumed, and realized in practice by relying on the Fiat-Shamir heuristic. In particular, this approach as well is not known to be breakable by quantum computers in polynomial time.

⁶Communication complexity in SCI is on the order of tens of megabytes long, compared with QSP based zk-SNARKs that are shorter than 300 Bytes.

1.3 Concrete degree, communication, and round complexity

In this section we *briefly* discuss the “size” of RS codes that would be needed for various practical applications and the effect of logarithmic round complexity on security. Due to space limitations, and because the focus of this paper is *theoretical* (within the information theoretic IOP model), we omit implementation details and point the interested reader to Appendix B and [BBC⁺16, BCGV16].

The message length of RS codes of degree $d = \rho \cdot N - 1$ is precisely d , so we start by recounting the range of degrees (message sizes) that seem practically relevant. Later we calculate the communication complexity arising from using the FRI protocol to argue proximity to codes of practically relevant block-lengths, and end by discussing the practical implications of an IOPP with $\log d$ rounds. Throughout this section $\rho = 1/8$ ($N = 8 \cdot d$) because this setting is used in prior [BBC⁺16] and future [BBHR17] works.

RS block-length of systems realized in code The recently realized IOP-based argument system called SCI (“Scalable Computational Integrity”) reduces computational statements, like “*the output of program P on input x equals y after T steps*” to a *pair* of RS-proximity testing problems. SCI uses an IOP version of the quasilinear PCP of [BS08], which could be replaced with FRI. Programs bench-marked by SCI were executed on a simple MIPS-like virtual machine called *TinyRAM* [BSCG⁺]. Generally speaking, RS degree increases in size with the number of TinyRAM machine cycles T . Figure 1.A plots the degree d as a function of T for a specific simple program, showing that $d \approx T \cdot 2^{21}$.

For crypto-currency applications requiring zero knowledge, block-length will be dominated by the type of cryptographic primitives required, and the number of times they are invoked within a computational statement. For instance, *ZK contingent payments* [Max11] require a single hash, and Zerocash’s *Pour circuit* [BCG⁺14] uses 64 hash invocations, leading in that work to RS codewords (over a prime field) with degree (=number of gates) approximately 2^{22} . Our new work in progress shows that a single hash invocation requires RS block-length between $2^{12} = 4096$ (for a Davies–Meyer hash based on AES128) to 2^{19} (for SHA2), meaning that degrees in the range $d \in [2^{12}, 2^{26}]$ are relevant for existing crypto-currency (ZK) applications [BBHR17].

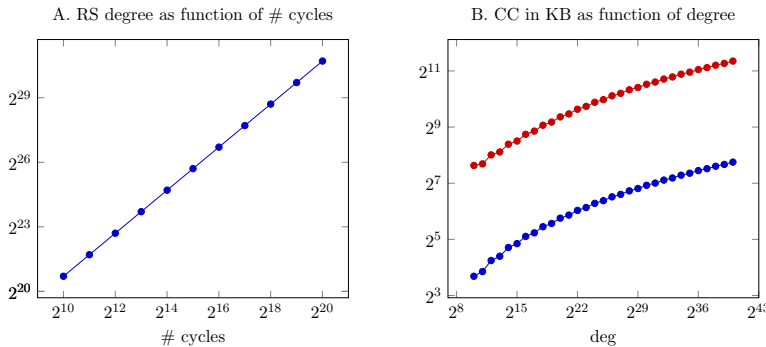


Figure 1: A. Degree of RS code arising from the *exhaustive subset sum* program [BBC⁺16, Appendix C], as a function of the number of TinyRAM machine cycles. B. Communication complexity (CC) as a function of degree, using $\lambda = 160$ bits, field size 2^{64} , soundness error $\epsilon = 2^{-80}$, and maximal proximity parameter $\delta = 1 - \rho$. The higher (red) graph corresponds to proven soundness (Theorem 3.3) and the lower (blue) corresponds to conjectured soundness (Conjecture 1.5). Both plots use code rate $\rho = 1/8$.

Estimated communication complexity and argument length The practical realization of interactive proof systems (see Section 1.2) into interactive argument systems [Kil92] and CS proofs [Mic00b] can be extended to the IOPP model, in which multiple rounds of interaction are used [BCS16]. Using Kilian’s scheme [Kil92], during the i th round the prover sends the root $\text{root}^{(i)}$ of a Merkle hash tree $\text{Tree}^{(i)}$ whose leaves are labeled by entries of $f^{(i)}$, and the verifier replies with randomness. Using Micali’s scheme [Mic00b], the (non-interactive) prover queries the random oracle with $\text{root}^{(i)}$ to “simulate” the verifier’s i th message. When verifier queries to $f^{(i)}$ are answered by the prover, each answer is accompanied by an *authentication path* (AP) that shows the query answer is consistent with $\text{root}^{(i)}$. Let $\text{CC}_{\delta,\epsilon}(N)$ denote the prover-side communication complexity (in bits) of an argument/CS proof realized by applying the Kilian/Micali scheme to FRI, where δ is the proximity parameter and ϵ is the error bound, i.e., words that are δ -far from the RS code are rejected with probability $< \epsilon$. Then

$$\text{CC}_{\delta,\epsilon}(N) = \mathbf{q}_{\delta,\epsilon} \cdot \log |\mathbb{F}| + \text{AP}_{\delta,\epsilon} \cdot \lambda \quad (2)$$

where $\mathbf{q}_{\delta,\epsilon}$ denotes total query complexity in the IOP model to reach soundness $\geq 1 - \epsilon$ for proximity parameter δ , $\text{AP}_{\delta,\epsilon}$ is the number of nodes in the sub-forest of the Merkle trees $\text{Tree}^{(0)}, \dots, \text{Tree}^{(r)}$ induced by all authentication paths, and λ is the number of output bits of the hash function used to construct the Merkle trees. In our preliminary results [BBHR17] we use $\lambda = 160$, $\epsilon = 2^{-80}$, $|\mathbb{F}| = 2^{64}$ and $\rho = 1/8$. Figure 1.B plots the communication complexity for this setting under the proven soundness of Theorem 1.3 and the (better) soundness of Conjecture 1.5. In both cases we use maximally large distance $\delta = 1 - \rho = 7/8$ to show the concrete difference in communication complexity between the proven and conjectured soundness. This plot also motivates the quest for improving the soundness analysis of Theorem 1.3.

Round complexity considerations Assuming that a crypto-currency block-chain serves as a time-stamping service for public messages and a public beacon of randomness, one may use block-chains to simulate verifier messages. Several block-chains (including Zcash) generate blocks every 2.5 minutes, which means that a FRI proof for $d = 2^k$ will take roughly $k \cdot \frac{5}{4}$ minutes to complete, or less than 1 hour⁷ for $d < 2^{40}$.

For fixed d , the round complexity stated in Theorem 1.3 is $\frac{1}{2} \log d$, but the more refined Theorem 3.3 gives a trade-off between query (\mathbf{q}) and round (\mathbf{r}) complexity, of the form $\mathbf{r} = \log d / \log \mathbf{q}$, allowing further reduction in round complexity in exchange for larger communication complexity.

Finally, the Random Oracle model used by Micali to “compress” interactive argument systems (like Kilian’s) into CS proofs applies equally to multi-round IOPs like FRI, with negligible impact on argument length; see [BCS16, Remark 1.6] for a detailed discussion. Practically speaking, those who treat hash functions like SHA2 as realizations of the RO model (a position taken by Bitcoin and other crypto-currency miners), might feel comfortable compiling IOP protocols like FRI into succinct non-interactive arguments, as described in [BCS16].

1.4 Related works

High-rate LTCs Locally testable codes (LTCs) are error correcting codes for which — by definition — prover complexity and proof length equal 0 (as stated for the case of RS codes by Rubinfeld

⁷Compare this with Bitcoin’s “best practice” of waiting 1 hour for confirmations, or 3 days required to clear standard checks.

and Sudan [RS92]); in other words, when focusing solely on prover complexity, LTCs offer an optimal solution (zero complexity). Nevertheless, as discussed in Section 1.2, the specific question of small prover complexity for RS codes is highly relevant because of its applications to practical ZK-IOPs.

Classical “direct” constructions of LTCs, such as the Hadamard code studied by Blum, Luby and Rubinfeld [BLR93] and the log N -variate RM codes used early PCP constructions [ALM⁺92, BFLS91] have sub-constant rate, thus lead to long proofs and large PCP prover complexity.

More recently, there has been remarkable progress on constructing locally testable codes (LTCs) with small query complexity and large soundness. Kopparty et al. obtained such codes with rate approaching 1 [KMRS16] and Gopi et al. presented LTCs that reach the Gilbert Varshamov bound [GKdO⁺17]. These LTCs have super-polylogarithmic query complexity. Additionally, in contrast to RS codes, we are not aware of PCP constructions with similar parameters nor do we know how to convert these LTCs into PCPs.

PCPs and IOPs: A number of recent works have considered PCP constructions with small proof length and query complexity. In addition to the aforementioned works on quasilinear PCPs, Moshkovitz and Raz constructed PCPs with optimally small query complexity (measured in bits) and proofs of length $N^{1+o(1)}$ [MR10], where N denotes the length of the NP statement (like a 3CNF) for which the PCP is constructed, achieving better soundness than Håstad’s result [Hås01]. A different line of works attempts to optimize the *bit-length* of PCP proofs; the state of the art, due to Ben-Sasson et al., achieves PCPs of bit-length $O(N)$ and query complexity N^ϵ [BKK⁺16]. In the IOP model, which generalizes PCPs by allowing more rounds of interaction, Ben-Sasson et al. presented a 2-round IOP with bit-length $O(N)$, constant query complexity (measured in bits) and constant soundness [BCG⁺16]. (Prover arithmetic complexity in all of these systems is super-linear.)

Soundness amplification: A number of results in the PCP literature have suggested techniques for improving soundness of general PCP constructions, including the parallel repetition theorem of Raz [Raz95], the gap amplification technique of Dinur [Din07] and direct-product testing, introduced by Goldreich and Safra [GS00] (cf. [DG08, IKW12]). These techniques lead to excellent soundness bounds with small query complexity. The concrete prover complexity of PCPs and PCPPs associated with these methods has not been studied in prior works but prover complexity is at least super-linear, and often polynomially large.

Doubly-efficient “proofs for muggles”: A recent line of works, initiated by Goldwasser, Kalai and Rothblum [GKR08], revisits the IP model which is equivalent to PSPACE [LFKN92, Sha92], focusing on *doubly efficient* systems in which the prover runs in polynomial time (as opposed to polynomial space, as in the aforementioned results) and verifier runs in nearly linear time. The state of the art along this line is due to Reingold et al. [RRR16], they construct doubly-efficient IP protocols with a constant number of rounds for a family of languages in P. Prover complexity in this line of works is at least super-linear, and typically polynomially large and verifier complexity is super-polylogarithmic, and often super-linear as well (cf. [CMT12, RRR16]).

2 Overview of the FRI IOPP and its soundness

In this section we consider the task of building an IOPP for a “smooth” RS code (defined below). We start in Section 2.1 by considering the *completeness* case, where we describe the interaction between the verifier and an *honest prover* attempting to prove membership in the RS code of a valid

codeword $f^{(0)}$. The IOPP protocol is explained in similarity to the Inverse Fast Fourier Transform (IFFT) [CT65]. Then, in Section 2.2, we consider the *soundness* case, where we assume $f^{(0)}$ is far in relative Hamming distance from the code and need to prove lower bounds on the verifier’s rejection probability. Soundness analysis is the most challenging aspect of our work (as it is for all prior PCPP/IOPP works). Our analysis uses the soundness analysis of the quasilinear RS-PCPP [BS08] for the case of “large” Hamming distance (beyond the unique decoding radius of the code), and presents a novel, tighter, analysis for “small” Hamming distance (below that radius).

