Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 126 (2026)

Exponentially Fewer-Server PIR from Sparser S-Decoding Polynomials

Aparna Gupte
MIT
agupte@mit.edu

Seyoon Ragavan
MIT
sragavan@mit.edu

October 15, 2026

Abstract

We show that under a plausible number-theoretic conjecture, for any constant s there exists an s-server private information retrieval (PIR) protocol that on an n-bit database requires communication $\exp(O((\log n)^{1/s}(\log\log n)^{1-1/s}))$. Previous constructions attaining the same communication required $2^{O(s)}$ servers. Our number-theoretic conjecture is implied by existing conjectures, namely the generalized repunit conjecture and Schinzel's hypothesis H (either one of these conjectures would suffice alone).

Our result builds on the "matching vector family + S-decoding polynomials" framework pioneered by Efremenko (STOC 2009) and recently refined by Ghasemi, Kopparty, and Sudan (STOC 2025). The main ingredient is a framework for constructing S-decoding polynomials with only $k+1$ nonzero coefficients modulo special products of k primes, resolving an open problem posed by Ghasemi and Kopparty (ITCS 2026). By the lower bound shown by Ghasemi and Kopparty, this is the minimum achievable sparsity. We also empirically validate our construction and make our result unconditional for all $s \leq 15$.

We also apply our techniques to regimes where s grows with n, showing under a stronger variant of our number-theoretic conjecture that the communication complexity of s-server matching-vector PIR can be superpolynomially reduced from the previous state of the art for any $s \leq \exp(o(\sqrt{\log\log n/\log\log\log n}))$.

The main result for $s = O(1)$ and its proof were discovered in a GPT-5.5 Pro conversation prompted by the authors.

Contents

1 Introduction **3**
- 1.1 Our Contributions . . . 4
- 1.2 The Number-Theoretic Conjecture . . . 5
- 1.3 Organization of the Paper . . . 6
- 1.4 Statement on AI Use . . . 6
- 1.5 Acknowledgements . . . 7

2 Technical Overview **7**
- 2.1 Decoding Polynomials and Root-of-Unity Grids . . . 7
- 2.2 Constructing Root-of-Unity Grids . . . 9
- 2.3 PIR with a Superconstant Number of Servers . . . 10
- 2.4 Heuristic Justification for Conjecture 5.2 . . . 13

ISSN 1433-8092

3 Preliminaries **13**
3.1 Matching Vector Families . . . 13
3.2 S-Decoding Polynomials . . . 14
3.3 Previous Constructions of PIR . . . 15
3.3.1 From Matching Vectors and Decoding Polynomials . . . 15
3.3.2 From Reed-Muller Codes . . . 16
3.4 Algebraic and Number-Theoretic Background . . . 16
3.4.1 Polynomials and Finite Fields . . . 16
3.4.2 Multiplicative Characters . . . 17
3.4.3 Galois Theory . . . 18
3.4.4 Cyclotomic Polynomials . . . 18
3.4.5 Prime Density Estimates . . . 18

4 New Constructions of S-Decoding Polynomials **20**
4.1 Equivalence of S-Decoding Polynomials and Root-of-Unity Grids . . . 20
4.1.1 Root-of-Unity Grid ⇒ Sparse Decoding Polynomial . . . 20
4.1.2 Sparse Decoding Polynomial ⇒ Root-of-Unity Grid . . . 21
4.2 Approach 1: Rough Repunits Imply Root-of-Unity Grids . . . 23
4.3 Approach 2: Prime Field Norms Imply Root-of-Unity Grids . . . 26
4.4 Comparison with [CFL+13] . . . 28
4.5 Stitching Together Simultaneous Root-of-Unity Grids . . . 28

5 Number-Theoretic Conjectures for Rough Repunits **29**
5.1 Constant # of Servers Regime . . . 29
5.2 Superconstant # of Servers Regime . . . 30
5.3 Heuristics Justifying Conjecture 5.3 . . . 33
5.3.1 Naive Dickman Random Factorization Heuristics . . . 33
5.3.2 Heuristics Accounting for the Cyclotomic Structure of $q^d - 1$ . . . 34

6 PIR from Our Decoding Polynomials **36**
6.1 Small Superconstant-Server Regime . . . 37
6.2 Intermediate-Server Regime . . . 39
6.3 Many-Server Regime . . . 42

7 Concrete Empirical Results for Small s **44**

8 Conclusion **47**
8.1 Related Work . . . 47
8.2 Open Questions . . . 48

A Proof of Corollary 3.4 **51**

B Proof of Theorem 3.6 **52**

C Proof of Lemma 3.29 **53**

1 Introduction

Private information retrieval (PIR), first introduced by Chor, Kushilevitz, Goldreich, and Sudan [CKGS98], is a protocol comprising a user who holds an index $i \in [n]$ and s non-colluding servers who hold a database $\mathsf{DB} \in \{0,1\}^n$. The user sends a question to each server, then each server replies with an answer. Correctness requires that after receiving the servers' answers, the user can recover the ith entry $\mathsf{DB}_i \in \{0,1\}$. The privacy requirement is that each server should individually learn nothing about the index i; formally, the marginal distribution of any one server's question should be completely independent of the index i. Subject to these constraints, a key motivating question in the field is to find protocols with as little total communication as possible. The trivial baseline communication is n, where one of the servers sends the entire database to the user.

In the forthcoming discussion, we will regard s as constant and n as growing; our work also addresses the regime where s grows with n, but the picture there is considerably more complicated for reasons we will discuss in Section 2.3; a summary of our results for growing s can be found in Figure 2. With that aside: a straightforward PIR construction based on Reed-Muller codes requires communication $O(n^{1/(s-1)})$. Using these ideas together with recursive composition led to protocols requiring communication $O(n^{1/(2s-1)})$ [Amb97] and later $n^{O(\log\log s/(s\log s))}$ [BIKR02]. Even with these improvements, the qualitative narrative remained that PIR with $O(1)$ servers seemed to require polynomial i.e., $n^{\Omega(1)}$ communication.

The story completely changed with Yekhanin's breakthrough work [Yek08, Rag07] which showed—assuming the infinitude of Mersenne primes—a 3-server PIR protocol achieving communication $n^{O(1/\log\log n)}$. Soon after, Efremenko [Efr09] showed unconditionally that with 3 servers, one can attain significantly smaller communication $\exp(\tilde{O}(\sqrt{\log n}))$. These works laid out a powerful two-ingredient template for constructing surprisingly low-communication PIR:

- Let $S \subseteq \mathbb{Z}_m$ with $0 \in S$. An S-matching vector family is a collection of vectors $\mathbf{u}_1, \ldots, \mathbf{u}_n, \mathbf{v}_1, \ldots, \mathbf{v}_n \in \mathbb{Z}_m^d$ such that $\langle \mathbf{u}_i, \mathbf{v}_j \rangle \bmod m \in S$ always, and moreover $\langle \mathbf{u}_i, \mathbf{v}_j \rangle \equiv 0 \pmod m \Leftrightarrow i = j$. A lower dimension d ultimately leads to a lower-communication PIR.

  For m a constant prime, a simple rank argument shows that this forces $d \geq n^{\Omega(1)}$. The surprise is that for m a product of $k > 1$ distinct primes, one can achieve this property with $d = \exp(\tilde{O}((\log n)^{1/k}))$. Moreover, the corresponding set $S := S_m^*$ (defined in Definition 3.2) has size 2^k. This construction of matching vector families is due to Grolmusz [Gro00] and Barrington, Beigel, and Rudich [BBR94], and has remained the state of the art since then.

- Let $\mathbb{F}$ be a finite field and $\gamma \in \mathbb{F}$ a primitive mth root of unity. An S-decoding polynomial is a polynomial $P : \mathbb{F} \to \mathbb{F}$ such that $P(1) = 1$ and $P(\gamma^j) = 0$ for all $j \in S \backslash \{0\}$. A sparser S-decoding polynomial (i.e., one with fewer nonzero coefficients) leads to a fewer-server PIR.

  One can always achieve a baseline sparsity of $|S_m^*|$ by interpolating a polynomial of degree $|S| - 1$ with the desired properties. Combining this with the matching vector family for m a product of two distinct primes would give a 4-server PIR with communication $\exp(\tilde{O}(\sqrt{\log n}))$.

  The observation by Efremenko is that one can sometimes do better than this baseline: for $m = 7 \times 73 = 511$ and $\mathbb{F} = \mathbb{F}_{512}$, there exists an S_m^*-decoding polynomial of sparsity 3. Later, this phenomenon was generalized by [CFL+13], who showed assuming a number-theoretic conjecture that there exist choices of m that are products of an even number k of distinct primes and admit S_m^*-decoding polynomials of sparsity $3^{k/2}$.

The history of matching-vector PIR took another decisive turn with the work of Dvir and Gopi [DG16], which showed that even two servers suffice to achieve communication $2^{\tilde{O}(\sqrt{\log n})}$. Their construction used matching vector families together with ideas based on derivatives akin to [WY05], but did not leverage S-decoding polynomials. Recently, a beautiful work of Ghasemi, Kopparty, and Sudan [GKS25] (with a later alternative presentation by [ABL25]) showed that the derivative techniques of [DG16, WY05] could be combined with the sparse S-decoding polynomial approach of [Efr09, CFL+13] to give the following unified picture of all these matching vector constructions:

Theorem 1.1 ([GKS25], see Theorem 3.8 for detailed statement)**.** *Assume $k = O(1)$. Let m be a product of k distinct primes and let $\mathbb{F}$ be a finite field with a primitive mth root of unity. Assume there exists an S_m^*-decoding polynomial over $\mathbb{F}$ of sparsity $\leq s$. Then there exists an s-server PIR scheme with communication* $\exp(O((\log n)^{1/(k+1)}(\log\log n)^{1-1/(k+1)}))$.

In particular, the above theorem was the first to show that there is a 3-server PIR protocol with communication $\exp(\tilde{O}((\log n)^{1/3}))$. Recall that the above theorem always admits a trivial instantiation with $s = 2^k$. On the other hand, plugging in the S-decoding polynomials found by [CFL+13] implies that—assuming the same number-theoretic conjecture—the above theorem can be instantiated with $s = 3^{k/2}$ if k is even and $s = 2 \cdot 3^{(k-1)/2}$ otherwise. This server dependence on k is polynomially better than trivial.

A natural approach to getting even better communication-server tradeoffs for PIR is therefore to find sparser S_m^*-decoding polynomials. Recently, Ghasemi and Kopparty [GK26] showed that such a polynomial must necessarily have sparsity $\geq k + 1$, leaving an exponential gap between the best upper and lower bounds. Ghasemi and Kopparty also explicitly raised finding constructions matching this lower bound as an open question. A similar open question was raised by Beimel and Lasri [BL26] in the language of secret share conversion (which we discuss more in Section 8.1).

1.1 Our Contributions

The main result of this work is to resolve both formulations of this open question [GK26, BL26]; we show under a number-theoretic conjecture that S_m^*-decoding polynomials of sparsity $k + 1$ exist for suitable primes $p_1, \ldots, p_k$ and $m = p_1 \ldots p_k$. In doing so, we completely close the previously exponential gap between the upper and lower bounds on the sparsity of an S_m^*-decoding polynomial.

Theorem 1.2 (Sparse decoding polynomials from Conjecture 1.1)**.** *Assuming Conjecture 1.1, for any constant k there exists m that is a product of k distinct primes, a finite field $\mathbb{F}$ with $m \mid |\mathbb{F}| - 1$, and an S_m^*-decoding polynomial of sparsity $k + 1$.*

Corollary 1.3 (Secret share conversion)**.** *Under the same conjecture as in Theorem 1.2, for any constant k there exists m that is the product of k distinct primes, such that there is a $(k + 1)$-party S_m^*-share conversion scheme.*

Plugging Theorem 1.2 into Theorem 1.1 implies:

Corollary 1.4 (PIR)**.** *Assume Conjecture 1.1. Then for any constant $s \geq 2$, there exists an s-server PIR scheme with communication* $\exp(O((\log n)^{1/s}(\log\log n)^{1-1/s}))$.