2.1 FRI overview and similarity to the Fast Fourier Transform (FFT)

Let $\omega^{(0)}$ generate a *smooth* multiplicative group of order $N = 2^n$ (see Remark 1.4), denoted $L^{(0)}$, that is contained in a field \mathbb{F} ; in signal processing applications $\omega^{(0)}$ is a complex root of unity of order 2^n and \mathbb{F} is the field of complex numbers (we shall use a different setting). Assume the prover claims that $f^{(0)} : L^{(0)} \rightarrow \mathbb{F}$ is a member of $\text{RS}[\mathbb{F}, L^{(0)}, \rho]$, i.e., $f^{(0)}$ is the evaluation of an unknown polynomial $P^{(0)}(X) \in \mathbb{F}[X]$, $\deg(P) < \rho 2^n$; for simplicity we assume $\rho = 2^{-\mathcal{R}}$ and \mathcal{R} is a positive integer. The task of the verifier is to distinguish between truisms ($f^{(0)} \equiv P^{(0)}$ for some low degree $P^{(0)}$) and cases where $f^{(0)}$ is far from all polynomials of degree $< \rho 2^n$. Recalling the IFFT, if $f^{(0)} \equiv P^{(0)}$ there exist polynomials $P_0^{(1)}, P_1^{(1)} \in \mathbb{F}[Y]$, $\deg(P_0^{(1)}, P_1^{(1)}) < \frac{1}{2}\rho 2^n$ such that

$$\forall x \in L^{(0)} \quad f^{(0)}(x) = P^{(0)}(x) = P_0^{(1)}(x^2) + x \cdot P_1^{(1)}(x^2),$$

or, letting $Q^{(1)}(X, Y) \triangleq P_0^{(1)}(Y) + X \cdot P_1^{(1)}(Y)$ and defining $q^{(0)}(X) \triangleq X^2$, we have

$$P^{(0)}(X) \equiv Q^{(1)}(X, Y) \pmod{Y - q^{(0)}(X)} \tag{3}$$

where $\deg_X(Q^{(1)}) < 2$ and $\deg_Y(Q^{(1)}) < \frac{1}{2}\rho 2^n$. The map $x \mapsto q^{(0)}(x)$ is 2-to-1 on $L^{(0)}$ and the output of this map is a multiplicative group of order 2^{n-1} that we shall denote by $L^{(1)}$. Moreover, for every $x^{(0)} \in \mathbb{F}$ and $y \in L^{(1)}$, the value of $Q(x^{(0)}, y)$ can be computed by querying two entries of $f^{(0)}$ because $\deg_X(Q) < 2$ (the two entries are the two roots of the polynomial $y - q^{(0)}(X)$).

Our verifier thus samples $x^{(0)} \in \mathbb{F}$ uniformly at random and requests the prover to send as its first oracle a function $f^{(1)} : L^{(1)} \rightarrow \mathbb{F}$ that is supposedly the evaluation of $Q^{(1)}(x^{(0)}, Y)$ on $L^{(1)}$. Assuming $f^{(0)} \in \text{RS}[\mathbb{F}, L^{(0)}, \rho]$, the discussion above shows that $f^{(1)} \in \text{RS}[\mathbb{F}, L^{(1)}, \rho]$. Notice that there exists a 3-query test for the consistency of $f^{(0)}$ and $f^{(1)}$, we call it the *round consistency test*:

1. sample a pair of distinct elements $s_0, s_1 \in L^{(0)}$ such that $s_0^2 = s_1^2 = y$; in other words, sample a uniform $y \in L^{(1)}$ and let s_0, s_1 be the two roots of the polynomial $y - X^2$;
2. query $f^{(0)}(s_0), f^{(0)}(s_1)$ and $f^{(1)}(y)$, denote the query answers by α_0, α_1 and β , respectively;
3. interpolate the “line” through (s_0, α_0) and (s_1, α_1) , i.e., find the polynomial $p(X)$ of degree at most 1 that satisfies $p(s_0) = \alpha_0$ and $p(s_1) = \alpha_1$; notice p is unique and well-defined because $s_0 \neq s_1$;
4. accept if and only if $p(x^{(0)}) = \beta$ and otherwise reject;

Tallying the costs of the first round, the verifier sends a single field element ($x^{(0)}$) and the prover responds with a message (oracle) $f^{(1)} : L^{(1)} \rightarrow \mathbb{F}$ evaluated on a domain that is half the size of $L^{(0)}$;

testing the consistency of $f^{(0)}$ and $f^{(1)}$ requires three field elements per test (repeating the test boosts soundness). We thus reduced a single proximity problem of size 2^n and rate ρ to a single analogous problem of size 2^{n-1} and same rate. Repeating the process for $r = n - \mathcal{R}$ rounds leads to a function $f^{(r)}$ that is supposedly of constant degree and evaluated over a domain of constant size $2^{\mathcal{R}}$, so at this point the prover sends the single constant that describes the function, and the verifier uses it as $f^{(r)}$ in the last round consistency test, the one that tests consistency of $f^{(r-1)}$ and $f^{(r)}$.

Applying inductive analysis to all r rounds, if $f^{(0)} \in \text{RS}[\mathbb{F}, L^{(0)}, \rho]$ (and the prover is honest) then all r round consistency tests pass with probability 1 and $f^{(r)}$ is indeed a constant function. In other words, the protocol we described has perfect completeness.

Differences between informal and actual protocol are mostly technical; we list them now. The field \mathbb{F} is finite and *binary*, i.e., of characteristic 2; nevertheless the construction and analysis can be immediately applied to RS codes evaluated over smooth multiplicative groups (of order 2^n), as explained informally above (cf. Remarks 1.4 and 3.1). In binary fields, the natural evaluation domains (like $L^{(0)}, L^{(1)}$ above) are cosets of *additive* groups (not multiplicative ones), i.e., $L^{(i)}$ is an affine shift of a linear space over \mathbb{F}_2 . The map $q^{(0)}(X) = X^2$ is *not* 2-to-1 on $L^{(0)}$ (in binary fields it is a 1-to-1 map, a Frobenius automorphism of \mathbb{F} over \mathbb{F}_2) so we use a different polynomial $q^{(0)}(X)$ that is many-to-one on $L^{(0)}$ and such that the set $L^{(1)} = \{y = q^{(0)}(x) \mid x \in L^{(0)}\}$ is a coset of an *additive* group, like $L^{(0)}$, but of smaller size ($|L^{(1)}| \ll |L^{(0)}|$); the polynomial $q^{(0)}$ is known as an *affine subspace* polynomial, belonging to the class of *linearized* polynomials (cf. Section 3.1). We use $q^{(0)}$ of degree 4 instead of 2 because this reduces the number of rounds from n to $n/2$ with no increase in total query complexity; notice that a similar reduction could be applied in the multiplicative setting by using $q^{(0)} = X^4$ (but we preferred simplicity to efficiency in the informal exposition above). Finally, the actual protocol performs all queries only after the prover has sent all of $f^{(1)}, \dots, f^{(r)}$. Thus, we construct a protocol with two phases. The first phase, called the **COMMIT** phase, involves r rounds. At the beginning of the i th round the prover has sent oracles $f^{(0)}, \dots, f^{(i-1)}$, and during this (i th) round the verifier samples and sends $x^{(i)}$ and the prover responds by sending the next oracle $f^{(i)}$. During the second phase, called the **QUERY** phase, the verifier applies the *round consistency test* to all r rounds. To save query complexity *and boost soundness*, the query made to $L^{(i)}$ is used to test *both* consistency of $f^{(i-1)}$ vs. $f^{(i)}$ *and* consistency of $f^{(i)}$ vs. $f^{(i+1)}$.

2.2 Soundness analysis — overview

Proof composition is a technique introduced by Arora and Safra [AS98] in the context of PCPs, adapted to PCPPs in [BGH⁺06, DR04] and optimized for the special case of the RS code in [BS08]. Informally, it reduces proximity testing problems over a large domain to similar proximity testing problems over significantly smaller domains. The process reducing $f^{(0)}$ to $f^{(1)}$ above is a special case of proof composition, and each invocation of it incurs two costs on behalf of the verifier. The first is the *query complexity* needed to check consistency of $f^{(0)}$ and $f^{(1)}$ (the “round consistency test”) and the second is the reduction in *distance*, which affects the *soundness* of the protocol. Assuming $f^{(0)}$ is $\delta^{(0)}$ -far from all codewords in relative Hamming distance, for proof composition to work one should prove that with high probability $f^{(1)}$ is $\delta^{(1)}$ -far from all codewords where $\delta^{(1)}$ depends on $\delta^{(0)}$; larger values of $\delta^{(1)}$ imply higher (better) soundness and smaller communication complexity. A benefit of the FRI protocol is that with high probability $\delta^{(1)} \geq (1 - o(1))\delta^{(0)}$, i.e.,

the reduction in distance in our protocol is *negligible*. In contrast, prior RS proximity PCPP and IOPP solutions follow the construction and analysis of [BS08] which in turn is based on the bivariate testing Theorem of Polischuk and Spielman [PS94] and incur a *constant multiplicative loss* in distance per round of proof composition ($\delta^{(1)} \leq \delta^{(0)}/2$). This loss limited the number of proof composition rounds to $\leq \log N$ and thus required replacing $q^{(0)}(X) = X^2$ with a higher degree polynomial, like $q^{(0)}(X) = X^{2^{n/2}}$. The higher degree of $q^{(0)}$ results in $Q^{(1)}(X, Y)$ having *balanced* X - and Y -degrees, namely

$$\deg_X(Q^{(1)}) \approx \deg_Y(Q^{(1)}) \approx 2^{n/2}.$$

Moving to $q^{(0)}(X)$ of *constant* degree as in FRI gives a *biased RS-IOPP* (because $\deg_X(Q^{(1)}) \ll \deg_Y(Q^{(1)})$). The main benefit of this bias is that one side of the recursive process (that of X) terminates immediately and consequently *removes* the constant multiplicative soundness loss incurred in prior works, replacing it with a negligible additive loss. More to the point, we show that for $\delta^{(0)}$ less than the unique decoding radius of the code ($\delta^{(0)} < (1 - \rho)/2$), with high probability (namely, $1 - \frac{O(1)}{|\mathbb{F}|}$) the sum of (i) the round consistency error and (ii) the “new” distance $\delta^{(1)}$ is at least as large as the “old” distance $\delta^{(0)}$. This statement is relatively straightforward to prove in case the prover is *honest*, i.e., when $f^{(1)}(y) = Q^{(1)}(x^{(0)}, y)$ for all $y \in L^{(1)}$ (in this case there is no round consistency error). The challenging part of the proof is to show this also holds for non-honest provers and arbitrary $f^{(1)}$; see Lemma 4.4 and Section 4.2 for more details.

Acknowledgements

We thank Justin Drake, Peter Manohar and Nicholas Spooner for helping clarify the presentation and for pointing out and correcting errors in earlier manuscripts.

3 FRI— detailed description and main properties

In this section we give a formal and detailed description of the FRI protocol, expanding on what was explained in the previous section. We start by providing additional needed definitions, followed by the description of the COMMIT and QUERY phases; we continue by listing the properties obtained by the protocol (Theorem 3.3), and conclude the section with a proof of Main Theorem 1.3. The next section is then devoted to the proof of Theorem 3.3.

3.1 Definitions and notation

Interpolant For a function $f : S \rightarrow \mathbb{F}$, $S \subset \mathbb{F}$, let interpolant^f denote the *interpolant* of f , defined as the unique polynomial $P(X) = \sum_{i=0}^{|S|-1} a_i X^i$ of degree less than $|S|$ whose evaluation on S equals $f|_S$, i.e., $\forall x \in S \ f(x) = P(x)$. We assume the interpolant $P(X)$ is represented as a formal sum, i.e., by the sequence of monomial coefficients $a_0, \dots, a_{|S|-1}$.

Subspace polynomials Given a set $L_0 \subset \mathbb{F}$ let $\text{Zero}_{L_0} \triangleq \prod_{x \in L_0} (X - x)$ be the unique non-zero monic polynomial of degree $|L_0|$ that vanishes on L_0 . When L_0 is an additive coset contained in a binary field, the polynomial $\text{Zero}_{L_0}(X)$ is an *affine subspace polynomial*, a special type of a *linearized polynomial* [Ore33, Ore34]. We shall use the following properties of such polynomials, referring the interested reader to [LN97, Chapter 3.4] for proofs and additional background:

1. The map $x \mapsto \text{Zero}_{L_0}(x)$ maps each additive coset S of L_0 to a single field element, which will be denoted by y_S
2. If $L \supset L_0$ are additive cosets, then $\text{Zero}_{L_0}(L) \triangleq \{\text{Zero}_{L_0}(z) \mid z \in L\}$ is an additive coset and $\dim(\text{Zero}_{L_0}(L)) = \dim(L) - \dim(L_0)$.

Subspace specification Henceforth, the letter L always denotes an additive coset in a binary field \mathbb{F} , we assume all mentioned additive cosets are specified by an additive shift $\alpha \in \mathbb{F}$ and a basis $\beta_1, \dots, \beta_k \in \mathbb{F}^k$ so that $L = \left\{ \alpha + \sum_{i=1}^k b_i \beta_i \mid b_1, \dots, b_k \in \mathbb{F}_2 \right\}$; we assume α and $\vec{\beta} = (\beta_1, \dots, \beta_k)$ are agreed upon by prover and verifier.

3.2 The COMMIT phase

The protocol is parameterized by an integer η ; to prove Theorem 1.3 we set $\eta = 2$ but in other settings a different value may be more beneficial. The number of rounds is $r \triangleq \lfloor \frac{k^{(0)} - \mathcal{R}}{\eta} \rfloor$ (recall that $\mathcal{R} = \log(1/\rho)$, where ρ is the rate). During the i th round of the COMMIT phase, $i \in \{0, \dots, r-1\}$, the verifier has oracle access to a function $f^{(i)} : L^{(i)} \rightarrow \mathbb{F}$, where $\dim(L^{(i)}) = k^{(i)} = k^{(0)} - \eta \cdot i$ submitted by the prover and the spaces $L^{(i)}$ are fixed in advance and, in particular, do not depend on verifier messages.

A single COMMIT round We assume the verifier and prover have also agreed upon a fixed “small” $L_0^{(i)} \subset L^{(i)}$, $\dim(L_0^{(i)}) = \eta$, one for each $i \in \{0, \dots, r-1\}$. Let $\mathcal{S}^{(i)}$ denote all cosets of $L_0^{(i)}$ in $L^{(i)}$. Let

$$q^{(i)}(X) \triangleq \text{Zero}_{L_0^{(i)}}(X)$$

be the subspace polynomial vanishing on $L_0^{(i)}$. Let $L^{(i+1)} \triangleq q^{(i)}(L^{(i)})$. The verifier's i th message is a uniformly random $x^{(i)} \in \mathbb{F}$. The prover's next message (or oracle) is $f^{(i+1)} : L^{(i+1)} \rightarrow \mathbb{F}$ computed for each $y_S \in L^{(i+1)}$, $y_S = q^{(i)}(S)$, $S \in \mathcal{S}^{(i)}$, by interpolating the function $f^{(i)}|_S$ to obtain a polynomial $P_S^{(i)}(X)$, $\deg(P_S^{(i)}) < 2^\eta$ and then setting $f^{(i+1)}(y_S) \triangleq P_S^{(i)}(x^{(i)})$.

Termination — COMMIT During the last round ($i = r$), the prover sends the interpolant $P^{(r)}(X) = \text{interpolant}^{f^{(r)}}$ of $f^{(r)}$ rather than $f^{(r)}$ itself. By this point $\deg(P^{(r)}(X)) < \rho \cdot |L^{(r)}| \leq 2^\eta$ (recall $\eta \in \mathbb{N}$ is a constant).

FRI-COMMIT:

Common input:

- Parameters \mathcal{R}, η, i , all are positive integers:
 - rate parameter \mathcal{R} : logarithm of RS code rate ($\rho = 2^{-\mathcal{R}}$)
 - localization parameter η : dimension of $L_0^{(i)}$ (i.e., $|L_0^{(i)}| = 2^\eta$); let $r \triangleq \lfloor \frac{k^{(0)} - \mathcal{R}}{\eta} \rfloor$ denote round complexity
 - $i \in \{0, \dots, r\}$: round counter
- A parametrization of $\text{RS}^{(i)} \triangleq \text{RS}[\mathbb{F}, L^{(i)}, \rho = 2^{-\mathcal{R}}]$, denote $k^{(i)} = \log_2 |L^{(i)}|$ (notice $k^{(i)} = \dim(L^{(i)})$);
- $L_0^{(i)} \subset L^{(i)}$, $\dim(L_0^{(i)}) = \eta$; let $q^{(i)}(X) = \text{Zero}_{L_0^{(i)}}(X)$ and denote $L^{(i+1)} = q^{(i)}(L^{(i)})$

Prover input: $f^{(i)} : L^{(i)} \rightarrow \mathbb{F}$, a purported codeword of $\text{RS}^{(i)}$ Loop: While $i \leq r$:

1. Verifier sends a uniformly random $x^{(i)} \in \mathbb{F}$
 2. Prover defines the function $f_{f^{(i)}, x^{(i)}}^{(i+1)}$ with domain $L^{(i+1)}$ thus, for each $y \in L^{(i+1)}$:
 - Let $S_y = \{x \in L^{(i)} \mid q^{(i)}(x) = y\}$ be the coset of $L_0^{(i)}$ mapped by $q^{(i)}$ to $\{y\}$;
 - $P_y^{(i)}(X) \triangleq \text{interpolant}^{f^{(i)}|_{S_y}}$;
 - $f_{f^{(i)}, x^{(i)}}^{(i+1)}(y) \triangleq P_y^{(i)}(x^{(i)})$;
 3. If $i = r$ then:
 - let $f^{(r)} = f_{f^{(r-1)}, x^{(r-1)}}^{(r)}$ for $f_{f^{(r-1)}, x^{(r-1)}}^{(r)}$ defined in step 2 above;
 - let $P^{(r)}(X) = \sum_{j \geq 0} a_j^{(r)} X^j \triangleq \text{interpolant}^{f^{(r)}}(X)$;
 - let $d = \rho \cdot |L^{(r)}| - 1$;
 - prover commits to first $d + 1$ coefficients of $P^{(r)}(X)$, namely, to $\langle a_0^{(r)}, \dots, a_d^{(r)} \rangle$
 - COMMIT phase terminates;
 4. Else ($i < r$):
 - let $f^{(i+1)} = f_{f^{(i)}, x^{(i)}}^{(i+1)}$ for $f_{f^{(i)}, x^{(i)}}^{(i+1)}$ defined in step 2 above;
 - prover commits to oracle $f^{(i+1)}$
 - both parties repeat the COMMIT protocol with common input
 - parameters $(\mathcal{R}, \eta, i + 1)$
 - a parametrization of $\text{RS}^{(i+1)} \triangleq \text{RS}[\mathbb{F}, L^{(i+1)}, \rho = 2^{-\mathcal{R}}]$ and $L_0^{(i+1)} \subset L^{(i+1)}$, $\dim(L_0^{(i+1)}) = \eta$
- and prover input $f^{(i+1)}$ defined at the beginning of this step;

Remark 3.1 (Adapting FRI to the family of smooth RS codes). *If $RS^{(0)}$ is a smooth code of block-length $2^{k^{(0)}}$ (i.e., $L^{(0)}$ is a multiplicative group of order $2^{k^{(0)}}$), the FRI protocol for $RS^{(0)}$ is obtained from the protocol above by applying the following modifications to the COMMIT and QUERY phases:*

- *define $L_0^{(i)}$ to be the multiplicative subgroup of $L^{(i)}$ of size 2^η , i.e., the set of roots of the polynomial $X^{2^\eta} - 1$*
- *define $q^{(i)}(X) = X^{2^\eta}$ (notice $q^{(i)}$ does not depend on i); observe $x \mapsto q^{(i)}(x)$ is a 2^η -to-1 map on $L^{(i)}$ and its image is a (smooth) multiplicative group;*
- *let $L^{(i+1)} = q^{(i)}(L^{(i)})$ (exactly as described in the protocol above), noticing $L^{(i+1)}$ is the (smooth) multiplicative group of order $2^{k^{(i+1)}} = 2^{k^{(i)} - \eta}$;*
- *terminology: interpret the words “coset” to mean “multiplicative coset” and the term “dimension” to mean “base-2 logarithm of group order” (e.g., $\dim(L^{(i)}) \triangleq \log_2 |L^{(i)}|$); replace the term “affine space” with “smooth group”;*

3.3 The QUERY phase

During this phase the prover does not participate and the verifier merely checks that the prover operated as specified above. Concretely, a single test of the verifier consists of sampling a uniformly random $s^{(0)} \in L^{(0)}$, and computing iteratively $s^{(i+1)} = q^{(i)}(s^{(i)})$, notice $s^{(i)} \in L^{(i)}$; let $S^{(i)}$ denote the unique coset of $L_0^{(i)}$ in which $s^{(i)}$ is contained. (Using the notation above, we have $s^{(i+1)} = y_{S^{(i)}}$.) The verifier now accepts if and only if for all $i < r$ it holds that $f^{(i+1)}(s^{(i+1)}) = \text{interpolant}^{f^{(i)}|_{S^{(i)}}}(x^{(i)})$. When $i < r$ the values $f^{(i)}(x)$ are queried directly by the verifier, and in the terminal case ($i = r$) the verifier queries all coefficients of $P^{(r)}(X)$ from the last prover message and interpolates this polynomial to reconstruct $f^{(r)}$.

FRI-QUERY:

verifier input:

- parameters \mathcal{R}, η as defined in the COMMIT phase
- repetition parameter ℓ
- sequence of rate- ρ RS-codes $\text{RS}^{(0)}, \dots, \text{RS}^{(r)}$, where $\text{RS}^{(i)} = \text{RS}[\mathbb{F}, L^{(i)}, \rho]$ and $\log_2 |L^{(i)}| = k^{(i)} = k^{(0)} - i \cdot \eta$; (notice $k^{(i)} = \dim(L^{(i)})$);
- sequence of affine spaces $L_0^{(0)}, \dots, L_0^{(r-1)}$, each $L_0^{(i)}$ is of dimension η and contained in $L^{(i)}$;
- transcript of verifier messages $x^{(0)}, \dots, x^{(r-1)} \in \mathbb{F}$
- access to oracles $f^{(0)}, \dots, f^{(r-1)}$
- access to last oracle $P^{(r)}(X) = \sum_{j=0}^d a_j^{(r)} X^j$ for $d = \rho \cdot |L^{(r)}| - 1$;

Terminal function reconstruction:

- query $a_0^{(r)}, \dots, a_d^{(r)}$; (a total of $d + 1 \leq 2^\eta$ queries)
- let $P'(X) \triangleq \sum_{j \leq d} a_j^{(r)} X^j$;
- let $f^{(r)}$ be the evaluation of $P'(X)$ on $L^{(r)}$; (notice $f^{(r)} \in \text{RS}^{(r)}$)

Repeat ℓ times: {

1. Sample uniformly random $s^{(0)} \in L^{(0)}$ and for $i = 0, \dots, r - 1$ let
 - $s^{(i+1)} = q^{(i)}(s^{(i)})$
 - $S^{(i)}$ be the coset of $L_0^{(i)}$ in $L^{(i)}$ that contains $s^{(i)}$
2. For $i = 0, \dots, r - 1$,
 - query $f^{(i)}$ on all of $S^{(i)}$; (a total of 2^η queries)
 - compute $P^{(i)}(X) \triangleq \text{interpolant}^{f^{(i)}|_{S^{(i)}}}$; (notice $\deg(P^{(i)}) < 2^\eta$)
3. **round consistency:** If for some $i \in \{0, \dots, r - 1\}$ it holds that

$$f^{(i+1)}(s^{(i+1)}) \neq P^{(i)}(x^{(i)}) \tag{4}$$

then reject and abort;

}

Return accept

3.4 Main properties of the FRI protocol

The following distance measure will be used in our soundness analysis. It is similar to the relative Hamming distance, only measured on blocks of symbols. Given a function $f : S \rightarrow \Sigma$ and $S' \subset S$ we denote by $f|_{S'}$ the *restriction* of f to domain S' . Given $g : S \rightarrow \Sigma$, let $f|_{S'} = g|_{S'}$ denote equality in the space $\Sigma^{S'}$, i.e., this equality holds iff for each $x \in S'$ we have $f(x) = g(x)$.

Definition 3.2 (Block-wise distance measure). *Let $\mathcal{S} = \{S_1, \dots, S_m\}$ be a partition of a set S and Σ be an alphabet. The relative \mathcal{S} -Hamming distance measure on Σ^S is defined for $f, g \in \Sigma^S$ as the relative Hamming distance over $\Sigma^{S_1} \times \dots \times \Sigma^{S_m}$,*

$$\Delta^{\mathcal{S}}(f, g) \triangleq \Pr_{i \in [m]} [f|_{S_i} \neq g|_{S_i}] = \frac{|\{i \in [m] \mid f|_{S_i} \neq g|_{S_i}\}|}{m}. \quad (5)$$

Thus, for $\mathcal{F} \subset \Sigma^S$ let $\Delta^{\mathcal{S}}(g, \mathcal{F}) = \min \{\Delta^{\mathcal{S}}(g, f) \mid f \in \mathcal{F}\}$.

In our soundness analysis of the FRI protocol we use the block-wise distance on $\mathbb{F}^{L^{(i)}}$ corresponding to the partition of $L^{(i)}$ to cosets of $L_0^{(i)}$; recall both $L^{(i)}, L_0^{(i)}$ are affine spaces with $L_0^{(i)} \subset L^{(i)}$, and $\mathcal{S}^{(i)}$ denotes the set of cosets of $L_0^{(i)}$ in $L^{(i)}$. To simplify notation we denote

$$\Delta^{(i)}(f, g) \triangleq \Delta^{\mathcal{S}^{(i)}}(f, g) \quad (6)$$

In words, $\Delta^{(i)}(\cdot, \cdot)$ measures the fraction of cosets of $L_0^{(i)}$ in $L^{(i)}$ on which f and g do not agree completely. Recalling $\text{RS}^{(i)}$ is a code of rate ρ we have

$$1 - \rho \geq \Delta^{(i)}(f^{(i)}, \text{RS}^{(i)}) \geq \Delta_{\text{H}}(f^{(i)}, \text{RS}^{(i)}) \quad (7)$$

where $\Delta_{\text{H}}(\cdot, \cdot)$ denotes relative Hamming distance. The first inequality holds because there always exists a polynomial of degree $< \rho|L^{(i)}|$ that agrees completely with $f^{(i)}$ on a ρ -fraction of cosets of $L_0^{(i)}$. The second inequality holds because if $f^{(i)}$ differs from $g \in \text{RS}^{(i)}$ on a δ -fraction of cosets in $\mathcal{S}^{(i)}$ then f and g differ on at most a δ -fraction of their entries because all cosets in $\mathcal{S}^{(i)}$ are of equal size.