In Section 7, we empirically validate our construction for small values of s, yielding the following corollary:

Communication	# of servers			
	[Efr09, CFL+13]	[DG16]	[GKS25]	This work
$\exp(\tilde{O}((\log n)^{1/2}))$	3	2	2	2
$\exp(\tilde{O}((\log n)^{1/3}))$	6	4	3	3
$\exp(\tilde{O}((\log n)^{1/4}))$	9	8	6	4
$\exp(\tilde{O}((\log n)^{1/5}))$	18	16	9	5
$\exp(\tilde{O}((\log n)^{1/6}))$	27	32	18	6
$\exp(\tilde{O}((\log n)^{1/7}))$	54	64	27	7
$\exp(\tilde{O}((\log n)^{1/8}))$	81	128	54	8
$\exp(\tilde{O}((\log n)^{1/9}))$	162	256	81	9

Table 1: For $r \in [2, 9]$, we tabulate the number of servers required for PIR with communication $\exp(\tilde{O}((\log n)^{1/r}))$. We compare our constructions with the previous works of [Efr09, CFL+13, DG16, GKS25]. For all r tabulated, our work requires r servers and [DG16] requires 2^{r-1} servers. For even r, [Efr09, CFL+13] requires $3^{r/2}$ servers while [GKS25] requires $2 \cdot 3^{(r-2)/2}$ servers. For odd r, [Efr09, CFL+13] requires $2 \cdot 3^{(r-1)/2}$ servers while [GKS25] requires $3^{(r-1)/2}$ servers.

Corollary 1.5 (PIR, unconditional)**.** *For any constant $s \in [2, 15]$, there unconditionally exists an s-server PIR with communication* $\exp(O((\log n)^{1/s}(\log\log n)^{1-1/s}))$.

Moreover, assuming the primality of 13 specific integers (listed out in Section 7) that pass probable prime tests, the above is true for $s \in [2, 24]$.

Finally, we note that our PIR results also admit the standard alternate interpretation in terms of locally-decodable codes [KT00, Tre04]:

Corollary 1.6 (Locally-decodable codes)**.** *Under the same Conjecture 1.1, for any constant s, there exists a smooth s-query locally-decodable code mapping n bits to codewords of length* $\exp(\exp(O((\log n)^{1/s}(\log\log n)^{1-1/s})))$ *over an alphabet of size* $\exp(\exp(O((\log n)^{1/s}(\log\log n)^{1-1/s})))$.

The above statements are purely asymptotic; as the statement of Corollary 1.5 suggests, our explicit constructions therein make use of probable primes $p_1, p_2, \ldots, p_k$ that have on the order of 10^6 bits and are difficult to deterministically test for primality. We tabulate the number of servers required by our construction in comparison with previous work for various communication regimes in Table 1.[1]

We next turn to a brief discussion about the number-theoretic conjecture our results rely on.

1.2 The Number-Theoretic Conjecture

Our theorems stated thus far hinge on the following conjecture:

[1] A careful reader might notice that our reported numbers for [Efr09, CFL+13] in even rows and [GKS25] in odd rows of Table 1 are better than those reported in those papers. For example, the number of servers claimed by [CFL+13] for even r is $8 \cdot 3^{(r-3)/2}$. The reason for this is that these papers made use of Theorem 3.6 due to [IS10] in the specific case where none of the m_i's are prime, as this was the variant originally stated by [IS10]. Looking more closely at the proof, this requirement is not necessary and setting one of the m_i's to be prime readily recovers the results that we claim from the constructions of [Efr09, CFL+13, GKS25] with essentially no modification. For completeness, we reprove Theorem 3.6 in this generality in Appendix B. We make no claim to originality in this observation; for example, the $r = 3$ case is briefly mentioned by [BL26].

Conjecture 1.1. *For every positive integer $k > 1$, there exists a prime $q \geq k+1$ for which there exist k distinct positive integers $d_1, \ldots, d_k \geq 6$ such that for all $i \in [k]$, $q^{d_i} - 1$ has a prime divisor $\geq (q^{d_i} - 1)^{1-1/(3k)}$.*

We note that it would suffice to ensure that (1) $d_i \geq 3k$; and (2) $(q^{d_i} - 1)/(q - 1)$ is prime. Integers of the form $(q^d - 1)/(q - 1)$ are sometimes referred to as *generalized (or base-q) repunits*, because their base q expansion consists of d 1's. Base-q repunit primes are a generalization of Mersenne primes, and their distribution is the subject of the generalized repunit conjecture [Bou09, Dub93], a generalization of the Mersenne prime conjecture. In particular, the generalized repunit conjecture states that there exist infinitely many base-q repunits for any prime q. The conjecture we need is only weaker; we only need $(q^{d_i} - 1)/(q - 1)$ to have a large prime divisor.

Finally, for reasons we will discuss in Section 2.3, Conjecture 1.1 will be insufficient for us in the regime where the number of servers is superconstant. What we roughly need is a variant of Conjecture 1.1 that also specifies how large q and the d_i's need to be (as a function of k). In particular, we would like to make use of a tighter heuristic analysis than just restricting attention to cases where $(q^{d_i} - 1)/(q - 1)$ is prime. We state this more complicated conjecture at the beginning of Section 5 and provide a heuristic justification there.

1.3 Organization of the Paper

This paper is organized as follows. In Section 2, we provide an overview of our techniques. In Section 3, we set up some notation and preliminaries. In Section 4, we prove the main technical results underlying Theorem 1.2. In Section 5, we provide additional justification for Conjecture 1.1 and also state and justify the stronger conjecture that we need for our results in the superconstant # of servers regime. In Section 6, we put these ingredients together to prove our results in the superconstant # of servers regime. In Section 7, we report empirical findings that allow us to make Theorem 1.2 unconditional for $k \leq 14$ and therefore our PIR theorem unconditional for $s \leq 15$ servers. In Section 8, we conclude with a survey of related work and some open questions.

1.4 Statement on AI Use

The main result (Theorem 1.2) and its proof (Lemmas 4.2 and 4.5) were discovered by OpenAI's GPT-5.5 Pro model in conversation with the authors. The authors first asked the LLM to conduct an empirical search for S_m^*-decoding polynomials achieving better-than-state-of-the-art sparsity for m being a product of k primes for small values of k (sparsity ≤ 5 for m a product of three primes, and sparsity ≤ 8 for m a product of four primes). After a few rounds of prompting, the model found a general framework to find many such polynomials. The model also stumbled upon a way of building decoding polynomials of sparsity ≤ 5 for m a product of four primes. On further prompting, the model generalized this framework to Lemma 4.2, its instantiation via Lemma 4.5, and therefore a proof of Theorem 1.2.

The authors then identified a number of next steps to produce a more complete picture of how far these techniques can be pushed to obtain better PIR. These steps were executed with the aid of GPT-5.5 Pro, GPT-5.6 Sol Pro, Codex 5.5, and Claude Opus 4.8. These were (a) Sections 5 and 6 (adapting these results to the case of super-constant s); (b) Lemma 4.4 (a converse of Lemma 4.2); and (c) the empirical validations of our number-theoretic conjectures detailed in Section 7.

All text in this paper was written by the authors, sometimes with the aid of draft proofs generated by GPT-5.5 Pro and Claude Opus 4.8. The authors manually verified all proofs. Claude Opus 4.8 and 5.5 were used to generate Figures 1 and 2, and for other editorial tasks.

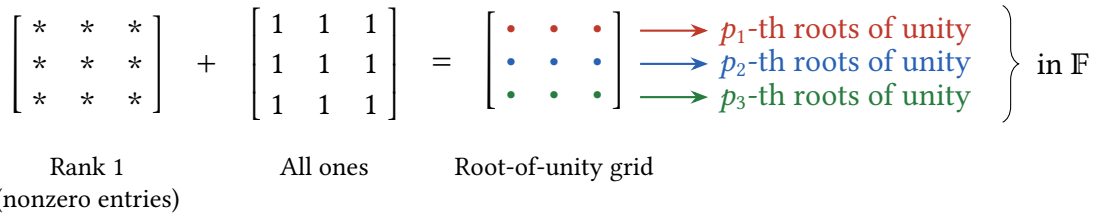


Figure 1: *Figure depicting a root-of-unity grid for a field* $\mathbb{F}$ *and modulus* $m = p_1 p_2 p_3$ *with* $m \mid |\mathbb{F}|-1$. *The entries with asterisks and bullets can be arbitrary subject to the stated conditions. The existence of a root-of-unity grid is necessary and sufficient for the existence of an* S_m^**-decoding polynomial over* $\mathbb{F}$ *with optimal sparsity* $k+1$. *A root-of-unity grid is essentially a* $k \times k$ *matrix which can be written as a sum of a rank 1 matrix with nonzero entries and the all 1's matrix, such that the* i*th row consists only of* p_i*th roots of unity over* $\mathbb{F}$.

1.5 Acknowledgements

The authors thank Vinod Vaikuntanathan for helpful suggestions and discussions throughout this project. SR also thanks Alexandra Henzinger and Edward Pyne for useful discussions about matching-vector PIR. The authors' access to ChatGPT Pro was supported by the UK AISI alignment project, and SR's access to Claude Max was supported by Jane Street.

2 Technical Overview

In Sections 2.1 and 2.2, we introduce our main technical machinery for building optimally-sparse decoding polynomials. In Section 2.3, we discuss how our results carry over to the superconstant-server regime. Finally, in Section 2.4, we briefly discuss our approach to justifying the stronger number-theoretic conjecture required for the superconstant-server setting.

We briefly introduce some relevant terminology. Let $m = p_1 \dots p_k$ be a product of k distinct primes. Then the subset $S_m^* \subseteq \mathbb{Z}_m$ consists of the 2^k elements that are congruent to either 0 or 1 modulo each of the p_i's. To handle this set, it will be convenient to define the CRT basis $c_1, \dots, c_k \in \mathbb{Z}_m$, where $c_i \equiv 1 \pmod{p_i}$ and $0 \pmod{p_j}$ for all $j \neq i$.

Let $\mathbb{F}$ be a finite field with a primitive mth root of unity γ. For a prime $p \mid |\mathbb{F}| - 1$, we will let $\mu_p(\mathbb{F})$ denote the multiplicative subgroup of pth roots of unity. Bearing this in mind, an S_m^*-decoding polynomial is a polynomial $P : \mathbb{F} \to \mathbb{F}$ such that $P(1) = 1$ and $P(\gamma^j) = 0$ for all $j \in S_m^* \backslash \{0\}$. The primary goal of this section is to show our construction (assuming a plausible number-theoretic conjecture) of $m, \mathbb{F}, P$ such that P is an S_m^*-decoding polynomial with only $k+1$ nonzero coefficients.

Before we proceed, we briefly remark that standard arguments (see Definition 3.5 and Lemma B.1) imply that (a) one need only require $P(1) \neq 0$; and (b) the existence of a decoding polynomial depends on $\mathbb{F}$ but not on the choice of γ within $\mathbb{F}$.

2.1 Decoding Polynomials and Root-of-Unity Grids

Starting point: the Ghasemi-Kopparty lower bound. Our starting point is the proof by Ghasemi and Kopparty [GK26, second proof of Theorem 3.6] that an S_m^*-decoding polynomial *cannot* have fewer than $k+1$ nonzero coefficients. Let us write the purported decoding polynomial as $P(X) = \sum_{j=1}^{\ell} \alpha_j X^{\beta_j}$ with $\ell < k+1$ nonzero coefficients. Then the idea behind the proof by [GK26] is a simple computation (which

we reproduce in Lemma 4.4) that yields elements $z_{i,j} \in \mu_{p_i}(\mathbb{F})$ for $i \in [k]$ and $j \in [\ell]$ such that:

$$\sum_{j=1}^{\ell} \alpha_j \prod_{i=1}^{k} (1 + X_i z_{i,j}^{-1}) = 1 \quad \forall X_1, \ldots, X_k \in \mathbb{F}.$$

If $\ell \leq k$, we can force the LHS to vanish with a "diagonalization" idea by setting $X_i = -z_{i,i}$ for all i, giving the desired contradiction.

Now what if $\ell = k + 1$? We can no longer set the X_i's to eliminate all $k + 1$ terms being summed on the LHS, however we can set them to eliminate any k of them. This gives some equations in the α_j's and $z_{i,j}$'s that can be combined (see Lemma 4.4) to arrive at the following necessary condition for a $(k + 1)$-sparse decoding polynomial to exist:

$$\begin{gathered}\text{There exist distinct } a_1, \ldots, a_k \in \mathbb{F}\backslash\{0\} \text{ and distinct } t_1, \ldots, t_{k+1} \in \mathbb{F} \text{ such that} \\ 1 + a_i t_j \in \mu_{p_i}(\mathbb{F}) \quad \forall i \in [k], j \in [k+1].\end{gathered}$$

We refer to this configuration as a *root-of-unity grid*. Note that we are free to set one of the t_j's to 0 to automatically satisfy all the conditions for that j. We visualize the definition of a root-of-unity grid for $k = 3$ in Figure 1.