The following theorem is a more detailed and precise version of Theorem 1.3.

Theorem 3.3 (Main properties of the FRI protocol). *The following properties hold when the FRI protocol is invoked on oracle $f^{(0)} : L^{(0)} \rightarrow \mathbb{F}$ with localization parameter η and rate parameter \mathcal{R} (and rate $\rho = 2^{-\mathcal{R}}$) such that $\rho|L^{(0)}| > 16$:*

1. **Completeness** *If $f^{(0)} \in \text{RS}^{(0)} \triangleq \text{RS}[\mathbb{F}, L^{(0)}, \rho = 2^{-\mathcal{R}}]$ and $f^{(1)}, \dots, f^{(r)}$ are computed by the prover specified in the COMMIT phase, then the FRI verifier outputs accept with probability 1.*

2. **Soundness** *Suppose $\delta^{(0)} \triangleq \Delta^{(0)}(f^{(0)}, \text{RS}^{(0)}) > 0$. Then with probability at least*

$$1 - \frac{3|L^{(0)}|}{|\mathbb{F}|} \quad (8)$$

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

$$\left(1 - \min \left\{ \delta^{(0)}, \frac{1 - 3\rho - 2^\eta / \sqrt{|L^{(0)}|}}{4} \right\}\right)^\ell \quad (9)$$

Consequently, the soundness of FRI is at least

$$\mathbf{s}^-(\delta^{(0)}) \triangleq 1 - \left(\frac{3|L^{(0)}|}{|\mathbb{F}|} + \left(1 - \min \left\{ \delta^{(0)}, \frac{1 - 3\rho - 2^\eta / \sqrt{|L^{(0)}|}}{4} \right\}\right)^\ell \right). \quad (10)$$

3. **Prover complexity** The i^{th} step of commit phase can be computed by a parallel random access machine (PRAM) with concurrent read and exclusive write (CREW) in $2\eta + 3$ cycles — each cycle involves a single arithmetic operation in \mathbb{F} — using $2|L^{(i)}| + \eta$ processors and a total of $4|L^{(i)}|$ arithmetic operations over \mathbb{F} .

Consequently, the total prover complexity is at most $6|L^{(0)}|$ arithmetic operations, which can be carried out in at most $4 \log |L^{(0)}|$ cycles on a PRAM-CREW with $2n + 3$ processors.

4. **Verifier complexity** Verifier communication during the COMMIT phase equals r field elements; query complexity (during QUERY phase) equals $\ell 2^\eta r = \ell 2^\eta \left(1 + \lfloor \frac{\log |L^{(0)}| - \mathcal{R}}{\eta} \rfloor\right)$. On a PRAM with exclusive read and exclusive write (EREW) with $\ell r \cdot 2^\eta$ processors, the verifier's decision is obtained after $2\eta + 3 + \log \ell$ cycles and a total of $\ell \cdot r \cdot (6 \cdot 2^\eta + 6\eta + 6)$ arithmetic operations in \mathbb{F} .

For $\delta^{(0)} \leq \frac{1 - 3\rho - 2^\eta / \sqrt{|L^{(0)}|}}{4}$ the soundness lower bound nearly matches the trivial upper bound of rejection probability $\leq \delta^{(0)}$ which is obtained by a prover that sets $f^{(1)}$ to be the codeword of $\text{RS}^{(1)}$ that is closest to $f_{f^{(0)}, x^{(0)}}^{(1)}$. We conjecture that soundness, as a function of the block-wise distance measure of Definition 3.2, approaches $\delta^{(0)}$ for any distance value $\delta^{(0)}$, even when $\delta^{(0)} \approx 1 - \rho$. The following conjecture formalizes this, and implies Conjecture 1.5 by setting $\eta = 2$ and using Equation (7). We discuss the justification for this conjecture, and for the particular choice of Equation (11), in Section 4.5.

Conjecture 3.4. Using the notation of Theorem 3.3, for any $f : L^{(0)} \rightarrow \mathbb{F}, |L^{(0)}| = 2^{k^{(0)}}$ and $\epsilon \in [0, \delta^{(0)}]$, with probability at least

$$1 - \frac{(k^{(0)})^2 \cdot 2^\eta}{\epsilon \cdot \eta^2 \cdot |\mathbb{F}|} \quad (11)$$

over the randomness of the verifier during the COMMIT phase, the probability of rejection during the QUERY phase is at least $1 - \delta^{(0)}(1 - \epsilon)$. Setting $\epsilon = \frac{k^{(0)} 2^{\eta/2}}{\eta \sqrt{|\mathbb{F}|}}$, the overall rejection probability of $f^{(0)}$ by the FRI verifier is at least $\delta^{(0)} - \frac{2k^{(0)} 2^{\eta/2}}{\eta \sqrt{|\mathbb{F}|}}$.

3.5 Proof of Main Theorem 1.3

We now show that the Theorem above indeed proves our Main Theorem 1.3 and the remainder of the paper is devoted to proving Theorem 3.3.

Proof of Main Theorem 1.3. Apply Theorem 3.3 with $N = |L^{(0)}|$ and $k = k^{(0)} = \dim(L^{(0)})$. Fix $\eta = 2$ and $\ell = 1$. Prover complexity follows immediately from Theorem 3.3, part 3. By construction $\dim(L^{(i)}) = \dim(L^{(0)}) - i\eta$ and thus, using the geometric series formula, the total proof length is

$$|L^{(1)}| + \dots + |L^{(r)}| = |L^{(0)}| \cdot \sum_{i=1}^r \frac{1}{2^{i\eta}} = |L^{(0)}| \cdot \sum_{i=1}^r 4^{-i} < |L^{(0)}|/3.$$

Round complexity is

$$r = \lfloor (k - \mathcal{R})/\eta \rfloor \leq k/2 - 1$$

the last inequality follows because $\mathcal{R} \geq \eta$. This completes the proof of the first bullet (“prover”) of Theorem 1.3.

Moving to the second bullet (“verifier”), query complexity is at most $2 \log N$ for our selection of $\eta = 2$ and the resulting value of r . The decision complexity of the verifier follows immediately from Theorem 3.3, part 4 using the setting of r , η and ℓ . This completes the proof of the second bullet.

The lower bound on soundness (third bullet) follows from (10) by setting $\ell = 1$; although (10) is stated for the blockwise distance measure, the same bound holds also with respect to the relative Hamming distance measure; this follows from (7).

The parallelization bullet follows from Theorem 3.3, part 3. This completes the proof of Main Theorem 1.3. \square

4 Proof of Theorem 3.3

We prove the items of Theorem 3.3 in the order stated there. The main technical challenge is that of proving soundness lower bounds in Section 4.2.

4.1 Completeness — Part 1

The proof of the completeness claim follows from the following lemma.

Lemma 4.1 (Inductive argument). *If $f^{(i)} \in \text{RS}^{(i)}$ then for all $x^{(i)} \in \mathbb{F}$ it holds that $f_{f^{(i)}, x^{(i)}}^{(i+1)} \in \text{RS}^{(i+1)}$.*

We complete the proof of completeness assuming the lemma above, then prove the lemma.

Proof of Theorem 3.3, item 1 (perfect completeness). If one applies the prover specified in the COMMIT phase to an arbitrary function, then for any $i < r - 1$ all round consistency tests pass because the equality (4) checked by the verifier is fulfilled by the construction of $f^{(i+1)}$ from $f^{(i)}$ described in step 2 of the COMMIT phase.

Thus we need only prove that the round consistency test passes also for $i = r - 1$. By assumption $\delta^{(0)} = 0$ so $f^{(0)} \in \text{RS}^{(0)}$. Applying Lemma 4.1 inductively shows that $f^{(r)} \in \text{RS}^{(r)}$ which means that its interpolant is a polynomial of degree $< \rho \cdot |L^{(r)}|$. In this case the function $f'^{(r)}$ extracted from

$P^{(r)}(X)$ in the “terminal function reconstruction” step is indeed equal to $f^{(r)}$ and hence all round consistency tests pass for $i = r - 1$ as well. This completes the proof. \square

4.1.1 Proof of Lemma 4.1

For our proof, we need the following claim from [BS08, Section 6] and repeat its proof for self-containment. We use capitalized letters like X, Y to denote formal variables and non-capitalized ones like x, y to denote field elements.

Claim 4.2. *For every $f^{(i)} : L^{(i)} \rightarrow \mathbb{F}$ there exists $Q^{(i)}(X, Y) \in \mathbb{F}[X, Y]$ satisfying*

1. $f^{(i)}(x) = Q^{(i)}(x, q^{(i)}(x))$ for all $x \in L^{(i)}$
2. $\deg_X(Q^{(i)}) < |L_0^{(i)}|$
3. If $f^{(i)} \in \text{RS}[\mathbb{F}, L^{(i)}, \rho]$ then $\deg_Y(Q^{(i)}) < \rho |L^{(i+1)}|$

Proof. Let $P^{(i)} = \text{interpolant}^{f^{(i)}}$. Let $\mathbb{F}[X, Y]$ denote the ring of bivariate polynomials over \mathbb{F} ; order monomials first according to total degree, then according to X -degree. Let

$$Q^{(i)}(X, Y) = P^{(i)}(X) \pmod{Y - q^{(i)}(X)} \quad (12)$$

be the remainder from dividing $P^{(i)}(X)$ by $Y - q^{(i)}(X)$. By definition, there exists $R(X, Y) \in \mathbb{F}[X, Y]$ such that

$$P^{(i)}(X) = Q^{(i)}(X, Y) + (Y - q^{(i)}(X)) \cdot R(X, Y).$$

For $x \in L^{(i)}$ and $y = q^{(i)}(x)$ the rightmost summand above vanishes, hence $P^{(i)}(x)$ equals $Q^{(i)}(x, y) = Q^{(i)}(x, q^{(i)}(x))$, implying item 1.

By the ordering chosen for monomials, the remainder Q defined in (12) satisfies

$$\deg_X(Q^{(i)}(X, Y)) < \deg(q^{(i)}) = |L_0^{(i)}|$$

and hence item 2 holds.

Finally, by the rules of division and the chosen monomial ordering,

$$\deg_Y(Q^{(i)}) = \lfloor \frac{\deg(P^{(i)})}{\deg(q^{(i)})} \rfloor = \lfloor \frac{\deg(P^{(i)})}{|L_0^{(i)}|} \rfloor < \rho |L^{(i+1)}|$$

The inequality follows because $|L^{(i+1)}| = |L^{(i)}|/|L_0^{(i)}|$ and $f^{(i)} \in \text{RS}^{(i)}$, implying $\deg(P^{(i)}) < \rho |L^{(i)}|$. We conclude item 3 holds and this proves the claim. \square

Proof of Lemma 4.1. We use the notation from Claim 4.2. From item 3 of that claim it follows that for any $x^{(i)}$ we have $\deg_Y(Q^{(i)}) < \rho \cdot |L^{(i+1)}|$. We will thus prove

$$\forall y \in L^{(i+1)}, f^{(i+1)}(y) = Q^{(i)}(x^{(i)}, y) \quad (13)$$

and this implies $\deg(f^{(i+1)}) \leq \deg_Y(Q^{(i)}) < \rho \cdot |L^{(i+1)}|$, as required.

To prove (13) fix $y \in L^{(i+1)}$ and let $S_y \in \mathcal{S}^{(i)}$ satisfy $q^{(i)}(S_y) = \{y\}$. By construction of $f^{(i+1)}$ we have

$$f^{(i+1)}(y) = \text{interpolant}^{f^{(i)}|_{S_y}}(x^{(i)}). \quad (14)$$

By Claim 4.2, item 1,

$$\forall x \in S_y, \quad f^{(i)}(x) = P^{(i)}(x) = Q^{(i)}(x, y) \quad (15)$$

And because $\deg_X(Q^{(i)}) < |S_y|$, due to Claim 4.2, item 2, we conclude that

$$\text{interpolant}^{f^{(i)}|_{S_y}}(X) = Q^{(i)}(X, y) \quad (16)$$

as formal polynomials in X , hence evaluating both polynomials on $x^{(i)}$ gives equal values. Combining this with (14) and (15) proves that (13) holds, and this completes the proof. \square

4.2 Soundness — Part 2

Soundness analysis is typically the most challenging aspect of proximity testing protocols; our case is no different. First we provide a few needed definitions, and continue in Section 4.2.2 with a statement of two main lemmas (Lemmas 4.3 and 4.4) that imply soundness. After completing the proof of soundness in that section, we prove the two main lemmas in Sections 4.2.3 and 4.2.4.