Sufficiency of root-of-unity grids. Perhaps surprisingly, we can also show that the converse is true! A root-of-unity grid is a necessary and sufficient condition for an S_m^*-decoding polynomial of sparsity $k + 1$ to exist. Again, we skip over the relevant computations and state the high-level workings of the proof, which we present in Lemma 4.2. For bits $u_1, \ldots u_k \in \{0, 1\}$, consider the univariate polynomial:

$$f_{u_1,\ldots,u_k}(X) := \prod_{i=1}^{k} (1 + a_i X)^{1-u_i}, \tag{1}$$

and let $s \in S_m^*$ be written as $\sum_{i\in[k]} c_i u_i$. Then we can explicitly write out a polynomial P of sparsity $k + 1$ such that:

$$P(\gamma^s) = \mathsf{Coeff}_{X^k}[f_{u_1,\ldots,u_k}]. \tag{2}$$

The way this polynomial is constructed is by unpacking how Lagrange interpolation would recover $\mathsf{Coeff}_{X^k}[f_{u_1,\ldots,u_k}]$ from the evaluations of $f_{u_1,\ldots,u_k}$ at the points $t_1, \ldots, t_{k+1}$. (We will see a simpler computation with some similar features—but only justifying the trivial claim that there exists a decoding polynomial of sparsity 2^k—in the next part of this overview. At this point, we are done; by construction, $f_{u_1,\ldots,u_k}$ has degree k if and only if all u_i's are 0, which is in turn equivalent to $s = 0$.

New perspectives on known decoding polynomials. Our construction also gives new perspectives on existing constructions of decoding polynomials that we already had. For example, there is a trivial 2^k-term decoding polynomial for every $m, \mathbb{F}$ of the form $\prod_{i=1}^{k}(X^{c_i} - \gamma^{c_i})$. In the above terms, one can lay out the this decoding polynomial as follows:

- Let $c_1, \ldots, c_k$ in $\mathbb{Z}_m$ be the CRT basis i.e., $c_i \equiv 1 \pmod{p_i}$ and $\equiv 0 \pmod{p_j}$ for all $j \neq i$.
- Then for any bits $u_1, \ldots, u_k \in \{0, 1\}$, we now consider the multilinear and multivariate polynomial (compare with the univariate, degree $\leq k$ polynomial in Equation (1)):

$$f_{u_1,\ldots,u_k}(X_1, \ldots, X_k) = \prod_{i=1}^{k} (1 + (\gamma^{c_i} - 1)X_i)^{1-u_i}, \tag{3}$$

- Indeed, in this case if we let $j = \sum_{i=1}^{k} c_i u_i$, the interpolation identity stated in Lemma 3.12 can be shown to yield $\mathsf{Coeff}_{X_1 X_2 \dots X_k}[f_{u_1,\dots,u_k}] = P(\gamma^j)$, where $P(X) = (-1)^k X^{-\sum_{i\in[k]} c_i} \prod_{i=1}^{k} (X^{c_i} - \gamma^{c_i}) = (-1)^k \prod_{i=1}^{k} (1 - \gamma^{c_i} X^{-c_i}) = (-1)^k \prod_{i=1}^{k} \left(1 - \gamma^{c_i(1-u_i)}\right)$ is, up to some prefactors, exactly the trivial 2^k-term decoding polynomial.

Secondly, there are known constructions of 3-term decoders modulo products of two primes [Efr09, CFL+13]. Our results clarify that the mechanism enabling these seemingly ad hoc constructions is a root-of-unity grid: in more detail, they work with fields that have distinct nonzero elements a_1, a_2 and distinct nonzero elements t_1, t_2 such that in the matrix

$$\begin{bmatrix} 1 & 1 \\ 1 + a_1 t_1 & 1 + a_2 t_1 \\ 1 + a_1 t_2 & 1 + a_2 t_2 \end{bmatrix},$$

all entries in the left column are p_1th roots of unity and all entries in the right column are p_2th roots of unity. We will say some more about the $k = 2$ case in the next section.

2.2 Constructing Root-of-Unity Grids

So far, we have completely characterized when there exists an S^*_m-decoding polynomial over sparsity $k + 1$, through our notion of a root-of-unity grid. For what choices of $m = p_1 \cdots p_k$ and $\mathbb{F}$ does such a grid exist? Prior work [Efr09, CFL+13] gave two constructions of such grids in the $k = 2$ case:

- Efremenko [Efr09] showed directly that there is an $S^*_{7\cdot73}$-decoding polynomial over $\mathbb{F}_{2^9}$ of sparsity 3.
- [CFL+13] provided a family of constructions for moduli $m = p_1 p_2$ such that $m + 1 = 2^d$ for a prime d.

For general k, we put forth two approaches to constructing root-of-unity grids. The first is based on the probabilistic method and is our main workhorse for the theorems in this work. The second approach is more algebraic and uses some light Galois theory; although we do not use it for our main results, it has the attractive feature that it provides a more general type of construction that includes Efremenko's original $m = 7 \cdot 73$ example. Finally, we explain why neither of our approaches seems to capture the construction by [CFL+13]. We view this as clear evidence that further improvements on our results (particular for the case of superconstant s) should be possible by finding different mechanisms for building root-of-unity grids; we will discuss this more in Section 8.2.

Approach 1: the probabilistic method (Section 4.2). One strategy for understanding this question would be to assume that the characteristic q of $\mathbb{F}$ is $\geq k+1$, and fix the elements $t_1, \dots, t_{k+1}$ to be $1, \dots, k-1, k, 0$. Let $T = \{0, 1, \dots, k\}$. Now imagine selecting each a_i independently and uniformly at random over $\mathbb{F}$ (in this overview, we elide technicalities about field elements that we consider possibly being 0). What is the probability that $1 + a_i T \subseteq \mu_{p_i}(\mathbb{F})$? This approach is attractive because the criteria that each a_i needs to satisfy are independent. We reason about this probability as follows:

- For each $j \in [k]$, $1 + a_i t_j$ is roughly a uniform and random element of $\mathbb{F}$.
- Therefore, the probability that it is an element of μ_{p_i} is roughly $p_i/|\mathbb{F}|$.
- Pretending for a moment that this event is independent for each j, the probability that a_i will work for us is roughly $(p_i/\mathbb{F})^k$. This is too small to inspire hope that such an a_i exists; for this, we would want a probability $\geq 1/|\mathbb{F}|$.

This should not be surprising; it is already known that even in the $k = 2$ case, combinations of m and $\mathbb{F}$ that admit a 3-term decoding polynomial (equivalently, a root-of-unity grid) are rare [CFL+13]. To fix the above idea, we make the following tweak: for each $i \in [k]$, let d_i be minimal such that $q^{d_i} - 1$ is divisible by p_i. We pick $\mathbb{F}$ so that it contains a (unique) subfield isomorphic to $\mathbb{F}_{q^{d_i}}$, and *we will only sample a_i from this subfield.* This changes things dramatically; now the heuristic probability will be $\approx (p_i/q^{d_i})^k$, which can now be much larger. In particular, if we choose d_i such that $\frac{q^{d_i}-1}{q-1}$ is prime (such a number is known as a *generalized (or base q) repunit* because its base q representation consists of *repeated units*—d_i 1's), we can take p_i to be this prime and now the heuristic probability will be $\approx q^{-k}$. If we set all the d_i's to be $> k$, this will be $> q^{-d_i}$ and we will have succeeded.

At this point, we have heuristically argued that Conjecture 1.1 ought to imply the existence of $m, \mathbb{F}$ for each k that admit a root-of-unity grid. Fortunately, this intuition can be made rigorous using the Weil bounds on character sums. In fact, we can accommodate milder forms of the conjecture—instead of requiring $\frac{q^d-1}{q-1}$ to be prime, we are happy whenever $q^d - 1$ has a large prime divisor; we merely discuss the generalized repunit prime variant for its simplicity. We refer the reader to Section 4.2 for details.

Approach 2: kernels of the norm map (Section 4.3). There is a second, more algebraic approach to building root-of-unity grids, that has the added benefit of capturing Efremenko's original example of a 3-sparse decoding polynomial with $m = 7 \times 73$ over $\mathbb{F}_{2^9}$. This approach manufactures root-of-unity grid rows $1 + aT$ directly from the norm map $N_{\mathbb{F}_{Q^d}/\mathbb{F}_Q} : \mathbb{F}_{Q^d}^\times \to \mathbb{F}_Q^\times$ that maps x to the product of all its Galois conjugates $x^{1+Q+\ldots+Q^{d-1}}$ (see Section 3.4.3 for a review of the very light Galois theory that we use). If $p := \frac{Q^d-1}{Q-1} = 1 + Q + \ldots + Q^{d-1}$ is prime, then μ_p is exactly the kernel of the norm map $N_{\mathbb{F}_{Q^d}/\mathbb{F}_Q}$. Therefore, asking for an a such that $1 + aT \subseteq \mu_p$ for fixed $T \subseteq \mathbb{F}_Q$ is equivalent to asking for an a such that $1 + aT \subseteq \ker N_{\mathbb{F}_{Q^d}/\mathbb{F}_Q}$. This can be achieved by choosing a to be a root of a carefully designed irreducible polynomial. See Section 4.3 for more details.

Comparison with [CFL+13] (Section 4.4). The construction of [CFL+13] gives a different source of root-of-unity grids when $k = 2$ that our theorems do not subsume. We will discuss this more in Section 4.4, but for now we note that we view this as evidence that root-of-unity grids can arise from mechanisms beyond the two sufficient conditions isolated in this paper. The equivalence between $(k+1)$-sparse decoding polynomials and root-of-unity grids gives a tight structural characterization of what must be constructed, but it does not yet give a sharp arithmetic characterization of when a given pair $(m, \mathbb{F})$ admits such a grid. In particular, the construction of [CFL+13] appears to exploit additional multiplicative structure in finite fields that is not captured either by our probabilistic-method approach or by the norm-kernel approach.

2.3 PIR with a Superconstant Number of Servers

To understand how the story changes depending on whether the number of servers is constant or superconstant, let us examine the workflow laid out by [GKS25] to go from an S_m^*-decoding polynomial of sparsity s over $\mathbb{F}$ of characteristic q, where $m = p_1 p_2 \cdots p_k$:

- Construct an S_{qm}^*-matching vector family of size n over $\mathbb{Z}_{qm}$ of some dimension d. This will go via the construction of [BBR94, Gro00, DGY11].
- Then there exists a PIR scheme with s servers and total communication $O(sd \log |\mathbb{F}|)$.

We summarize this in Theorem 3.8. When the number of servers s is constant, we can also take $m, |\mathbb{F}|$ to be constant so that the communication is just $O(d)$. The matching vector construction of [Gro00] works with any product of $k+1$ primes and achieves a dimension of

$$d = \exp(O_m((\log n)^{1/(k+1)}(\log\log n)^{1-1/(k+1)})).$$

The contents of this section are succinctly captured by the phase diagram in Figure 2.

Few-server regime (Theorem 6.1 and Green Shaded Region in Figure 2). Now suppose that the number of servers s grows. This necessarily forces the number of primes k to grow, which means that m cannot be constant anymore. Consequently, the hidden dependence on m in the $O_m(\cdot)$ is actually important now; very roughly (see Corollary 3.4 for a precise formulation), this dependence hides a multiplicative factor of $m^{1/(k+1)} \geq \Omega(k \log k)$ (since m is a product of k distinct primes). This is a limitation of any PIR protocol that uses the matching vector constructions of [Gro00]; no matter how many servers we have, the communication cannot drop below $\exp((\log\log n)^{2+o(1)})$. This is in stark contrast to Reed-Muller PIR, where the communication can be pushed down to $\mathrm{polylog}(n)$ with $O(\log n/\log\log n)$ servers.

The upshot of this is that when building matching-vector PIR for a superconstant number of servers, we need to pay attention to how m grows with n. For example, the construction by [DG16] uses a decoding polynomial of trivial sparsity 2^k and does not require any structure on the k primes, so the optimal choice is to take the first k primes and their product m will grow like $(k\log k)^k$. In our setting, we can get away with sparsity $k+1$ but at the expense of requiring considerable number-theoretic structure of the primes. This forces the primes to be larger; we heuristically quantify how much larger in Conjecture 5.2 (which we will discuss in Section 2.4). Under this conjecture, our exponential reduction in the number of servers persists only up to $k \sim (\log\log n)^{1/3}/(\log\log\log n)^{2/3}$. The takeaway of this is the following:

> For regimes of communication complexity CC where
> $\log\log \mathsf{CC} \geq \Omega((\log\log n)^{2/3}(\log\log\log n)^{2/3})$, our construction (assuming Conjecture 5.2)
> needs exponentially fewer servers than constructions in prior work.

Many-server regime (Theorem 3.9 and Violet Curve in Figure 2). On the other extreme, once the number of servers is $\geq \exp(\Omega(\sqrt{\log\log n/\log\log\log n}))$, Conjecture 5.2 and our application of it in Section 6 suggest that our methods do not offer any improvement over prior constructions of matching-vector PIR by [Efr09, GKS25]. As we remark in Section 8.2, we believe that different approaches to building exponentially sparser S-decoding polynomials should be possible, which may help improve the state of affairs in this many-server regime. To summarize:

> For regimes of communication complexity CC where $\log\log \mathsf{CC} \leq O(\sqrt{\log\log n \cdot \log\log\log n})$,
> our construction does not appear to improve on prior matching-vector PIR [Efr09].

Intermediate-server regime (Theorem 6.2 and Blue Curve Between Green Shaded Region and Violet Curve in Figure 2). When the number of servers lies between $(\log\log n)^{1/3}/(\log\log\log n)^{2/3}$ and $\exp(O(\sqrt{\log\log n/\log\log\log n}))$, we can smoothly interpolate between the two extremal approaches detailed so far to find a sweet spot. To do this, we will need to compose a few different products of k primes, each of which admits a $(k+1)$-sparse decoding polynomial. This composition technique is standard, due to Itoh and Suzuki [IS10], and employed by [CFL+13] to build the sparsest-known decoding polynomials prior to our work; we restate their result in Theorem 3.6. We summarize as follows:

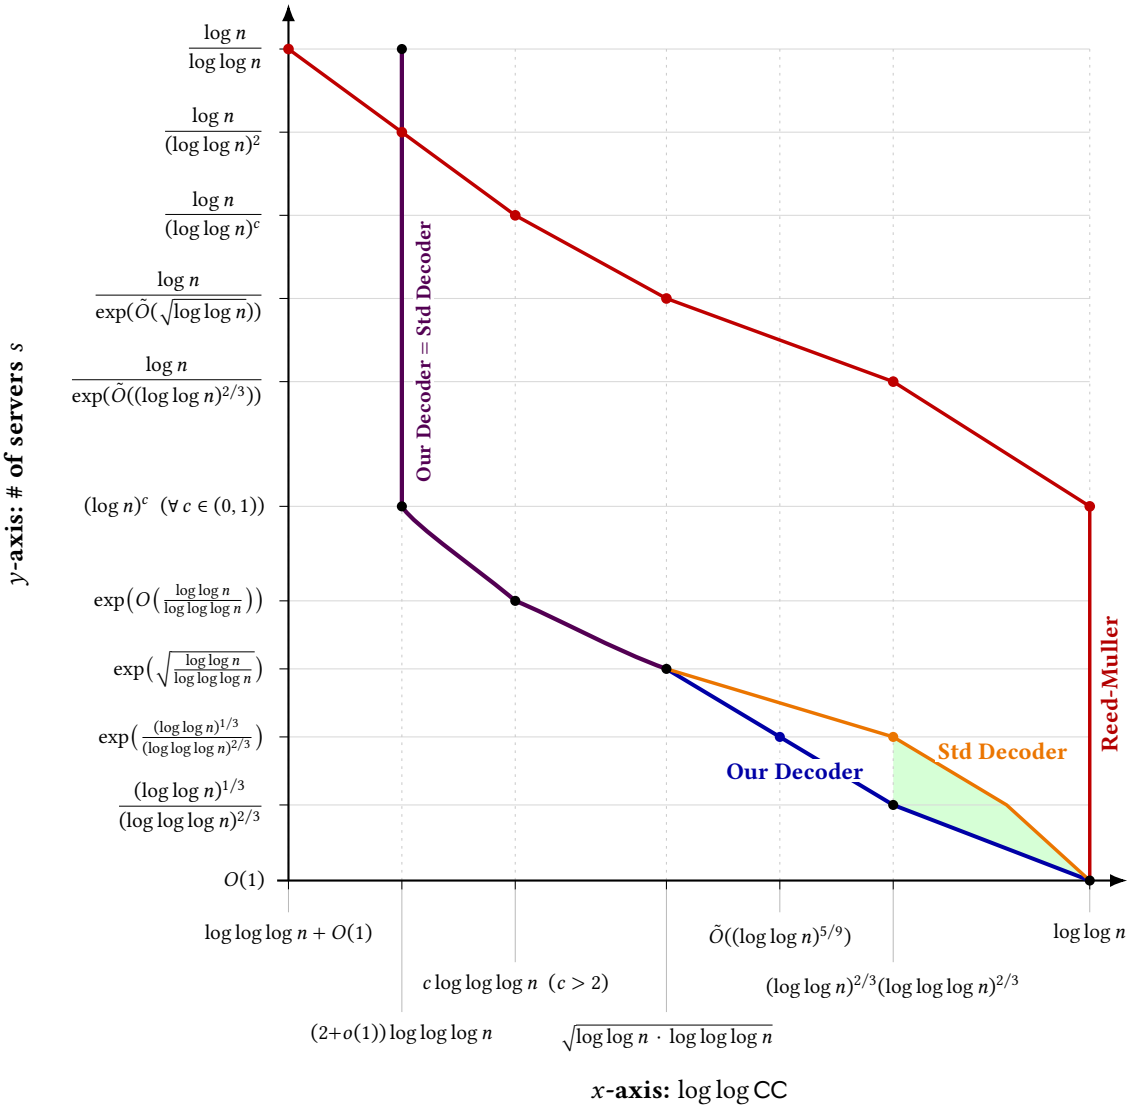


Figure 2: *Qualitative phase diagram comparing various PIR constructions across different asymptotic regimes. The y-axis is indexed by the number of servers s (as a function of the database size n). The x-axis is indexed by* $\log\log \mathsf{CC}$*, where* CC *is the communication complexity. The leftmost extreme of* $\log\log\log n + O(1)$ *corresponds to* $\mathrm{polylog}(n)$ *communication and the rightmost extreme captures* $\mathrm{poly}(n)$ *communication and also any* $2^{(\log n)^{O(1)}}$ *communication.*

Red curve: PIR based on Reed-Muller codes ([CKGS98], restated in Theorem 3.10). Orange and violet curves: PIR prior to our work based on matching-vector families and S_m^**-decoding polynomials ([Efr09], restated in Theorem 3.9). Blue and violet curves: PIR in our work—assuming Conjecture 5.2—by constructing sparser* S_m^**-decoding polynomials, stated in Theorems 6.1 (green shaded region) and 6.2 (between green region and violet curve). Green shaded region: regimes of communication complexity where we manage to reduce the server count exponentially compared to [Efr09] while preserving the communication.*

As indicated in the diagram, we obtain superpolynomial reductions in communication complexity whenever $s \le \exp(o(\sqrt{\log\log n / \log\log\log n}))$*, and exponential reductions in the server count whenever* $\log\log \mathsf{CC} \ge (\log\log n)^{2/3}(\log\log\log n)^{2/3}$.

For regimes of s, n where $s \leq \exp(o(\sqrt{\log\log n / \log\log\log n}))$, our construction (assuming Conjecture 5.2) attains superpolynomially smaller communication complexity than prior matching-vector PIR [Efr09].

2.4 Heuristic Justification for Conjecture 5.2

In the superconstant server regime, as discussed above, it turns out to be optimal to stitch many root-of-unity grids of different sizes over the same characteristic into a single decoding polynomial. Conjecture 5.2 estimates the cost $(|\mathbb{F}|, m)$ of this decoding polynomial, given the sizes of the individual root-of-unity grids we want to stitch together. Taking approach 1 and constructing root-of-unity grids from repunits with large prime factors, estimating this costs boils down to estimating the density of integers d such that $q^d - 1$ has large prime factors. That is, one can reduce Conjecture 5.2 to a much simpler conjecture about the density of such integers d (Conjecture 5.3).

The density estimate in Conjecture 5.3 is based on the Dickman heuristic that studies the probability that a random integer has a large prime factor. Treating $q^d - 1$ as a random integer, however, ignores a lot of structure—for example, if d is even, $q^d - 1 = (q^{d/2} - 1)(q^{d/2} + 1)$, and so $q^d - 1$ cannot have a prime factor $\gg q^{d/2}$. One can restrict attention to prime d to minimize such effects, but even this ignores the fact that the largest prime factors of $q^d - 1$ can come from a special arithmetic progression. We also account for this structure by heuristically estimating the density of such primes, and corroborate the original estimate given by naive Dickman heuristic.

3 Preliminaries

Notation. For a polynomial $P \in \mathbb{F}[X]$, we let $\mathsf{Coeff}_{X^d}(P) \in \mathbb{F}$ denote the coefficient of X^d in P. We use $\mathbb{Z}_m = \mathbb{Z}/m\mathbb{Z}$ to denote the integers modulo m. See Section 3.4 for more algebraic and number-theoretic notation and terminology we use.

3.1 Matching Vector Families

We next define matching-vector families and verify some useful facts about them. Let $a, p_1, p_2, \ldots, p_t$ be distinct primes. Let $m > 1$ be a positive integer, and let $\mathbb{Z}_m = \mathbb{Z}/m\mathbb{Z}$ denote the ring of integers modulo m.

Definition 3.1 (S-matching-vector families)**.** *Let d be a positive integer and $S \subseteq \mathbb{Z}_m$ some subset containing 0. We say that two collections of vectors (U, V), where $U = (\mathbf{u}_1, \ldots, \mathbf{u}_N) \in (\mathbb{Z}_m^d)^N$ and $V = (\mathbf{v}_1, \ldots, \mathbf{v}_N) \in (\mathbb{Z}_m^d)^N$, form an S-matching-vector family over $\mathbb{Z}_m^d$ of size N, if:*

1. *for every $i, j \in [N]$, it holds that $\langle \mathbf{u}_i, \mathbf{v}_j \rangle \bmod m \in S$, and*
2. *for every $i, j \in [N]$, we have $i = j$ if and only if $\langle \mathbf{u}_i, \mathbf{v}_j \rangle \equiv 0 \pmod m$.*

Definition 3.2 (Canonical set)**.** *Let $\ell \geq 1$ and suppose $p_1, \ldots, p_\ell$ are distinct primes and $m = p_1 p_2 \ldots p_\ell$. Then the canonical set of $\mathbb{Z}_m$ is defined to be the set $S_m^* \subseteq \mathbb{Z}_m$ of residues that are either 0 or 1 mod p_i for all $i \in [\ell]$. (Hence we have $|S_m^*| = 2^\ell$.)*

Theorem 3.3 ([BBR94, Gro00, DGY11])**.** *Let m be the product of $\ell > 1$ distinct primes $p_1, \ldots, p_\ell$, and let S be the canonical set of $\mathbb{Z}_m$. Let h, w be positive integers. Then there exists a S_m^*-matching-vector family over $\mathbb{Z}_m^d$ of size N, where:*

$$N := \binom{h}{w}, \text{ and}$$

$$d := \sum_{j=0}^{\lfloor (mw)^{1/\ell} \rfloor} \binom{h}{j}.$$

Proof Sketch. This is immediate from setting parameters suitably in [DGY11, Lemma 11]. For each $i \in [\ell]$, we take $e_i \geq 0$ to be minimal such that $p_i^{e_i} > (mw)^{1/\ell}/p_i$. This implies that $\prod_i p_i^{e_i} > mw/\prod_i p_i = w$,[2] and moreover for any i we have $p_i^{e_i} \leq (mw)^{1/\ell}$. □

The following corollary is straightforward; we defer its proof to Appendix A. It so happens that our main results all only use the second of the two cases stated (where $\eta \geq 2^{\Omega(\ell^2)}$); however, we also include an asymptotically optimal expression for d in the other case for completeness.

Corollary 3.4. *Let $\eta = \frac{(\log n)^{\ell-1}}{m}$ and assume $\eta \geq 1$. Then for any n, there exists a S_m^*-matching-vector family over $\mathbb{Z}_m^d$ of size $\geq n$, where:*

$$d = \begin{cases} \exp\left(O\left(\log n \cdot \frac{1+\frac{1}{\ell}\log \eta}{\eta^{1/\ell}}\right)\right), & \text{if } \eta \leq 2^{O(\ell^2)}, \\ \exp\left(O\left(\log n \cdot \frac{(\log \eta)^{1-1/\ell}}{\ell \eta^{1/\ell}}\right)\right), & \text{if } \eta \geq 2^{\Omega(\ell^2)}. \end{cases}$$

3.2 S-Decoding Polynomials

Another useful ingredient for us will be that of S-decoding polynomials, a notion introduced by [Efr09]:

Definition 3.5 (S-decoding polynomials)**.** *Let $m \geq 1$ be a positive integer, $S \subseteq \mathbb{Z}_m$ a subset containing 0, and $\mathbb{F}$ a field that contains some primitive mth root of unity γ_m. A polynomial $P \in \mathbb{F}[X]$ is called an $(\mathbb{F}, m, S)$-decoding polynomial with respect to γ_m if:*

1. *for every $j \in S \backslash \{0\}$, we have $P(\gamma_m^j) = 0$, and*
2. *$P(\gamma_m^0) = P(1) \neq 0$. (This can be strengthened to $P(1) = 1$ without loss of generality by scaling P by a suitable constant factor.)*

The sparsity of P is the number of nonzero coefficients it has after reducing the exponents modulo m. Often, $\mathbb{F}$ and m are clear from the context, and we write S-decoding polynomial.