4.2.1 Definitions — round consistency error and distortion set

Given oracles $f^{(i)}$ and $f^{(i+1)}$ produced in response to verifier randomness $x^{(i)}$, we shall use the following terms and notation:

- **inner-layer distance** The i th *inner-layer distance* is the $\Delta^{(i)}$ -distance of $f^{(i)}$ from $\text{RS}^{(i)}$,

$$\delta^{(i)} \triangleq \Delta^{(i)}(f^{(i)}, \text{RS}^{(i)})$$

- **round error** For $i > 0$, the i th *round error set* is the subset of $L^{(i)}$ defined by

$$A_{\text{err}}^{(i)}(f^{(i)}, f^{(i-1)}, x^{(i-1)}) \triangleq \left\{ y_S^{(i)} \in L^{(i)} \mid \text{interpolant}^{f^{(i-1)}|_S}(x^{(i-1)}) \neq f^{(i)}(y_S^{(i)}) \right\}$$

and the i th *round error* $\text{err}^{(i)}$ is the probability that the round consistency test rejects $f^{(i)}$ and $f^{(i-1)}$,

$$\text{err}^{(i)}(f^{(i)}, f^{(i-1)}, x^{(i-1)}) \triangleq \frac{|A_{\text{err}}^{(i)}|}{|L^{(i)}|}$$

- **closest codeword** Let $\bar{f}^{(i)}$ denote the $\text{RS}^{(i)}$ -codeword that is closest to $f^{(i)}$ in the $\Delta^{(i)}(\cdot)$ -measure, breaking ties arbitrarily. Let $\mathcal{S}_B^{(i)} \subset \mathcal{S}^{(i)}$ denote the set of “bad” cosets on which $f^{(i)}$ and $\bar{f}^{(i)}$ disagree,

$$\mathcal{S}_B^{(i)} = \left\{ S \in \mathcal{S}^{(i)} \mid f^{(i)}|_S \neq \bar{f}^{(i)}|_S \right\}. \quad (17)$$

Let $D^{(i)} = \bigcup_{S \in \mathcal{S}_B^{(i)}} S$ denote the subset of $L^{(i)}$ of elements that belong to some “bad” coset.

Notice that $\delta^{(i)} < (1 - \rho)/2$ implies uniqueness of $\bar{f}^{(i)}$, $\mathcal{S}_B^{(i)}$ and $D^{(i)}$.

- **distortion set** For $\epsilon > 0$ the *distortion set* of $f^{(i)}$ is

$$B[f^{(i)}; \epsilon] \triangleq \left\{ x^{(i)} \in \mathbb{F} \mid \Delta_{\text{H}}(f_{f^{(i)}, x^{(i)}}^{(i+1)}, \text{RS}^{(i+1)}) < \epsilon \right\} \quad (18)$$

Notice the use of the Hamming distance measure above.

4.2.2 Proof of soundness

The following pair of lemmas will be needed to complete the analysis of soundness.

Lemma 4.3 (Soundness above unique decoding radius). *For any $\epsilon \geq \frac{2^\eta}{|\mathbb{F}|}$ and $f^{(i)}$ such that $\delta^{(i)} > 0$*

$$\Pr_{x^{(i)} \in \mathbb{F}} \left[x^{(i)} \in B \left[f^{(i)}; \frac{1}{2} \cdot \left(\delta^{(i)}(1 - \epsilon) - \rho \right) \right] \right] \leq \frac{2^\eta}{\epsilon |\mathbb{F}|} \quad (19)$$

Lemma 4.4 (Soundness within unique decoding radius). *If $\delta^{(i)} < (1 - \rho)/2$ then*

$$\Pr_{x^{(i)} \in \mathbb{F}} \left[x^{(i)} \in B \left[f^{(i)}; \delta^{(i)} \right] \right] \leq \frac{|L^{(i)}|}{|\mathbb{F}|}. \quad (20)$$

Moreover, suppose that for $i < r$ the sequences $\vec{f} = (f^{(i)}, \dots, f^{(r)})$ and $\vec{x} = (x^{(i)}, \dots, x^{(r-1)})$ satisfy

1. for all $j \in \{i, \dots, r\}$ we have $\delta^{(j)} < \frac{1-\rho}{2}$
2. for all $j \in \{i, \dots, r-1\}$ we have $\vec{f}^{(j+1)} = f_{\vec{f}^{(j)}, x^{(j)}}^{(j+1)}$
3. for all $j \in \{i, \dots, r-1\}$ we have $x^{(j)} \notin B[f^{(j)}; \delta^{(j)}]$

then

$$\Pr_{s^{(i)} \in D^{(i)}} \left[\text{QUERY}(\vec{f}, \vec{x}) = \text{reject} \right] = 1 \quad (21)$$

and consequently

$$\Pr_{s^{(i)} \in L^{(i)}} \left[\text{QUERY}(\vec{f}, \vec{x}) = \text{reject} \right] \geq \frac{|D^{(i)}|}{|L^{(i)}|} = \delta^{(i)} \quad (22)$$

We are ready to prove the soundness of the protocol, in three steps. First, we define a sequence of “bad” events $E^{(0)}, \dots, E^{(r-1)}$ that may occur (only) during the COMMIT phase. Second, we bound from above the probability that some bad event occurs by $\frac{3|L^{(0)}|}{|\mathbb{F}|}$, as stated in (8). Third and last, assuming no bad event occurs, we bound from below the probability of the verifier rejecting during the QUERY phase, proving this rejection probability is at least as stated in (9). Details follow.

Proof of Theorem 3.3, item 2 (soundness). Set $\epsilon = \frac{2^\eta}{|L^{(r/2)}|}$; for simplicity we assume r is even (using $\epsilon = \frac{2^\eta}{|L^{(\lceil r/2 \rceil)}|}$ gives the same bounds but with a slightly messier analysis).

Part I — A sequence of bad events The i th bad event $E^{(i)}$ is defined thus:

- **large distance:** If $\delta^{(i)} \geq \frac{1-\rho}{2}$ then $E^{(i)}$ is the event

$$x^{(i)} \in B \left[f^{(i)}; \frac{1}{2} \cdot \left(\delta^{(i)}(1 - \epsilon) - \rho \right) \right]$$

- **small distance:** If $\delta^{(i)} < \frac{1-\rho}{2}$ then $E^{(i)}$ is the event

$$x^{(i)} \in B \left[f^{(i)}, \delta^{(i)} \right]$$

Assuming that event $E^{(i)}$ does not hold implies that for $\delta^{(i)} < \frac{1-\rho}{2}$,

$$\Delta^{(i+1)} \left(f_{f^{(i)}, x^{(i)}}^{(i+1)}, \text{RS}^{(i+1)} \right) \geq \delta^{(i)} \quad (23)$$

and for $\delta^{(i)} \geq \frac{1-\rho}{2}$,

$$\Delta^{(i+1)} \left(f_{f^{(i)}, x^{(i)}}^{(i+1)}, \text{RS}^{(i+1)} \right) \geq \frac{1}{2} \cdot \left(\delta^{(i)}(1-\epsilon) - \rho \right) \geq \frac{(1-\rho)(1-\epsilon)}{4} - \frac{\rho}{2} \geq \frac{1-3\rho-\epsilon}{4}. \quad (24)$$

We use δ_0 to denote the rightmost term of (24) and summarize this by saying that when no $E^{(i)}$ holds we have

$$\Delta^{(i+1)} \left(f_{f^{(i)}, x^{(i)}}^{(i+1)}, \text{RS}^{(i+1)} \right) \geq \min \left\{ \delta^{(i)}, \delta_0 \right\} \quad (25)$$

Part II — bounding the probability of a bad event occurring By Lemmas 4.3 and 4.4, and by our choice of ϵ we have

$$\Pr \left[E^{(i)} \right] \leq \max \left\{ \frac{2^\eta}{\epsilon^{|\mathbb{F}|}}, \frac{|L^{(i)}|}{|\mathbb{F}|} \right\} \leq \begin{cases} \frac{|L^{(i)}|}{|\mathbb{F}|} & i \leq r/2 \\ \frac{|L^{(r/2)}|}{|\mathbb{F}|} & i > r/2 \end{cases}$$

so the probability that none of $E^{(0)}, \dots, E^{(r-1)}$ hold is at least

$$\begin{aligned} \Pr \left[\bigwedge_{i=1}^{r-1} \neg E^{(i)} \right] &\geq 1 - \left(\sum_{i \leq r/2} \frac{|L^{(i)}|}{|\mathbb{F}|} + \frac{r}{2} \frac{|L^{(r/2)}|}{|\mathbb{F}|} \right) \\ &\geq 1 - \left(|L^{(0)}| \sum_{i=0}^{r/2} \left(\frac{1}{2^\eta} \right)^i + \left(\frac{\log |L^{(0)}| - \mathcal{R}}{\eta} \cdot \frac{|L^{(0)}|}{2^{\eta \frac{\log |L^{(0)}| - \mathcal{R}}{2\eta}}} \cdot \frac{1}{|\mathbb{F}|} \right) \right) \\ &\geq 1 - \frac{1}{|\mathbb{F}|} \left(2|L^{(0)}| + \log(\rho|L^{(0)}|) \cdot \sqrt{|L^{(0)}|/\rho} \right) > 1 - 3 \frac{|L^{(0)}|}{|\mathbb{F}|}. \end{aligned}$$

The second inequality above uses the definition of r , and the third uses the geometric sum formula; the last inequality holds because $\log(\rho|L^{(0)}|) \cdot \sqrt{|L^{(0)}|/\rho} < |L^{(0)}|$ when $\rho|L^{(0)}| > 16$. We continue with the proof assuming no $E^{(i)}$ holds.

Part III — bounding soundness when no bad event occurred There are two cases to consider. The first and simpler case is when the sequences $\vec{f} = (f^{(0)}, \dots, f^{(r)})$ and $\vec{x} = (x^{(0)}, \dots, x^{(r-1)})$ satisfy the three assumptions of Lemma 4.4; in this case Lemma 4.4 immediately gives the desired lower bound of $\delta^{(0)}$ on rejection probability.

The other case is when the sequences \vec{f} and \vec{x} do not satisfy all conditions of Lemma 4.4. It cannot be the case that both assumptions 1 and 2 hold while assumption 3 fails, because that would imply that some event $E^{(i)}$ holds, contradicting our earlier assumption. Thus, it must be the case that either assumption 1 or 2 of that lemma fails to hold for the sequences \vec{f}, \vec{x} , so there exists some $i \in \{0, \dots, r-1\}$ for which either

1. $\delta^{(i)} \geq \frac{1-\rho}{2}$, or

2. $\delta^{(i)} < \frac{1-\rho}{2}$ and $\bar{f}^{(i+1)} \neq f_{\bar{f}^{(i)},x^{(i)}}^{(i+1)}$

Abusing notation, let $i < r$ be the largest integer satisfying either of the above conditions. Notice $D^{(i+1)}$ is uniquely defined because $\delta^{(i+1)} < \frac{1-\rho}{2}$ and hence $\bar{f}^{(i+1)}$ is unique. The following claim says that the honest prover's $(i+1)$ message is at least δ_0 from $\bar{f}^{(i+1)}$ in relative Hamming distance.