Remark. *Strictly speaking, the definition of P as a decoding polynomial should also depend on the choice of root γ_m. However, a straightforward argument (see Lemma B.1) shows that the existence of P within a certain sparsity is independent of the particular γ_m that is chosen. So we will abuse terminology slightly and omit γ_m when referring to S-decoding polynomials.*

We begin by recalling a standard observation about trivially sparse decoding polynomials always existing [Efr09]:

Claim 3.1. *Let Q, m be positive integers with m prime, Q a prime power, and $m \mid Q-1$. (Note that $S_m^* = \{0, 1\}$.) Then there exists an $(\mathbb{F}_Q, m, S_m^*)$-decoding polynomial of sparsity 2.*

Proof. Let $\gamma \in \mathbb{F}_Q$ be any primitive mth root of unity. We can simply take $P(X) := X - \gamma$, which has sparsity 2. We have $P(\gamma^1) = 0$ and $P(\gamma^0) = P(1) = 1 - \gamma \neq 0$, as desired. □

[2] As written, [DGY11, Lemma 11] requires $p_i^{e_i} > w^{1/\ell}$ for all i, but inspecting the proof implies that it suffices to have $\prod_{i \in [\ell]} p_i^{e_i} > w$. Their proof does not explicitly handle the $e_i = 0$ case but it is straightforward to check that allowing $e_i = 0$ does not affect the proof.

The following theorem is due to [IS10]; for completeness, we reproduce its proof in Appendix B.

Theorem 3.6 (Composition theorem [IS10])**.** *Let $q, m_1, \dots, m_r$ be pairwise coprime squarefree positive integers with q prime. Assume that for every $i \in [r]$ there is an $(\mathbb{F}_{q^{c_i}}, m_i, S^*_{m_i})$-decoding polynomial of sparsity d_i for some integer $c_i \geq 1$. Then there is also an $(\mathbb{F}_{q^{\mathrm{lcm}(c_1,\dots,c_r)}}, m_1 \dots m_r, S^*_{m_1 \dots m_r})$-decoding polynomial of sparsity $d_1 \dots d_r$.*

Corollary 3.7. *Let $\ell \geq 5$ be a positive integer. Then there exists m that is a product of ℓ distinct primes and a field $\mathbb{F}$ of characteristic 2 such that there exists an S^*_m-decoding polynomial over $\mathbb{F}$ of sparsity $\leq 2^\ell$. Moreover, we have the following bounds:*

$$m \leq (\ell \log \ell)^\ell$$
$$\log |\mathbb{F}| \leq \exp\left(O\left(\ell \log \ell\right)\right).$$

Proof. We apply Theorem 3.6 with $q = 2$, $r = \ell$, and $m_1, \dots, m_\ell$ the first ℓ odd primes. For each $i \in [\ell]$, let c_i be the multiplicative order of 2 modulo m_i. Then by Claim 3.1 there exists an $(\mathbb{F}_{2^{c_i}}, m_i, S^*_{m_i})$-decoding polynomial of sparsity 2 for every $i \in [\ell]$. So Theorem 3.6 tells us that there exists an $(\mathbb{F}_{2^{\mathrm{lcm}(c_1,\dots,c_\ell)}}, m_1 \dots m_\ell, S^*_{m_1 \dots m_\ell})$-decoding polynomial of sparsity $\leq 2^\ell$.

It remains to establish the stated bounds on $m := m_1 \dots m_\ell$ and $|\mathbb{F}| = 2^{\mathrm{lcm}(c_1,\dots,c_\ell)}$. First, m is simply the product of the first ℓ odd primes, which is $\leq (\ell \log \ell)^\ell$ by Lemma 3.24. Secondly, $\mathrm{lcm}(c_1, \dots, c_\ell)$ is exactly the multiplicative order of 2 modulo m. It must therefore be at most m. Consequently, we have:

$$\begin{aligned}
\log |\mathbb{F}| &\leq O(\mathrm{lcm}(c_1, \dots, c_\ell)) \\
&\leq O(m) \\
&\leq O((\ell \log \ell)^\ell) \\
&= \exp(O(\ell \log \ell)),
\end{aligned}$$

as claimed. □

3.3 Previous Constructions of PIR

3.3.1 From Matching Vectors and Decoding Polynomials

Below we state the main theorem of [GKS25], reworded in such a way that it robustly holds even when the number of servers and the modulus m are superconstant.

Theorem 3.8 (Essentially Theorem 2.3 of [GKS25])**.** *Let m be a product of k distinct primes, let q be a prime with $\gcd(q, m) = 1$, and let $\mathbb{F}$ be a finite field of characteristic q that contains a primitive mth root of unity. Assume the following two ingredients:*

- *an S^*_{qm}-matching-vector family over $\mathbb{Z}^d_{qm}$ of size n; and*
- *an S^*_m-decoding polynomial over $\mathbb{F}$ with sparsity $\leq T$.*

Then there is a T-server PIR scheme for a database of size n with total communication $O(T d \log |\mathbb{F}|)$. In more detail, each query from the user is an element of $\mathbb{Z}^d_m$, and each answer from each server is an element of $\mathbb{F}^{d+1}$.

Proof Sketch. As written, the construction by [GKS25] has the queries live in $\mathbb{F}^d$ and the answers live in $\mathbb{F}^{d+1}$, but on closer inspection every coordinate of the users' queries is an mth root of unity so it can be implicitly represented as an element of $\mathbb{Z}_m$. □

Plugging in the standard matching-vector family construction of [BBR94, Gro00] and the trivial 2^k-sparse decoding polynomial leads to the below theorem [Efr09, DGY11]. For completeness, we re-derive this communication bound in Section 6.3.

Theorem 3.9 (Many-server matching-vector PIR [Efr09, DGY11])**.** *For any n, s such that $32 \leq s \leq \log n$, there exists an s-server PIR for a database of size n achieving total communication:*

$$\exp\left(O\left((\log n)^{\frac{1}{\lceil \log_2 s \rceil + 1}} \cdot (\log\log n)^{1 - \frac{1}{\lceil \log_2 s \rceil + 1}} \cdot \log s \cdot \log\log s\right)\right).$$

This can in turn be more coarsely bounded by:

$$\exp\left(\exp\left(O\left(\log\log\log n + \frac{\log\log n}{\log s}\right)\right)\right).$$

3.3.2 From Reed-Muller Codes

Another approach to PIR is that based on Reed-Muller codes [CKGS98, BIKR02, WY05]. We summarize the performance of Reed-Muller PIR below:

Theorem 3.10 (Reed-Muller-based PIR [CKGS98, BIKR02, WY05])**.** *For any n, s, there exists an s-server PIR protocol achieving total communication:*

$$\mathsf{CC} = \min\left(2^{\tilde{O}(s)} \cdot n^{\frac{2\log\log s}{s\log s}}, s^2 \log s \cdot n^{\frac{1}{2s-1}}\right).$$

The former expression is due to [BIKR02] and is preferable for $s \leq \sqrt{\log n} \cdot \mathrm{poly}(\log\log n)$; the latter expression is due to [CKGS98, WY05] and is preferable for $s \geq \sqrt{\log n} \cdot \mathrm{poly}(\log\log n)$.

We note that although the comparison between [BIKR02] and [WY05] is somewhat complicated, for the level of coarseness of our comparisons in Figure 2, the improvements by [BIKR02] in some regimes will not matter. We could have even used the original $\mathrm{poly}(s) \cdot n^{1/(s-1)}$ communication construction by [CKGS98] without any impact on the contents of Figure 2. Our statement of the results by [BIKR02, WY05] is only for completeness.

3.4 Algebraic and Number-Theoretic Background

Notation and terminology. For a positive integer n, we write $P^+(n)$ to denote the largest prime factor of n. For integers a, m such that $\gcd(a, m) = 1$, we define $\mathrm{ord}_m(a)$ to be the smallest positive integer d such that $a^d \equiv 1 \pmod{m}$. It is well-known that $a^d \equiv 1 \pmod{m} \Leftrightarrow \mathrm{ord}_m(a) \mid d$. For a positive real number x, we write $\pi(x)$ to count the number of prime integers $0 < n \leq x$. We use asymptotic notation $O(\cdot)$ to bound the *absolute value* of quantities; for example we write $\varepsilon \leq O(1/n)$ to mean $|\varepsilon| \leq O(1/n)$.

3.4.1 Polynomials and Finite Fields

Lemma 3.11 (Lagrange interpolation)**.** *Let $P \in \mathbb{F}[X]$ be a polynomial of degree $\leq k$. Let $t_1, \ldots, t_{k+1} \in \mathbb{F}$ be distinct points. Then we have:*

$$P(X) = \sum_{j=1}^{k+1} P(t_j) \cdot \frac{\prod_{\ell \neq j}(X - t_\ell)}{\prod_{\ell \neq j}(t_j - t_\ell)}.$$

In particular, equating coefficients of X^k implies:

$$\mathsf{Coeff}_{X^k}(P) = \sum_{j=1}^{k+1} P(t_j) \cdot \frac{1}{\prod_{\ell \neq j}(t_j - t_\ell)}.$$

Lemma 3.12 (Multilinear polynomial interpolation)**.** *Let $P \in \mathbb{F}[X_1, \ldots, X_k]$ be a multilinear polynomial (i.e., it has individual degree ≤ 1). Then we have:*

$$\mathsf{Coeff}_{X_1 X_2 \ldots X_k}[P] = \sum_{\mathbf{b} \in \{0,1\}^k} (-1)^{k-|\mathbf{b}|} P(\mathbf{b}),$$

where $|\mathbf{b}|$ denotes the Hamming weight of $\mathbf{b}$.

Definition 3.13 (Multiplicative structure of finite fields)**.** *Let $\mathbb{F} = \mathbb{F}_Q$ be a finite field of order Q. We denote by $\mathbb{F}^\times$ either the set $\mathbb{F}\backslash\{0\}$ or the multiplicative group that these elements form, depending on context. $\mathbb{F}^\times$ is a cyclic group of order $Q-1$. For $\mathbb{F}$ a finite field and m a positive integer dividing $|\mathbb{F}| - 1$, we let $\mu_m(\mathbb{F})$ denote the cyclic subgroup of $\mathbb{F}^\times$ consisting of mth roots of unity.*

3.4.2 Multiplicative Characters

Definition 3.14 (Multiplicative character)**.** *For a finite field $\mathbb{F}_Q$, a multiplicative character is a group homomorphism $\chi : \mathbb{F}_Q^\times \to \mathbb{C}$. The collection of multiplicative characters forms a cyclic group of order $Q-1$, which we denote by $\widehat{\mathbb{F}_Q^\times}$. The order of a particular character χ is its order in this group. The trivial multiplicative character is the identity element of $\widehat{\mathbb{F}_Q^\times}$, and it maps every nonzero element to 1.*

We will sometimes abuse notation and extend χ to a function on all of $\mathbb{F}$ by defining $\chi(0) = 0$. Note that this extension still satisfies $\chi(xy) = \chi(x)\chi(y)$ for all $x, y \in \mathbb{F}$.

Definition 3.15 (Annihilator)**.** *For m, Q such that $m | Q-1$, we define the annihilator subgroup $\mu_m^\perp(\mathbb{F}) \subseteq \widehat{\mathbb{F}_Q^\times}$ by:*

$$\mu_m^\perp(\mathbb{F}) := \{\chi \in \widehat{\mathbb{F}_Q^\times} : \chi(x) = 1 \text{ for all } x \in \mu_m(\mathbb{F}_Q)\}.$$

$\mu_m^\perp(\mathbb{F})$ is cyclic and consists of $(Q-1)/m$ elements.

Lemma 3.16. *For any $x \in \mathbb{F}_Q$, we have:*

$$\sum_{\chi \in \mu_m^\perp(\mathbb{F}_Q)} \chi(x) = \begin{cases} (Q-1)/m, & \text{if } x \in \mu_m(\mathbb{F}_Q), \\ 0, & \text{otherwise.} \end{cases}$$

Proof. If $x \in \mu_m(\mathbb{F}_Q)$, then by definition we have $\chi(x) = 1$ for all $\chi \in \mu_m^\perp(\mathbb{F}_Q)$, so the sum will be $(Q-1)/m$. If $x = 0$, then we have $\chi(x) = 0$ for all χ, so the sum will indeed be 0. For $x \notin \mu_m(\mathbb{F}_Q) \cup \{0\}$, note that for any $\chi' \in \mu_m^\perp(\mathbb{F}_Q)$ we have:

$$\begin{aligned} \chi'(x) \sum_{\chi \in \mu_m^\perp(\mathbb{F}_Q)} \chi(x) &= \sum_{\chi \in \mu_m^\perp(\mathbb{F}_Q)} (\chi'\chi)(x) = \sum_{\chi \in \mu_m^\perp(\mathbb{F}_Q)} \chi(x) \\ \Rightarrow (\chi'(x) - 1) \sum_{\chi \in \mu_m^\perp(\mathbb{F}_Q)} \chi(x) &= 0. \end{aligned}$$

Since $x \notin \mu_m(\mathbb{F}_Q)$, there will exist a choice of χ' such that $\chi'(x) \neq 1$, so it follows that the sum must be 0. □

Lemma 3.17 (Weil bound, Theorem 2C' on Page 43 of [Sch76]). *Let $\mathbb{F}_Q$ be a finite field, and let χ be a nontrivial multiplicative character of $\mathbb{F}_Q^\times$ of order $h > 1$. Let $f \in \mathbb{F}_Q[X]$ be a nonconstant polynomial that is not of the form $a \cdot g(X)^h$ for some $a \in \mathbb{F}_Q^\times$ and $g \in \mathbb{F}_Q[X]$. Suppose that k is the number of distinct roots of f in the algebraic closure $\overline{\mathbb{F}_Q}$. Then the following holds:*

$$\left| \sum_{\lambda \in \mathbb{F}_Q} \chi(f(\lambda)) \right| \leq (k-1)\sqrt{Q}.$$

3.4.3 Galois Theory

Throughout this section, let Q be a prime power and $d > 1$ an integer. All definitions and facts here are with respect to the field extension $\mathbb{F}_{Q^d}/\mathbb{F}_Q$.

Definition 3.18 (Automorphisms). *A mapping $\sigma : \mathbb{F}_{Q^d} \to \mathbb{F}_{Q^d}$ is said to be an* automorphism *if it is a field homomorphism, bijective, and satisfies $\sigma(x) = x$ for every $x \in \mathbb{F}_Q$. Such automorphisms form a group which we denote by* $\mathrm{Aut}(\mathbb{F}_{Q^d}/\mathbb{F}_Q)$.

This group has order exactly d. In fact, every automorphism has the form $\sigma(x) = x^{Q^i}$ for some $i \in [0, d-1]$.

Definition 3.19 (Norm). *For an element $\alpha \in \mathbb{F}_{Q^d}$, we define the norm $N_{\mathbb{F}_{Q^d}/\mathbb{F}_Q}(\alpha)$ to be:*

$$N_{\mathbb{F}_{Q^d}/\mathbb{F}_Q}(\alpha) = \prod_{\sigma \in \mathrm{Aut}(\mathbb{F}_{Q^d}/\mathbb{F}_Q)} \sigma(\alpha) = \alpha^{1+Q+\ldots+Q^{d-1}}.$$

In particular, when $\alpha \in \mathbb{F}_Q$, we have $N_{\mathbb{F}_{Q^d}/\mathbb{F}_Q}(\alpha) = \alpha^d$.

Lemma 3.20. *For any $\alpha \in \mathbb{F}_{Q^d}^\times$, we have $N_{\mathbb{F}_{Q^d}/\mathbb{F}_Q}(\alpha) \in \mathbb{F}_Q^\times$. In fact, $N_{\mathbb{F}_{Q^d}/\mathbb{F}_Q}$ is a surjective group homomorphism from $\mathbb{F}_{Q^d}^\times$ to $\mathbb{F}_Q^\times$.*

3.4.4 Cyclotomic Polynomials

Definition 3.21 (Cyclotomic polynomials). *Let n be a positive integer. The nth cyclotomic polynomial is the unique monic irreducible polynomial with integer coefficients that is a divisor of $x^n - 1$ and is not a divisor of $x^k - 1$ for any $k < n$.*

Fact 3.22.

$$x^n - 1 = \prod_{d \mid n} \Phi_d(x).$$

3.4.5 Prime Density Estimates

Lemma 3.23 (Bertrand, Chebyshev). *If $n > 1$, then there is at least one prime p such that $n < p < 2n$.*

Lemma 3.24. *For any $\ell \geq 5$, the product of the first ℓ odd primes is $\leq (\ell \log \ell)^\ell$.*

Definition 3.25 (Rough numbers). *Let $\varepsilon \in (0, 1/2)$. Recall that for a positive integer n, we define $P^+(n)$ to be the largest prime divisor of n. We say that a positive integer n is ε-rough if $P^+(n) \geq n^{(1-\varepsilon)}$.*[3]

[3]We note that this is non-standard terminology; usually, roughness refers to the property where *all* prime divisors of the number are large, not just the largest prime divisor. When $\varepsilon < 1/2$, the standard definition is essentially meaningless because it is equivalent to n being prime.

Lemma 3.26. *Let q be a prime, and let $d \geq 6$ be a positive integer such that $q^d - 1$ is ε rough for some $0 < \varepsilon \leq 1/6$. Then, if $p = P^+(q^d - 1)$, we must have that $\mathrm{ord}_p(q) = d$.*

Proof. Observe that $\mathrm{ord}_p(q) | d$. Suppose for contradiction that $\mathrm{ord}_p(q) = e < d$. Then, $e \leq d/2$, so $p | q^e - 1 \Rightarrow p \leq q^e - 1 \leq q^{d/2} - 1$, and so $p \leq q^e - 1 \leq q^{d/2} - 1 \leq q^{d/2}$. This is a contradiction, since by hypothesis, we must have that

$$p \geq (q^d - 1)^{(1-\varepsilon)} \geq (q^d - 1)^{5/6} > q^{d/2},$$

where the last inequality follows from the following calculation:

$$\begin{aligned}
\frac{(q^d - 1)^{5/6}}{q^{d/2}} &= \left(q^{2d/5} - \frac{1}{q^{3d/5}} \right)^{5/6} \\
&\geq \left(q^{2d/5} - 1 \right)^{5/6} && \text{since } q > 1 \\
&\geq (2^{12/5} - 1)^{5/6} && \text{since } q \geq 2, d \geq 6 \\
&\geq 3.27^{5/6} > 1.
\end{aligned}$$

□

Lemma 3.27. *Let q be a prime. Suppose that $d, d' \geq 6$ are distinct positive integers such that $q^d - 1$ and $q^{d'} - 1$ are ε-rough, for some $0 < \varepsilon \leq 1/6$. Then,*

$$P^+(q^d - 1) \neq P^+(q^{d'} - 1).$$

Proof. Let $p = P^+(q^d - 1)$ and $p' = P^+(q^{d'} - 1)$. Lemma 3.26 shows that $\mathrm{ord}_p(q) = d$, and similarly $\mathrm{ord}_{p'}(q) = d'$. Since the order is well-defined, we are done. □

Lemma 3.28 ([Dic30], Equation (1) of [Mor14])**.** *For every $\alpha > 0$,*

$$\lim_{N \to \infty} \Pr_{n \leftarrow [1,N]} [P^+(n) \leq N^\alpha] = \rho(1/\alpha),$$

where ρ is the Dickman-de Bruijn function, and $\rho(1/\alpha) = 1 - \log(1/\alpha)$ for $1/2 \leq \alpha \leq 1$.

The following strengthening of Mertens theorem can be proven assuming the Generalized Riemann Hypothesis (See Section C).

Lemma 3.29 (Strong Mertens' Theorem)**.** *Assume the GRH for all Dirichlet L-functions. There is a universal constant $C > 0$ such that for all integers $m \geq 1$, and a with $\gcd(a, m) = 1$, and all real numbers $y \geq x \geq 8$ such that $m \leq x$,*

$$\left| \left(\sum_{\substack{\text{prime } p \in [x,y] \\ p = a \bmod m}} \frac{1}{p} \right) - \frac{1}{\varphi(m)} (\log \log y - \log \log x) \right| \leq C \frac{\log x}{\sqrt{x}}.$$

4 New Constructions of S-Decoding Polynomials

In this section, we first reduce the task of constructing a sparse decoding polynomial to the (equivalent) task of constructing a primitive we call "root-of-unity grid" (Definition 4.1). Section 4.1 shows this equivalence. We then give two approaches to constructing these root-of-unity grids: (1) sequences of rough repunits (Section 4.2) and (2) prime field norms (Section 4.3). The relevant algebraic and number-theoretic background can be found in Section 3.4, and will be referenced when required. Finally, Section 4.5 shows how to stitch many root-of-unity grids together to obtain a sparse decoding polynomial. This will be relevant in the regime of superconstant number of servers.

4.1 Equivalence of S-Decoding Polynomials and Root-of-Unity Grids

Definition 4.1 (Root-of-unity grid)**.** *Let $m = p_1 \dots p_k$ be a product of k distinct primes, and let $\mathbb{F}$ be a finite field with $m|(|\mathbb{F}| - 1)$.*[4] *A $(\mathbb{F}, m)$-root-of-unity grid is a collection of elements $a_1, \dots, a_k \in \mathbb{F}\backslash\{0\}$ and distinct elements $t_1, \dots, t_{k+1} \in \mathbb{F}$ such that*

$$1 + a_i t_j \in \mu_{p_i}(\mathbb{F}),$$

for all $i \in [k]$, $j \in [k+1]$.

The following straightforward claim is not necessary for our results, but we present it for intuition.

Claim 4.1. *In any root-of-unity grid, the elements $a_1, \dots, a_k$ must be distinct.*

Proof. If they are not, then there exist different $i, i' \in [k]$ such that $1 + a_i t_j = 1 + a_{i'} t_j$ for all $j \in [k+1]$. This common element lies in $\mu_{p_i}(\mathbb{F}) \cap \mu_{p_{i'}}(\mathbb{F})$, so it must be 1. This forces $t_j = 0$ for all $j \in [k+1]$ which contradicts the requirement that the t_j's are distinct. □

4.1.1 Root-of-Unity Grid ⇒ Sparse Decoding Polynomial

Lemma 4.2. *Let $m = p_1 \cdots p_k$ be a product of k distinct primes, and let $\mathbb{F}$ be a finite field with $m|(|\mathbb{F}| - 1)$. Suppose that there exists a $(\mathbb{F}, m)$-root-of-unity grid. Then there exists an S^*_m-decoding polynomial over $\mathbb{F}$ of sparsity $k+1$, for some primitive mth root of unity γ.*

Proof. Let $\gamma \in \mathbb{F}$ be a primitive mth root of unity, and let $a_1, \dots, a_k \in \mathbb{F}\backslash\{0\}$ and $t_1, \dots, t_{k+1} \in \mathbb{F}$ form a $(\mathbb{F}, m)$-root-of-unity grid. That is, $1 + a_i t_j \in \mu_{p_i}(\mathbb{F})$ for all $i \in [k]$, $j \in [k+1]$. For $j \in [k+1]$, define the Lagrange weights

$$w_j := \frac{1}{\prod_{\ell \neq j}(t_j - t_\ell)}.$$

Let $c_1, \dots, c_k \in \mathbb{Z}_m$ form a CRT basis for m i.e., we have $c_i \bmod p_\ell = 1$ if $\ell = i$ and 0 otherwise. For each $i \in [k]$, let $\gamma_i := \gamma^{c_i}$. γ_i is a generator of $\mu_{p_i}(\mathbb{F})$, so by the root-of-unity grid condition there exist $e_{i,j}$ such that $\gamma_i^{e_{i,j}} = (1 + a_i t_j)^{-1}$. Then, define

$$P(X) := \sum_{j=1}^{k+1} \left(w_j \cdot \prod_{i=1}^{k} (1 + a_i t_j) \right) \cdot X^{\sum_{\ell=1}^{k} e_{\ell,j} \cdot c_\ell}.$$

[4] Therefore there exists a primitive mth root of unity.