Claim 4.5.

$$\Delta_{\text{H}}\left(\bar{f}^{(i+1)}, f_{\bar{f}^{(i)},x^{(i)}}^{(i+1)}\right) \geq \delta_0$$

Proof. If $\delta^{(i)} \geq \frac{1-\rho}{2}$ then the assumption that $E^{(i)}$ doesn't occur means that $\Delta_{\text{H}}\left(f_{f^{(i)},x^{(i)}}^{(i+1)}, \text{RS}^{(i+1)}\right) \geq \delta_0$ and the claim clearly holds. Otherwise we are in the case that both $\delta^{(i)} < \frac{1-\rho}{2}$ and $\bar{f}^{(i+1)} \neq f_{\bar{f}^{(i)},x^{(i)}}^{(i+1)}$ hold. To simplify exposition denote $f_{\bar{f}^{(i)},x^{(i)}}^{(i+1)}$ by g . Both $\bar{f}^{(i+1)}$ and g belong to $\text{RS}^{(i+1)}$, hence are at least $(1-\rho)$ -far from each other. The triangle inequality gives

$$1 - \rho \leq \Delta_{\text{H}}\left(\bar{f}^{(i+1)}, g\right) \leq \Delta_{\text{H}}\left(\bar{f}^{(i+1)}, f_{f^{(i)},x^{(i)}}^{(i+1)}\right) + \Delta_{\text{H}}\left(f_{f^{(i)},x^{(i)}}^{(i+1)}, g\right) \quad (26)$$

By the assumption $\delta^{(i)} < \frac{1-\rho}{2}$ we have $\Delta_{\text{H}}\left(f_{f^{(i)},x^{(i)}}^{(i+1)}, g\right) < \frac{1-\rho}{2}$. Rearranging (26) gives

$$\Delta_{\text{H}}\left(\bar{f}^{(i+1)}, f_{f^{(i)},x^{(i)}}^{(i+1)}\right) \geq \Delta_{\text{H}}\left(\bar{f}^{(i+1)}, g\right) - \Delta_{\text{H}}\left(f_{f^{(i)},x^{(i)}}^{(i+1)}, g\right) > (1-\rho) - \frac{1-\rho}{2} > \frac{1-\rho}{2} > \delta_0.$$

This completes the proof. \square

Our next claim is

Claim 4.6.

$$\frac{|A_{\text{err}}^{(i+1)} \cup D^{(i+1)}|}{|L^{(i+1)}|} \geq \Delta_{\text{H}}\left(\bar{f}^{(i+1)}, f_{f^{(i)},x^{(i)}}^{(i+1)}\right). \quad (27)$$

Proof. For all $x \notin A_{\text{err}}^{(i+1)} \cup D^{(i+1)}$ we have

$$\bar{f}^{(i+1)}(x) = f^{(i+1)}(x) = f_{f^{(i)},x^{(i)}}^{(i+1)}(x)$$

because the first equality holds for $x \notin D^{(i+1)}$ and the second for $x \notin A_{\text{err}}^{(i+1)}$. But

$$\Pr_{x \in L^{(i+1)}} \left[\bar{f}^{(i+1)}(x) \neq f_{f^{(i)},x^{(i)}}^{(i+1)}(x) \right] = \Delta_{\text{H}}\left(f_{f^{(i)},x^{(i)}}^{(i+1)}, \bar{f}^{(i+1)}\right)$$

so the claim holds. \square

Combining Claims 4.5 and 4.6 gives

$$\frac{|A_{\text{err}}^{(i+1)} \cup D^{(i+1)}|}{|L^{(i+1)}|} \geq \delta_0.$$

Consider $s^{(i+1)}$ used during the QUERY phase. If $s^{(i+1)} \in A_{\text{err}}^{(i+1)}$ then the QUERY test rejects by definition. If $i+1 = r$ then $D^{(i+1)} = \emptyset$ by definition because $f^{(r)} \in \text{RS}^{(r)}$ so in this case we have already shown the rejection probability is at least δ_0 . Otherwise we are in the case that $i+1 < r$ and

by choice of i , the sequences $(f^{(i+1)}, \dots, f^{(r)})$ and $(x^{(i+1)}, \dots, x^{(r-1)})$ which are both non-empty, satisfy all three assumptions of Lemma 4.4. By the conclusion of that lemma, if $s^{(i+1)} \in D^{(i+1)}$ then the QUERY phase rejects, cf. (21). We have shown that the probability of error is at least the probability that $s^{(i+1)}$ belongs to $A_{\text{err}}^{(i+1)} \cup D^{(i+1)}$ and this probability is at least δ_0 , completing the proof of soundness. \square

4.2.3 Unique decoding radius — Proof of Lemma 4.4

Proof of Lemma 4.4. Recall $\bar{f}^{(i)}$ and $\mathcal{S}_B(f^{(i)})$ are uniquely defined because $\delta^{(i)} < \frac{1-\rho}{2}$. For a bad coset $S \in \mathcal{S}_B(f^{(i)})$ let

$$X_S^{(i)} \triangleq \left\{ x^{(i)} \in \mathbb{F} \mid \text{interpolant}^{f^{(i)}|_S}(x^{(i)}) = \text{interpolant}^{\bar{f}^{(i)}|_S}(x^{(i)}) \right\} \quad (28)$$

be the set of “misleading” values of $x^{(i)}$ on which $\text{interpolant}^{f^{(i)}|_S}$ and $\text{interpolant}^{\bar{f}^{(i)}|_S}$ agree, even though the two are distinct low-degree polynomials. We claim that

$$B[f^{(i)}, \delta^{(i)}] = \bigcup_{S \in \mathcal{S}_B(f^{(i)})} X_S^{(i)} \quad (29)$$

Indeed, by Lemma 4.1 we conclude that $f_{\bar{f}^{(i)}, x^{(i)}}^{(i+1)} \in \text{RS}^{(i+1)}$. For all $S \notin \mathcal{S}_B(f^{(i)})$ with $y_S = q^{(i)}(S)$ we have $f_{f^{(i)}, x^{(i)}}^{(i+1)}(y_S) = f_{\bar{f}^{(i)}, x^{(i)}}^{(i+1)}(y_S)$. Since $\delta^{(i)}$ is smaller than the unique decoding radius it follows that $f_{\bar{f}^{(i)}, x^{(i)}}^{(i+1)}$ is the $\text{RS}^{(i+1)}$ -codeword closest to $f_{f^{(i)}, x^{(i)}}^{(i+1)}$ in Hamming distance. Therefore, $\Delta_{\text{H}}(f_{f^{(i)}, x^{(i)}}^{(i+1)}, \text{RS}^{(i+1)}) = \Delta_{\text{H}}(f_{f^{(i)}, x^{(i)}}^{(i+1)}, f_{\bar{f}^{(i)}, x^{(i)}}^{(i+1)})$ and the two functions agree on y_S if and only if either $S \notin \mathcal{S}_B(f^{(i)})$ or $S \in \mathcal{S}_B(f^{(i)})$ and $x^{(i)} \in X_S^{(i)}$. This shows that $f_{f^{(i)}, x^{(i)}}^{(i+1)}$ disagrees with the (unique) closest $\text{RS}^{(i+1)}$ -codeword, which is $f_{\bar{f}^{(i)}, x^{(i)}}^{(i+1)}$, on all of $\{y_S \mid S \in \mathcal{S}_B(f^{(i)})\}$ if and only if $x^{(i)} \notin \bigcup_{S \in \mathcal{S}_B(f^{(i)})} X_S^{(i)}$, proving (29).

With this equality in hand, we bound the right hand side of (29). Indeed, $\text{interpolant}^{f^{(i)}|_S}$ and $\text{interpolant}^{\bar{f}^{(i)}|_S}$ are distinct polynomials of degree less than $|S|$, so $|X_S| < |S|$ and hence

$$\left| B[f^{(i)}, \delta^{(i)}] \right| = \left| \bigcup_{S \in \mathcal{S}_B(f^{(i)})} X_S \right| < |S| \cdot |\mathcal{S}_B(f^{(i)})| \leq |L^{(i)}|,$$

and this proves (20).

Next, consider the sequences \vec{f}, \vec{x} assumed in the Lemma. Assume for simplicity that $\bar{f}^{(i)}$ is the zero function evaluated over $L^{(i)}$, by subtracting $\bar{f}^{(i)}$ from $f^{(i)}$ if this is not the case; denote by $\mathbf{0}|_{L^{(i)}}$ this function. Then

$$f_{\bar{f}^{(i)}, x^{(i)}}^{(i+1)} = f_{\mathbf{0}|_{L^{(i)}}, x^{(i)}}^{(i+1)} = \mathbf{0}|_{L^{(i+1)}}$$

so by assumption 2 of Lemma 4.4 we have $\bar{f}^{(i+1)} = \mathbf{0}|_{L^{(i+1)}}$ and similarly by induction we have $\bar{f}^{(j)} = \mathbf{0}|_{L^{(j)}}$ for all $j \in \{i, \dots, r\}$. In particular, $f^{(r)} = \mathbf{0}|_{L^{(r)}}$.

Consider the sequence $(s^{(i)}, \dots, s^{(r)})$ defined in the QUERY phase, where $s^{(i)} \in D^{(i)}$. Let j denote the largest integer such that $s^{(j)} \in D^{(j)}$. This j is well defined because $s^{(i)} \in D^{(i)}$. Notice

$j < r$ because by definition $f^{(r)} = \mathbf{0}|_{L^{(r)}}$ so $D^{(r)} = \emptyset$. Assumption 3 of Lemma 4.4 together with (29) implies that $x^{(j)} \notin \cup_{S \in \mathcal{S}^{(j)}} X_S^{(j)}$, so $f_{f^{(j)}, x^{(j)}}^{(j+1)}(s^{(j+1)}) \neq 0$. But by definition of j we have $f^{(j+1)}(s^{(j+1)}) = \bar{f}^{(i+1)}(s^{(j+1)}) = 0$. We conclude

$$f_{f^{(j)}, x^{(j)}}^{(j+1)}(s^{(j+1)}) \neq f^{(j+1)}(s^{(j+1)})$$

which means that QUERY rejects the sequence $(s^{(i)}, \dots, s^{(r)})$. This proves (21), and thus implies (22) because $\delta^{(i)} = |D^{(i)}|/|L^{(i)}|$. \square

4.2.4 Beyond unique decoding radius — Proof of Lemma 4.3

To prove Lemma 4.3 we need the following improved version of Lemma 4.2.18 from [Spi95]. See Appendix A for a proof sketch.

Lemma 4.7. *Let $E(X, Y)$ be a polynomial of degree $(\alpha m, \delta n)$ and $P(X, Y)$ a polynomial of degree $((\alpha + \epsilon)m, (\delta + \rho)n)$. If there exist distinct x_1, \dots, x_m such that $E(x_i, Y)|P(x_i, Y)$ and y_1, \dots, y_n such that $E(X, y_i)|P(X, y_i)$ and*

$$1 > \max \left\{ \delta + \rho, 2\alpha + \epsilon + \frac{\rho}{\delta} \right\} \quad (30)$$

then $E(X, Y)|P(X, Y)$.

Proof of Lemma 4.3. We shall prove the contrapositive, namely, if for some $\epsilon \geq \frac{2^\eta}{|\mathbb{F}|}$

$$\frac{|B[f^{(i)}; \frac{1}{2} \cdot (\delta(1 - \epsilon) - \rho)]|}{|\mathbb{F}|} > \frac{2^\eta}{\epsilon|\mathbb{F}|} \quad (31)$$

then

$$\delta^{(i)} = \Delta^{(i)}(f^{(i)}, \text{RS}^{(i)}) < \delta. \quad (32)$$

We fix a few constants: Let $n = |L^{(i+1)}|$, $\alpha = \frac{1}{2}(1 - \epsilon - \frac{\rho}{\delta})$, $\delta' = \delta\alpha$, $B = B[f^{(i)}; \delta']$, and $m = |B|$. By definition, for every $x \in B$ we have $\Delta_{\text{H}}(f_{f^{(i)}, x}^{(i+1)}, \text{RS}^{(i+1)}) < \delta'$. Recall $\bar{f}_{f^{(i)}, x}^{(i+1)} \in \text{RS}^{(i+1)}$ is the codeword closest to $f_{f^{(i)}, x}^{(i+1)}$, breaking ties arbitrarily.

Let $C(X, Y)$ be the polynomial with $\deg_X(C) < m$, $\deg_Y(C) < \rho n$ that agrees with $\bar{f}_{f^{(i)}, x}^{(i+1)}$ for each $x \in B$; this polynomial exists because, by definition, $\bar{f}_{f^{(i)}, x}^{(i+1)}$ is an evaluation of a polynomial of degree less than ρn . Let $Q^{(i)}$ be the polynomial corresponding to $f^{(i)}$ from Claim 4.2 as defined in (12) and recall from item 2 of that claim that $\deg_X(Q^{(i)}) < |L_0^{(i)}|$; By definition $|L_0^{(i)}| = 2^\eta$ and by assumption above $2^\eta < \epsilon m$, so $\deg_X(Q^{(i)}) < \epsilon m$. From item 1 of Claim 4.2 we deduce that for all $x \in \mathbb{F}$ and $y \in L^{(i+1)}$ we have $Q^{(i)}(x, y) = f_{f^{(i)}, x}^{(i+1)}(y)$. By Equations (18) and (31),

$$\Pr_{x \in B, y \in L^{(i+1)}} [C(x, y) \neq Q^{(i)}(x, y)] < \delta'. \quad (33)$$

By construction $\alpha\delta \geq \delta'$, so there exists a non-zero polynomial

$$E(X, Y), \quad \deg_X(E) \leq \alpha m, \deg_Y(E) \leq \delta n$$

that vanishes on all points (x, y) where $x \in B, y \in L^{(i+1)}$ and $C(x, y) \neq Q^{(i)}(x, y)$. The polynomial E is known as the *error locator polynomial* [Sud92] because its zeros cover the set of error locations, where Q deviates from being a low-degree polynomial.

Since $\deg_Y(C) < \rho|L^{(i+1)}|$ and $\deg_X(Q^{(i)}) < 2^\eta < \epsilon m$, by [Spi95, Chapter 4] there exists a polynomial $P(X, Y)$ satisfying

$$\deg_X(P) < (\epsilon + \alpha)m \text{ and } \deg_Y(P) < (\rho + \delta)n \quad (34)$$

such that

$$\forall x \in B, y \in L^{(i+1)} P(x, y) = C(x, y) \cdot E(x, y) = Q^{(i)}(x, y) \cdot E(x, y) \quad (35)$$

Let $\alpha \triangleq \frac{\deg_X(P)}{m} - \epsilon$ and $\rho' = \frac{\deg_Y(P)}{n} - \delta$; Equation (34) implies $\alpha' < \alpha$ and $\rho' < \rho$.