One can check that $P(X)$ has sparsity exactly $k+1$ (no two terms have the same power of X because $t_1, \dots, t_{k+1}$ are distinct so the tuples $\{(e_{1,j}, e_{2,j}, \dots, e_{k,j}) : j \in [k+1]\}$ are also distinct). We now claim that $P(X)$ is an S_m^*-decoding polynomial. Indeed, consider any $x \in S_m^*$. We may write $x = \sum_{i=1}^k u_i c_i$ where $u_i \in \{0,1\}$ for all i. Then we have:

$$\begin{aligned}
P(\gamma^x) &= \sum_{j=1}^{k+1} \left(w_j \cdot \prod_{i=1}^{k} (1 + a_i t_j) \right) \cdot \gamma^{x(\sum_{\ell=1}^k e_{\ell,j} \cdot c_\ell)} \\
&= \sum_{j=1}^{k+1} \left(w_j \cdot \prod_{i=1}^{k} (1 + a_i t_j) \right) \cdot \gamma^{\sum_{\ell,\ell'=1}^k e_{\ell,j} \cdot u_{\ell'} \cdot c_\ell \cdot c_{\ell'}} && \text{writing } x = \sum u_{\ell'} c_{\ell'} \\
&= \sum_{j=1}^{k+1} \left(w_j \cdot \prod_{i=1}^{k} (1 + a_i t_j) \right) \cdot \gamma^{\sum_{\ell=1}^k e_{\ell,j} \cdot u_\ell \cdot c_\ell} && \text{since } c_\ell c_{\ell'} \equiv \mathbb{1}[\ell = \ell'] c_\ell \pmod m \\
&= \sum_{j=1}^{k+1} \left(w_j \cdot \prod_{i=1}^{k} (1 + a_i t_j) \right) \cdot \prod_{\ell=1}^{k} \gamma_\ell^{e_{\ell,j} \cdot u_\ell} && \text{by definition of } \gamma_\ell \\
&= \sum_{j=1}^{k+1} w_j \cdot \prod_{i=1}^{k} (1 + a_i t_j)^{1-u_i} && \text{by definition of } e_{\ell,j}.
\end{aligned}$$

Now define the univariate polynomial $f_{u_1,\dots,u_k}(X) = \prod_{i=1}^k (1 + a_i X)^{1-u_i}$. It has degree $\le k$. Applying Lemma 3.11 to $f_{u_1,\dots,u_k}$ and the evaluation points $t_1, \dots, t_{k+1}$ implies:

$$\begin{aligned}
\mathsf{Coeff}_{X^k}(f_{u_1,\dots,u_k}) &= \sum_{j=1}^{k+1} f_{u_1,\dots,u_k}(t_j) \cdot \frac{1}{\prod_{\ell \neq j} (t_j - t_\ell)} \\
&= \sum_{j=1}^{k+1} w_j f_{u_1,\dots,u_k}(t_j) \\
&= P(\gamma^x),
\end{aligned}$$

by the previous computation. Now we have two cases:

- If $x \neq 0$, then we have $(u_1, \dots, u_k) \neq (0, \dots, 0)$, so $f_{u_1,\dots,u_k}$ has degree $\le k-1$, implying $P(\gamma^x) = 0$.
- If $x = 0$, then we have $(u_1, \dots, u_k) = (0, \dots, 0)$ so we have $P(\gamma^x) = \mathsf{Coeff}_{X^k}(f_{u_1,\dots,u_k}) = a_1 a_2 \dots a_k \neq 0$.

This completes the proof. □

4.1.2 Sparse Decoding Polynomial ⇒ Root-of-Unity Grid

We begin with a straightforward preliminary lemma:

Lemma 4.3. *Assume $k \ge 2$. Let $\{z_{i,j} \in \mathbb{F} \setminus \{0\} : i \in [k], j \in [k+1]\}$ be such that for any distinct $\ell, a, b \in [k+1]$ and distinct $r, s \in [k]$ we have:*

$$(z_{r,\ell} - z_{r,a})(z_{s,\ell} - z_{s,b}) = (z_{r,\ell} - z_{r,b})(z_{s,\ell} - z_{s,a}) \neq 0.$$

Then there exist $a_1, \dots, a_k \in \mathbb{F} \setminus \{0\}$, and distinct $t_1, \dots, t_{k+1} \in \mathbb{F}$ such that $z_{i,j} = z_{i,1} + a_i t_j$ for all $i \in [k], j \in [k+1]$.

Proof. For any $r \in [k], a \in [k+1]$, define $M_{r,a} = z_{r,a} - z_{r,1}$. Then for any distinct $a, b \in [2, k+1]$ and $r, s \in [k]$, we have:

$$\begin{aligned} M_{r,a}M_{s,b} &= (z_{r,a} - z_{r,1})(z_{s,b} - z_{s,1}) \\ &= (z_{r,b} - z_{r,1})(z_{s,a} - z_{s,1}) \\ &= M_{r,b}M_{s,a}. \end{aligned}$$

Moreover, we have $M_{r,a} = 0 \Leftrightarrow a = 1$. Consequently, the above remains true if either $a = 1$ or $b = 1$; both sides will simply be zero. It follows M is a rank 1 matrix i.e., there exist $a_1, \ldots, a_k, t_2, \ldots, t_{k+1} \in \mathbb{F}\backslash\{0\}$ and $t_1 = 0$ such that $M_{i,j} = a_i t_j$ for all $i \in [k], j \in [k+1]$. Consequently, we have $z_{i,j} = z_{i,1} + a_i t_j$. It remains to check the nonzero and distinctness conditions:

- The a_i's are nonzero by construction.
- Finally, if $t_j = t_{j'}$ this would imply $z_{i,j} = z_{i,j'} \Rightarrow j = j'$. So the t_j's are also distinct.

□

The following converse of Lemma 4.2 builds on the second approach by [GK26] to show that there does not exist an S_m^*-decoding polynomial over $\mathbb{F}$ of sparsity $\leq k$:

Lemma 4.4. *Assume $k \geq 2$. Let $\mathbb{F}, m$ (with m a product of k distinct primes) be such that there is an S_m^*-decoding polynomial over $\mathbb{F}$ of sparsity $k+1$. Then there exists a root-of-unity grid.*

Proof. Again, let $\gamma \in \mathbb{F}$ be a primitive mth root of unity, $c_1, \ldots, c_k \in \mathbb{Z}_m$ a CRT basis for m, and for each $i \in [k]$ let $\gamma_i = \gamma^{c_i}$. Let us write the S-decoding polynomial as:

$$P(X) = \sum_{j=1}^{k+1} \alpha_j X^{\beta_j},$$

for $\alpha_j \neq 0 \forall j \in [k+1]$. Assume without loss of generality that $P(1) = 1$ (by scaling). In the following, let $z_{i,j} = \gamma_i^{-\beta_j} \in \mu_{p_i}(\mathbb{F})$ for every $i \in [k], j \in [k+1]$. Then the S-decoding property tells us that for any $(u_1, \ldots, u_k) \in \{0,1\}^k$, we have:

$$\sum_{j=1}^{k+1} \alpha_j \prod_{i=1}^{k} \gamma_i^{u_i \beta_j} = P\left(\prod_{i=1}^{k} \gamma_i^{u_i}\right) = \begin{cases} 1, & \text{if } u_i = 0 \forall i \\ 0, & \text{otherwise.} \end{cases}$$

To analyze this condition, introduce formal variables $X_1, \ldots, X_k$ and notice that:

$$\begin{aligned} \sum_{j=1}^{k+1} \alpha_j \prod_{i=1}^{k} (1 + X_i z_{i,j}^{-1}) &= \sum_{j=1}^{k+1} \alpha_j \left(\sum_{u_1,\ldots,u_k \in \{0,1\}} \prod_{i=1}^{k} X_i^{u_i} \gamma_i^{u_i \beta_j} \right) \\ &= \sum_{u_1,\ldots,u_k \in \{0,1\}} \left(\prod_{i=1}^{k} X_i^{u_i} \right) \cdot \sum_{j=1}^{k+1} \alpha_j \prod_{i=1}^{k} \gamma_i^{u_i \beta_j} \\ &= 1. \end{aligned} \tag{4}$$

The idea is to argue that $\{z_{i,j} : i \in [k], j \in [k+1]\}$ forms our desired root-of-unity grid (up to some rescaling). To this end, consider any permutation $\sigma : [k+1] \to [k+1]$. Then we can substitute $X_i = -z_{i,\sigma(i)}$ into (4). The only term on the LHS that will survive is that corresponding to $j = \sigma(k+1)$, yielding:

$$\begin{aligned} \alpha_{\sigma(k+1)} \prod_{i=1}^{k} \left(1 - \frac{z_{i,\sigma(i)}}{z_{i,\sigma(k+1)}}\right) &= 1 \\ \Rightarrow \prod_{i=1}^{k} (z_{i,\sigma(k+1)} - z_{i,\sigma(i)}) &= \frac{1}{\alpha_{\sigma(k+1)}} \prod_{i=1}^{k} z_{i,\sigma(k+1)}. \end{aligned} \tag{5}$$

We begin by noting that the RHS of (5) is clearly nonzero, so this must also be true of the LHS. Consequently, for any $i \in [k]$ and distinct $j, j' \in [k+1]$ we have $z_{i,j} \neq z_{i,j'}$. It follows that $\beta_j \not\equiv \beta_{j'} \pmod{p_i}$, so in particular the β_j's are distinct modulo m.

Now, fix distinct $\ell, a, b \in [k+1]$ and distinct $r, s \in [k]$. We consider a permutation σ on $[k+1]$ such that $\sigma(k+1) = \ell, \sigma(r) = a, \sigma(s) = b$. We also define σ' to be identical to σ except we swap to get $\sigma'(r) = b, \sigma'(s) = a$. Then (5) implies that:

$$\begin{aligned} \prod_{i=1}^{k} (z_{i,\ell} - z_{i,\sigma(i)}) &= \frac{1}{\alpha_\ell} \prod_{i=1}^{k} z_{i,\ell} \\ &= \prod_{i=1}^{k} (z_{i,\ell} - z_{i,\sigma'(i)}) \\ \Rightarrow (z_{r,\ell} - z_{r,a})(z_{s,\ell} - z_{s,b}) &= (z_{r,\ell} - z_{r,b})(z_{s,\ell} - z_{s,a}). \end{aligned} \tag{6}$$

At this point, we apply Lemma 4.3 to deduce that there exist $a'_1, \ldots, a'_k \in \mathbb{F}\backslash\{0\}$ and distinct $t_1, \ldots, t_{k+1} \in \mathbb{F}$ such that $z_{i,j} = z_{i,1} + a'_i t_j$ for all $i \in [k], j \in [k+1]$. To construct our root of unity grid, we will take $a_i = a'_i / z_{i,1}$ for all $i \in [k]$ and the same $t_1, \ldots, t_{k+1}$. We check the desired conditions:

- For each i, j, we have $1 + a_i t_j = (z_{i,1} + a'_i t_j)/z_{i,1} = z_{i,j}/z_{i,1} \in \mu_{p_i}(\mathbb{F})$.
- We have $a'_i \neq 0 \Rightarrow a_i \neq 0$ for all i.
- We are given by Lemma 4.3 that $t_1, \ldots, t_{k+1}$ are distinct.

The conclusion follows. □

4.2 Approach 1: Rough Repunits Imply Root-of-Unity Grids

In this section, we show in that to construct a root-of-unity grid, it suffices to find a sequence of rough repunits,[5] that is, numbers of the form $q^{d_1} - 1, \ldots, q^{d_k} - 1$ such that the largest prime factor of $q^{d_i} - 1$ is at least $(q^{d_i} - 1)^{1-1/(3k)}$ (see Corollary 4.6). One can intuit the results of this section in terms of the probabilistic method, as we sketched in Section 2.2.

Lemma 4.5 (Sufficient conditions for a root-of-unity grid)**.** *Let $k \geq 2$ be an integer, and let q be a prime. Suppose there exist positive integers $d_1, \ldots, d_k$ and pairwise distinct primes $p_1, \ldots, p_k$ such that the following hold:*

[5] Repunits in base q are numbers of the form $\frac{q^d-1}{q-1}$, not $q^d - 1$, but we do not make this distinction since a rough repunit (that is, a repunit with a large prime factor) will automatically give us a rough number of the form $q^d - 1$ and vice versa.

- $p_i \mid q^{d_i} - 1$ *for all* $i \in [k]$;
- $q^{\gcd(d_1,\dots,d_k)} \geq k+1$;
- *for all* $i \in [k]$, *we have that*

$$\frac{q^{d_i} - k}{h_i^k} > (k-1)\sqrt{q^{d_i}}(1 - 1/h_i^k) + 1,$$

where $h_i := \frac{q^{d_i}-1}{p_i}$.