We conclude from (34), (35) that for every row $y \in L^{(i+1)}$ we have $E(X, y)|P(X, y)$ and similarly for every column $x \in B$ we have $E(x, Y)|P(x, Y)$. By (7) we have $\delta + \rho < 1$. By Equation (34) and the definition of α we also have $2\alpha' + \epsilon + \rho'/\delta < 1$. So the assumption (30) of Lemma 4.7 holds. By the conclusion of that lemma $E(X, Y)|P(X, Y)$ as polynomials in the ring $\mathbb{F}[X, Y]$. Let $Q \equiv P/E$. We conclude Q agrees with $Q^{(i)}$ on every row $y \in L^{(i+1)}$ such that $E(X, y)$ is non-zero. By the bound on $\deg_Y(E)$, the fraction of such rows is at least $1 - \delta$. In other words $f^{(i)}$ agrees with some polynomial of degree $\rho|L^{(i)}|$ on more than a $(1 - \delta)$ -fraction of cosets of $L_0^{(i)}$ in $L^{(i)}$, implying (32) and completing the proof of the Lemma. \square

4.3 Prover complexity — Part 3

Consider the computation performed by the prover during the i th step of the protocol. At this point the prover has already committed to $f^{(i)} : L^{(i)} \rightarrow \mathbb{F}$ and has received $x^{(i)} \in \mathbb{F}$ and needs to compute $f^{(i+1)} : L^{(i+1)} \rightarrow \mathbb{F}$ as explained in steps 2–4 of the COMMIT protocol.

During step 2, for each distinct coset $S_y \in \mathcal{S}^{(i)}$, the prover needs to interpolate the polynomial $P_y(X)$ defined there and evaluate it on $x^{(i)}$ to obtain $f^{(i+1)}(y) = P_y(x^{(i)})$. If $x^{(i)} \in S_y$ then $f^{(i+1)}(y) = f^{(i)}(x^{(i)})$ and the computation terminates in single step. Otherwise, using Lagrange interpolation,

$$P_y(x^{(i)}) = \sum_{\alpha \in S_y} f^{(i)}(\alpha) \cdot \frac{\prod_{\beta \in S_y} (\beta - x^{(i)})}{\prod_{\beta \in S_y \setminus \{\alpha\}} (\beta - \alpha)} = \text{Zero}_{S_y}(x^{(i)}) \cdot \sum_{\alpha \in S_y} \frac{f^{(i)}(\alpha)}{c_\alpha \cdot (x^{(i)} - \alpha)} \quad (36)$$

where $c_\alpha = \prod_{\gamma \in S_y \setminus \{\alpha\}} (\gamma - \alpha)$ can be precomputed in advance on a PRAM with sufficiently many processors because c_α does not depend on $f^{(i)}, x^{(i)}$. The polynomial Zero_{S_y} is linearized and has $\eta + 1$ terms, hence can be evaluated on $x^{(i)}$ via repeated squaring in $3\eta + 3$ cycles in the PRAM-CREW model (each cycle is a single arithmetic operation in \mathbb{F}), using η processors and a total of $3\eta + 3$ arithmetic operations: (i) $\eta + 1$ squarings to obtain the relevant powers of $x^{(i)}$, (ii) $\eta + 1$ multiplications and (iii) $\eta + 1$ additions to evaluate the polynomial once the powers of $x^{(i)}$ are known. The summation on the right hand side of (36) has 2^η terms so it can be computed separately in parallel using 2^η processors and $\eta + 2$ cycles for a total of $4 \cdot 2^{\eta+2}$ arithmetic operations. The total PRAM-CREW number of cycles is $2\eta + O(1)$ using 2^η processors and at most $4(2^\eta + \eta + 1)$ arithmetic operations.

The calculation above refers to a single $y \in L^{(i+1)}$. Summing over all such y shows that the i th step requires a total of $|L^{(i)}| + \eta$ processors, and is computed in $3\eta + 1$ cycles using a total of $3(|L^{(i)}| + \eta + 1)$ arithmetic operations. For $i = r$ the function $f^{(r)}$ — which is evaluated on $L^{(r)}, |L^{(r)}| \leq 2^\eta$ — needs to be interpolated; this can be done using 2^η processors in 3η cycles because $\deg(P^{(r)}) < |L^{(r)}| \leq 2^\eta$ (details omitted).

Summing over all r steps completes the proof of this part.

Remark 4.8 (Arithmetic complexity for smooth RS codes). *For smooth codes, prover complexity is somewhat smaller than mentioned above, because $\text{Zero}_{S_y} = X^{2^\eta} - \zeta$ for a constant ζ depending on S_y ; i.e., Zero_{S_y} has only 2 terms as opposed to $\eta + 1$ terms in the additive case. Similar savings are obtained for the verifier arithmetic complexity (discussed next) in the smooth case. Notice, however, that Theorem 1.3 sets $\eta = 2$ and hence the difference between arithmetic complexity in the additive and smooth cases is minor.*

4.4 Verifier complexity — Part 4

The unit of measurement for communication and query complexity is field elements of \mathbb{F} . During the COMMIT phase the verifier sends a total of $r \leq (k^{(0)} - \mathcal{R})/\eta$ field elements. During the QUERY phase, the verifier precomputes $q^{(i)}(s^{(i)})$ where $q^{(i)}$ is a linearized polynomial with $\eta + 1$ terms whose coefficients are precomputed because they are independent of all verifier messages $x^{(i)}$ and prover oracles $f^{(i)}$. Using the explanation above (for evaluating Zero_{S_y}), each $q^{(i)}(s^{(i)})$ evaluation costs $3(\eta + 1)$ cycles and arithmetic operations.

Having specified the query set $s^{(i)}$, the verifier now receives a total of $\ell \cdot r \cdot 2^\eta$ field elements as answers, and solves $\ell \cdot r$ interpolation and evaluation problems of the kind described in (36). Using the explanation provided in Section 4.3 we conclude that the verifier’s work can be performed on a PRAM with exclusive read and write (EREW) using $6(\eta + 1)$ cycles for each round, requiring $\ell \cdot r \cdot 2^\eta$ processors and a total of $6\ell \cdot r \cdot (2^\eta + \eta + 1) \leq 6\ell \cdot \frac{k^{(0)} - \mathcal{R}}{\eta} \cdot (2^\eta + \eta + 1)$ arithmetic operations.

4.5 Discussion of Conjecture 3.4

In this section we justify our Conjecture 3.4, by describing a scenario that could lead to rejection probability that is close to it, and conjecturing that this scenario is the “worst possible”. We use the notation from the previous Section 4.2. For $\epsilon \leq \delta^{(0)}$ set $\epsilon' = \frac{\epsilon}{r} = \frac{\epsilon\eta}{k^{(0)}}$. Consider $f^{(0)}$ that is $\delta^{(0)}$ -far from $\mathbf{0}$ and furthermore, partition the bad cosets $\mathcal{S}_B(f^{(0)})$ into $t = 1/\epsilon'$ sets $\mathcal{S}'_1, \dots, \mathcal{S}'_t$ such that each set \mathcal{S}'_i is associated with a set of roots $R_i^{(1)} \subset \mathbb{F}, |R_i^{(1)}| = 2^\eta - 1$, all these sets $R_i^{(1)}$ are distinct, and for $S \in \mathcal{S}'_i$ the set of roots of interpolant $f^{(0)}|_S$ is precisely the set $R_i^{(1)}$.

Let $R^{(1)} = \cup_{i=1}^t R_i^{(1)}$. If $x^{(0)} \in R^{(1)}$ then $f_{f^{(0)}, x^{(0)}}^{(1)}$ is $(\delta^{(0)} - \epsilon')$ -far from $\mathbf{0}$ (and from $\text{RS}^{(1)}$) in Hamming distance. By carefully selecting the sets $\mathcal{S}'_1, \dots, \mathcal{S}'_t$, one can achieve blockwise distance $(\delta^{(0)} - \epsilon')$ as well. For this $f^{(0)}$ we have

$$\Pr \left[\delta^{(1)} \leq \delta^{(0)} - \epsilon' \right] \leq \frac{2^\eta - 1}{\epsilon' |\mathbb{F}|}.$$

Assuming similar partitions of $\mathcal{S}_B^{(i)}$ and root sets $R^{(i)}$ exist for all $i < r$, a union bound gives

$$\Pr \left[\bigvee_{i < r} \left(\delta^{(i+1)} \leq \delta^{(i)} - \epsilon' \right) \right] \leq \frac{r 2^\eta}{\epsilon' |\mathbb{F}|} = \frac{r^2 \cdot 2^\eta}{\epsilon |\mathbb{F}|} = \frac{(k^{(0)})^2 \cdot 2^\eta}{\epsilon \cdot \eta^2 |\mathbb{F}|}.$$

If none of the events $\delta^{(i+1)} \leq \delta^{(i)} - \epsilon'$ hold, then we conclude $\delta^{(r-1)} > \delta^{(0)} - r \cdot \epsilon' = \delta^{(0)} - \epsilon$ and the verifier rejects during the QUERY phase with probability at least $\delta^{(0)} - \epsilon$.

References

- [ALM⁺92] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and hardness of approximation problems. In *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science*, pages 14–23, 1992.
- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998. Preliminary version in FOCS '92.
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: a new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998. Preliminary version in FOCS '92.
- [Bab85] László Babai. Trading group theory for randomness. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*, STOC '85, pages 421–429, 1985.
- [BBC⁺16] Eli Ben-Sasson, Iddo Bentov, Alessandro Chiesa, Ariel Gabizon, Daniel Genkin, Matan Hamilis, Evgenya Pergament, Michael Riabzev, Mark Silberstein, Eran Tromer, and Madars Virza. Computational integrity with a public random string from quasi-linear PCPs. *IACR Cryptology ePrint Archive*, 2016:646, 2016.
- [BBGR16a] Eli Ben-Sasson, Iddo Bentov, Ariel Gabizon, and Michael Riabzev. A security analysis of probabilistically checkable proofs. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:149, 2016.
- [BBGR16b] Eli Ben-Sasson, Iddo Bentov, Ariel Gabizon, and Michael Riabzev. A security analysis of probabilistically checkable proofs. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:149, 2016.
- [BBHR17] Eli Ben-Sasson, Iddo Bentov, Ynon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity, 2017. Unpublished manuscript.
- [BCCT12] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, pages 326–349, 2012.
- [BCF⁺16] Eli Ben-Sasson, Alessandro Chiesa, Michael A. Forbes, Ariel Gabizon, Michael Riabzev, and Nicholas Spooner. On probabilistic checking in perfect zero knowledge. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:156, 2016.
- [BCG⁺13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In *Proceedings of the 33rd Annual International Cryptology Conference*, CRYPTO '13, pages 90–108, 2013.
- [BCG⁺14] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from Bitcoin. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy*, SP '14, 2014.
- [BCG⁺16] Eli Ben-Sasson, Alessandro Chiesa, Ariel Gabizon, Michael Riabzev, and Nicholas Spooner. Short interactive oracle proofs with constant query complexity, via composition and sumcheck. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:46, 2016.

- [BCGT13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, and Eran Tromer. On the concrete efficiency of probabilistically-checkable proofs. In *Proceedings of the 45th ACM Symposium on the Theory of Computing*, STOC '13, pages 585–594, 2013.
- [BCGV16] Eli Ben-Sasson, Alessandro Chiesa, Ariel Gabizon, and Madars Virza. Quasilinear-size zero knowledge from linear-algebraic PCPs. In *Proceedings of the 13th Theory of Cryptography Conference*, TCC '16, pages 33–64, 2016.
- [BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. *Interactive Oracle Proofs*, pages 31–60. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.
- [BFL90] László Babai, Lance Fortnow, and Carsten Lund. Nondeterministic exponential time has two-prover interactive protocols. In *Proceedings of the 31st Annual Symposium on Foundations of Computer Science*, SFCS '90, pages 16–25, 1990.
- [BFLS91] László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. Checking computations in polylogarithmic time. In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, STOC '91, pages 21–32, 1991.
- [BGH⁺05] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil Vadhan. Short PCPs verifiable in polylogarithmic time. In *Proceedings of the 20th Annual IEEE Conference on Computational Complexity*, CCC '05, pages 120–134, 2005.
- [BGH⁺06] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil P. Vadhan. Robust PCPs of proximity, shorter PCPs, and applications to coding. *SIAM Journal on Computing*, 36(4):889–974, 2006.
- [BKK⁺16] Eli Ben-Sasson, Yohay Kaplan, Swastik Kopparty, Or Meir, and Henning Stichtenoth. Constant rate pcps for circuit-sat with sublinear query complexity. *J. ACM*, 63(4):32:1–32:57, 2016.
- [BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Comput. Syst. Sci.*, 47(3):549–595, 1993.
- [BS08] Eli Ben-Sasson and Madhu Sudan. Short PCPs with polylog query complexity. *SIAM Journal on Computing*, 38(2):551–607, 2008. Preliminary version appeared in STOC '05.
- [BSCG⁺] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. Tinyram architecture specification v2. 00, 2013. URL: <http://scipr-lab.org/tinyram>.
- [CMT12] Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Practical verified computation with streaming interactive proofs. In *Proceedings of the 4th Symposium on Innovations in Theoretical Computer Science*, ITCS '12, pages 90–112, 2012.
- [CT65] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of computation*, 19(90):297–301, 1965.
- [DG08] Irit Dinur and Elazar Goldenberg. Locally testing direct product in the low error range. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 613–622, 2008.
- [Din07] Irit Dinur. The PCP theorem by gap amplification. *Journal of the ACM*, 54(3):12, 2007.
- [DR04] Irit Dinur and Omer Reingold. Assignment testers: Towards a combinatorial proof of the PCP theorem. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '04, pages 155–164, 2004.
- [FS86] Amos Fiat and Adi Shamir. How to prove yourself: practical solutions to identification and signature problems. In *Proceedings of the 6th Annual International Cryptology Conference*, CRYPTO '86, pages 186–194, 1986.

- [FS95] K. Friedl and M. Sudan. Some improvements to total degree tests. In *Proceedings Third Israel Symposium on the Theory of Computing and Systems*, pages 190–198, Jan 1995.
- [GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In *Proceedings of the 32nd Annual International Conference on Theory and Application of Cryptographic Techniques*, EUROCRYPT '13, pages 626–645, 2013.
- [GKdO⁺17] Sivakanth Gopi, Swastik Kopparty, Rafael Mendes de Oliveira, Noga Ron-Zewi, and Shubhangi Saraf. Locally testable and locally correctable codes approaching the gilbert-varshamov bound. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 2073–2091, 2017.
- [GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for Muggles. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, STOC '08, pages 113–122, 2008.
- [GS00] Oded Goldreich and Shmuel Safra. A combinatorial consistency lemma with application to proving the pcp theorem. *SIAM Journal on Computing*, 29(4):1132–1154, 2000.
- [Hr as01] Johan Hr astad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.
- [HBHW17] Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox, March 2017.
- [IKW12] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. New direct-product testers and 2-query pcps. *SIAM Journal on Computing*, 41(6):1722–1768, 2012.
- [IMSX15] Yuval Ishai, Mohammad Mahmoody, Amit Sahai, and David Xiao. On zero-knowledge PCPs: Limitations, simplifications, and applications, 2015. Available at <http://www.cs.virginia.edu/~mohammad/files/papers/ZKPCPs-Full.pdf>.
- [IW14] Yuval Ishai and Mor Weiss. Probabilistically checkable proofs of proximity with zero-knowledge. In *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*, pages 121–145, 2014.
- [Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, STOC '92, pages 723–732, 1992.
- [KMRS16] Swastik Kopparty, Or Meir, Noga Ron-Zewi, and Shubhangi Saraf. High-rate locally-correctable and locally-testable codes with sub-polynomial query complexity. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 202–215, 2016.
- [KPT97] Joe Kilian, Erez Petrank, and Gábor Tardos. Probabilistically checkable proofs with zero knowledge. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, STOC '97, pages 496–505, 1997.
- [KR08] Yael Kalai and Ran Raz. Interactive PCP. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, ICALP '08, pages 536–547, 2008.
- [LFKN92] Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992.
- [LN97] Rudolf Lidl and Harald Niederreiter. *Finite Fields*. Cambridge University Press, second edition edition, 1997.
- [Max11] Gregory Maxwell. Zero knowledge contingent payment, 2011. [Online; accessed 13-October-2017].

- [Mic00a] Silvio Micali. Computationally sound proofs. *SIAM Journal on Computing*, 30(4):1253–1298, 2000. Preliminary version appeared in FOCS '94.
- [Mic00b] Silvio Micali. Computationally sound proofs. *SIAM J. Comput.*, 30(4):1253–1298, 2000.
- [Mie09] Thilo Mie. Short PCPPs verifiable in polylogarithmic time with $o(1)$ queries. *Annals of Mathematics and Artificial Intelligence*, 56:313–338, 2009.
- [MR10] Dana Moshkovitz and Ran Raz. Two-query PCP with subconstant error. *J. ACM*, 57(5):29:1–29:29, 2010.
- [Ore33] O. Ore. On a special class of polynomials. *Trans. Amer. Math. Soc.*, 35(3):559–584, 1933.
- [Ore34] O. Ore. Contributions to the theory of finite fields. *Trans. Amer. Math. Soc.*, 36(2):243–274, 1934.
- [Pec16] M. Peck. A blockchain currency that beat s bitcoin on privacy [news]. *IEEE Spectrum*, 53(12):11–13, December 2016.
- [PS94] Alexander Polishchuk and Daniel A. Spielman. Nearly-linear size holographic proofs. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, STOC '94, pages 194–203, 1994.
- [PZ03] John Proos and Christof Zalka. Shor's discrete logarithm quantum algorithm for elliptic curves. *Quantum Info. Comput.*, 3(4):317–344, July 2003.
- [Raz95] Ran Raz. A parallel repetition theorem. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing*, STOC '95, pages 447–456, 1995.
- [RRR16] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 49–62, 2016.
- [RS60] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics*, 8(2):300–304, 1960.
- [RS92] Ronitt Rubinfeld and Madhu Sudan. Self-testing polynomial functions efficiently and over rational domains. In *Proceedings of the Third Annual ACM/SIGACT-SIAM Symposium on Discrete Algorithms, 27-29 January 1992, Orlando, Florida.*, pages 23–32, 1992.
- [Sha92] Adi Shamir. $IP = PSPACE$. *Journal of the ACM*, 39(4):869–877, 1992.
- [Sho94] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, Nov 1994.
- [Spi95] Daniel A. Spielman. *Computationally Efficient Error-Correcting Codes and Holographic Proofs*. PhD thesis, MIT, 1995.
- [Spi96] Daniel A. Spielman. Linear-time encodable and decodable error-correcting codes. *IEEE Transactions on Information Theory*, 42(6):1723–1731, 1996. Preliminary version appeared in STOC '95.
- [Sud92] Madhu Sudan. *Efficient Checking of Polynomials and Proofs and the Hardness of Approximation Problems*. PhD thesis, UC Berkeley, Berkeley, CA, USA, 1992. UMI Order No. GAX93-30747.
- [Val08] Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In *Proceedings of the 5th Conference on Theory of Cryptography*, TCC'08, pages 1–18, Berlin, Heidelberg, 2008. Springer-Verlag.
- [WB15] Michael Walfish and Andrew J. Blumberg. Verifying computations without reexecuting them. *Commun. ACM*, 58(2):74–84, 2015.

A Proof of Lemma 4.7

We restate Lemma 4.7 using the same notation as in [Spi96].

Lemma A.1. *Let $E(X, Y)$ be a polynomial of degree $(\alpha m, \beta n)$ and $P(X, Y)$ a polynomial of degree $((\alpha + \delta)m, (\beta + \epsilon)n)$. If there exist distinct x_1, \dots, x_m such that $E(x_i, Y) | P(x_i, Y)$ and y_1, \dots, y_n such that $E(X, y_i) | P(X, y_i)$ and*

$$1 > \max \left\{ \beta + \epsilon, 2\alpha + \delta + \frac{\epsilon}{\beta} \right\} \quad (37)$$

then $E(X, Y) | P(X, Y)$.

The difference between the version above and the original one is that (37) is replaced in [Spi96] with

$$1 > \alpha + \beta + \delta + \epsilon \quad (38)$$

So, to prove the statement above we guide the reader through the proof in [Spi96] (pages 97–98) using the notation there and point out the modifications needed to use (37) instead of (38). Details follow.

First, we do not assume $\beta \geq \alpha$ as there. Next, the inequality

$$\alpha + \beta + \delta + \epsilon \geq \frac{\alpha m - a}{m - a} + \frac{\delta m - a}{m - a} + \frac{\beta n - b}{n - b} + \frac{\epsilon n - b}{n - b}$$

there is replaced with the two inequalities

$$\beta + \epsilon \geq \frac{\beta n - b}{n - b} + \frac{\epsilon n - b}{n - b}$$

and

$$2\alpha + \delta + \frac{\epsilon}{\beta} \geq 2\frac{\alpha m - a}{m - a} + \frac{\delta m - a}{m - a} + \frac{\epsilon n - b}{\beta n - b}$$

which hold because each term on the left hand side is less than 1, as follows from (37).

Finally, before the very last inequality on page 98 there, replace both (i) the use of (38) and (ii) the assumption $\beta \geq \alpha$, with the assumption (37) to show

$$\beta mn > mn(\beta\alpha + \beta\delta + \alpha\beta + \alpha\epsilon)$$

To see that this inequality follows from (37), simply divide both sides by βmn . Therefore Lemma A.1, the restatement of Lemma 4.7, holds as claimed.

B Improved concrete efficiency

Ben-Sasson et al. defined the *concrete efficiency threshold (CET)* as a way to formalize the “practicality” of a PCPP construction [BCGT13], and we use the generalization of this definition to the IOPP setting below [BBGR16a]. The CET measure assigns a number (or ∞) to each IOPP system for a family of codes. Lower thresholds are considered better. The CET of an IOPP takes into account both (i) the query complexity $q_{\epsilon, \delta}$ needed to reject with probability ϵ words that are δ -far from the code, and (ii) the proof overhead, i.e., the ratio of total proof length to message-length.

Concrete efficiency threshold

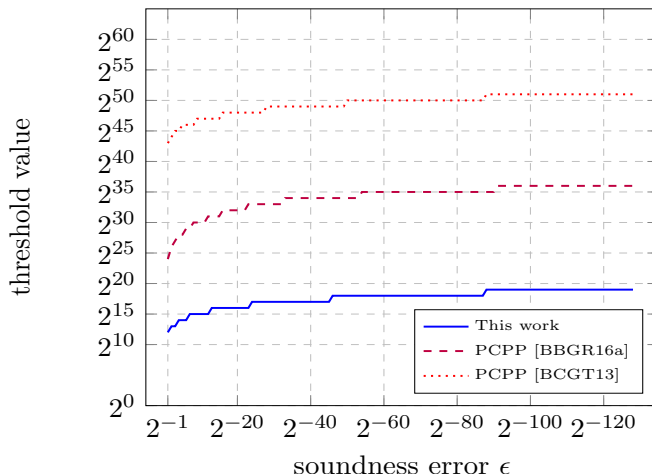


Figure 2: The concrete efficiency threshold for RS codes as a function of the soundness $\epsilon \in (2^{-1}, \dots, 2^{-128})$. We choose the same setting as [BCGT13], namely, code rate is $\rho = 1/8$ and the proximity parameter δ is a third of the code distance, i.e., $\delta = (1 - \rho)/3 = 7/24$.

Thus, to improve (i.e., decrease) the concrete efficiency threshold, one should build IOPPs that simultaneously decrease both $q_{\epsilon, \delta}$ and total proof length.

Recall that a PCPP is a 1-round IOPP, hence the following definition applies to PCPPs as a special case.

Definition B.1 (Concrete efficiency threshold). *Fix an IOPP system $S = (P, V)$ for a family of error correcting codes $\mathcal{C} = \{C_k\}$ where C_k has message-length k and block-length $N(k)$. Let $\ell(k)$ denote the IOPP proof length for C_k . Let $q_{\epsilon, \delta}(k)$ denote the minimal query complexity needed to obtain soundness error $\leq \epsilon$ for proximity parameter δ for C_k .*

A family of error correcting codes \mathcal{C} is said to have a concrete (soundness) efficiency threshold $t_{\epsilon, \delta}$ if for any code $C_k \in \mathcal{C}$, $k \geq t_{\epsilon, \delta}$ it holds that

$$q_{\epsilon, \delta}(k) \cdot \frac{N(k) + \ell(k)}{k} < k.$$

Figure 2 compares the concrete efficiency threshold of our system to prior published works on the subject [BCGT13, BBGR16a]. We vary the value of the soundness parameter ϵ , plotted on a double logarithmic scale. As seen there, the thresholds of our new system are significantly better (i.e., smaller) than the prior state of the art. We comment that for larger proximity parameters the soundness bounds conjectured earlier give even better (smaller) threshold values⁸.

⁸E.g., for the maximal value of $\delta = 1 - \rho$, the threshold derived from Conjecture 1.5 ranges between 2^{10} for soundness $\epsilon = 1/2$ to 2^{16} for soundness 2^{-128} .