Then, there exists a $(\mathbb{F}_{q^D}, m)$*-root-of-unity grid, where* $D = \operatorname{lcm}(d_1, \dots, d_k)$ *and* $m = p_1 \cdots p_k$.

Proof. Recall that a product $m = p_1 \cdots p_k$ of k primes is defined to have a root-of-unity grid if there exists a field $\mathbb{F}$ containing a primitive mth root of unity γ, $a_1, \dots, a_k \in \mathbb{F}^\times$, and distinct $t_1, \dots, t_{k+1} \in \mathbb{F}$ such that $1 + a_i t_j \in \mu_{p_i}(\mathbb{F})$ for all $i \in [k]$, $j \in [k+1]$. Note that without loss of generality, we can take $t_{k+1} = 0$.

Suppose m is such that there exist $q, d_1, \dots, d_k, p_1, \dots, p_k$ satisfying the preconditions of the statement. Fix arbitrary $T = \{t_1, \dots, t_k\} \subseteq \mathbb{F}_{q^{\gcd(d_1,\dots,d_k)}} \setminus \{0\}$. Such a T can be chosen since $q^{\gcd(d_1,\dots,d_k)} \geq k+1$. Applying Claim 4.2 with $Q = q^{\gcd(d_1,\dots,d_k)}$, $d = d_i/\gcd(d_1, \dots, d_k)$, and $p = p_i$ for each $i \in [k]$, we know that there exists $a_i \in \mathbb{F}^\times_{q^{d_i}}$ such that $1 + a_i T \subseteq \mu_{p_i}(\mathbb{F}_{q^{d_i}})$.

To finish, we take $\mathbb{F} = \mathbb{F}_{q^D}$, where $D = \operatorname{lcm}(d_1, \dots d_k)$, to be a common extension of all the $\mathbb{F}_{q^{d_i}}$ fields and the $a_1, \dots, a_k$ chosen above. $\mathbb{F}$ has a primitive mth root of unity since $p_i \mid q^{d_i} - 1 \Rightarrow p_i \mid q^D - 1 \forall i \Rightarrow m \mid q^D - 1 = |\mathbb{F}| - 1$. This yields the desired root-of-unity grid. □

Claim 4.2. *Let* $k \geq 2$ *be an integer, and let* $Q \geq k+1$ *be a prime power. Let* $T = \{t_1, \dots, t_k\} \subseteq \mathbb{F}_Q \setminus \{0\}$ *be an arbitrary set of* k *distinct nonzero elements. Suppose* d, p *are positive integers with* p *prime such that* $p \mid Q^d - 1$ *and*

$$\frac{Q^d - k}{h^k} > (k-1)\sqrt{Q^d}(1 - 1/h^k) + 1, \tag{7}$$

where $h := \frac{Q^d-1}{p}$. *Then, there exists some* $a \in \mathbb{F}_{Q^d} \setminus \{0\}$ *such that* $1 + at_j \in \mu_p(\mathbb{F}_{Q^d})$ *for every* $j \in [k]$.

Proof. We begin by disposing of the case where $h = 1$. In this case, $\mu_p(\mathbb{F}_{Q^d}) = \mathbb{F}^\times_{Q^d}$, so all we need to ensure is that $1 + at_j \neq 0$ for all $j \in [k]$. This rules out exactly k values of a, so it suffices to check that $Q^d - 1 > k$, which is immediate from (7). So assume from now on that $h > 1$.

For fixed p, d, the existence of some $a \in \mathbb{F}_{Q^d} \setminus \{0\}$ such that $1 + aT \subseteq \mu_p(\mathbb{F}_{Q^d})$ is equivalent to the integer

$$N_{p,d} := \#\{\lambda \in \mathbb{F}_{Q^d} : 1 + \lambda T \subseteq \mu_p(\mathbb{F}_{Q^d})\}$$

being strictly greater than 1 (it will be at least 1 due to the trivial case of $\lambda = 0$). Let $h := \frac{Q^d-1}{p}$, and let $\mu_p^\perp \leq \widehat{\mathbb{F}^\times_{Q^d}}$ be the annihilator of μ_p as defined in Definition 3.15. In the following, we will use the convention specified in Definition 3.14 and also define $\chi(0) = 0$ for any character χ on $\mathbb{F}^\times_{Q^d}$. Then $N_{p,d}$ can be expressed as the character sum:

$$N_{p,d} = \sum_{\lambda \in \mathbb{F}_{Q^d}} \prod_{j \in [k]} \mathbb{1}\left[1 + \lambda t_j \in \mu_p(\mathbb{F}_{Q^d})\right]$$

$$= \sum_{\lambda \in \mathbb{F}_{Q^d}} \prod_{j \in [k]} \left(\frac{1}{h} \sum_{\chi \in \mu_p^\perp} \chi(1 + \lambda t_j) \right) \text{ (Lemma 3.16)}$$

$$= \frac{1}{h^k} \sum_{\chi_1, \ldots, \chi_k \in \mu_p^\perp} \sum_{\lambda \in \mathbb{F}_{Q^d}} \left(\prod_{j \in [k]} \chi_j(1 + \lambda t_j) \right).$$

Then, the all-trivial character tuple $(\chi_1, \ldots, \chi_k) = (1, \ldots, 1)$ contributes $\frac{Q^d - k}{h^k}$ to the sum, because there are exactly k values of λ such that $1 + \lambda t_j = 0$ for some j.

Recall from Definition 3.15 that $\mu_p^\perp$ is a cyclic group of order h. So letting χ be a generator of $\mu_p^\perp$, we can rewrite $\chi_j = \chi^{r_j}$ for $r_j \in \{1, 2, \ldots, h\}$.

$$\begin{aligned} N_{p,d} &= \frac{Q^d - k}{h^k} + \frac{1}{h^k} \sum_{(r_1, \ldots, r_k) \in \{1, \ldots, h\}^k \setminus \{h^k\}} \sum_{\lambda \in \mathbb{F}_{Q^d}} \left(\prod_{j \in [k]} \chi^{r_j}(1 + \lambda t_j) \right) \\ &= \frac{Q^d - k}{h^k} + \frac{1}{h^k} \sum_{(r_1, \ldots, r_k) \in \{1, \ldots, h\}^k \setminus \{h^k\}} \sum_{\lambda \in \mathbb{F}_{Q^d}} \chi \left(\prod_{j \in [k]} (1 + \lambda t_j)^{r_j} \right) \\ &\geq \frac{Q^d - k}{h^k} - \frac{1}{h^k} \sum_{(r_1, \ldots, r_k) \in \{1, \ldots, h\}^k \setminus \{h^k\}} \left| \sum_{\lambda \in \mathbb{F}_{Q^d}} \chi(f_{r_1, \ldots, r_k}(\lambda)) \right| \end{aligned}$$

where $f_{r_1, \ldots, r_k}(\lambda) := \prod_{j \in [k]} (1 + \lambda t_j)^{r_j}$, and the above inequality follows from the triangle inequality. We now apply the Weil bound for character sums (Lemma 3.17) to $f_{r_1, \ldots, r_k}$ and χ. We check the preconditions:

- We are assuming that $h > 1$.
- χ is defined to be a generator of $\mu_p^\perp$, so it has order h.
- Note that the factors $(1 + \lambda t_j)^{r_j}$ are pairwise coprime (viewed as elements of $\mathbb{F}_Q[\lambda]$) by distinctness of the t_j's. Consequently, $f_{r_1, \ldots, r_k}$ can only be an hth power (up to scaling) if each of these factors is an hth power (up to scaling). For this to be true, we need $r_j \equiv 0 \pmod{h}$ for all j, which contradicts the fact that we are only applying the bound when $(r_1, \ldots, r_k) \neq (h, \ldots, h)$.
- Finally, it is clear from definition that $f_{r_1, \ldots, r_k}$ has k distinct roots.

Therefore applying Lemma 3.17 tells us that:

$$\begin{aligned} N_{p,d} &\geq \frac{Q^d - k}{h^k} - \frac{1}{h^k} \sum_{(r_1, \ldots, r_k) \in \{1, \ldots, h\}^k \setminus h^k} (k - 1)\sqrt{Q^d} \\ &\geq \frac{Q^d - k}{h^k} - \frac{h^k - 1}{h^k}(k - 1)\sqrt{Q^d} \\ &= \frac{Q^d - k}{h^k} - (k - 1)\sqrt{Q^d}(1 - 1/h^k), \end{aligned}$$

which by assumption is > 1, as desired. □

Corollary 4.6 (Rough repunits suffice for root-of-unity grids)**.** *Let $k \geq 2$ be an integer. Suppose there exists a prime q and pairwise distinct positive integers $d_1, \ldots, d_k$ such that the following hold:*

- *$q \geq k+1$ and $d_i \geq 6$ for all $i \in [k]$, and*
- *$q^{d_i} - 1$ is ε-rough, for $\varepsilon = 1/(3k)$ for all $i \in [k]$.*

Then, there exists a $(\mathbb{F}_{q^D}, m)$-root-of-unity grid where $D := \operatorname{lcm}(d_1, \ldots, d_k)$, $m := p_1 \cdots p_k$, and $p_i := P^+(q^{d_i} - 1)$.

Proof. By Lemma 3.27, we see that the primes p_i must be pairwise distinct. Since $q \geq k+1$, we automatically get that $q^{\gcd(d_1,\ldots,d_k)} \geq k+1$, so all that remains is to deduce the third bullet of the hypothesis in Lemma 4.5.

Since $p_i \geq (q^{d_i} - 1)^{1-1/(3k)}$ for all i, we get that

$$h_i := \frac{q^{d_i} - 1}{p_i} \leq (q^{d_i} - 1)^{1/(3k)} \leq q^{d_i/(3k)}.$$

Now, writing $x = q^{d_i/6} \geq k$, we have:

$$\begin{aligned}
\frac{q^{d_i} - k}{h_i^k} - (k-1)\sqrt{q^{d_i}}(1 - 1/h_i^k) &> \frac{q^{d_i} - k}{q^{d_i/3}} - (k-1)\sqrt{q^{d_i}} \\
&= x^4 - k/x^2 - (k-1)x^3 \\
&= x^3(x - (k-1)) - k/x^2 \\
&\geq k^3 - k/x^2 \\
&\geq k^3 - 1/k \\
&> 1.
\end{aligned}$$

Therefore, all the preconditions of Lemma 4.5 are met, and the statement follows. □

4.3 Approach 2: Prime Field Norms Imply Root-of-Unity Grids

In this section, we present a more algebraically-structured approach to constructing root-of-unity grids. This will not be used in any of our main results, but we show it for completeness and as an abstract explanation of why Efremenko's decoding polynomial for $m = 7 \cdot 73$ and $k = 2$ exists [Efr09] (which was not captured by the results of [CFL+13]).

Lemma 4.7. *Let Q be a prime power and $d > 1$ a positive integer such that $(Q^d - 1)/(Q - 1) = p$ is prime. Let $T = \{t_1, \ldots, t_k\} \subseteq \mathbb{F}_Q^\times$ be an arbitrary set of k distinct nonzero elements. Assume there exists a polynomial $P(X) \in \mathbb{F}_Q[X]$ such that:*

- *$\deg P = d$ and $P(0) = 1$;*
- *$P(t) = 1$ for every $t \in T$; and*
- *the polynomial $f(X) := X^d \cdot P(-1/X) \in \mathbb{F}_Q[X]$ is irreducible.*

Then there exists $a \in \mathbb{F}_{Q^d}^\times$ such that $1 + aT \subseteq \mu_p(\mathbb{F}_{Q^d})$.

Proof. Since $f \in \mathbb{F}_Q[X]$ is an irreducible polynomial over $\mathbb{F}_Q$ of degree d, we know that all its roots must appear in $\mathbb{F}_{Q^d}$. We also know that $f(0) \neq 0$. We will take $a \in \mathbb{F}_{Q^d}^\times$ to be any root of f and argue that it satisfies the given conditions. Firstly, note that $f \in \mathbb{F}_Q[X]$ so for any automorphism $\sigma \in \mathrm{Aut}(\mathbb{F}_{Q^d}/\mathbb{F}_Q)$ we have $f(\sigma(a)) = 0$. So $a, a^Q, a^{Q^2}, \ldots, a^{Q^{d-1}}$ are all roots of f. Moreover, they must be distinct; if some two of these are equal then powering both sides implies that there is a $d' < d$ such that $a^{Q^{d'}} = a$, which means that $a \in \mathbb{F}_{Q^{d'}}$ which contradicts that fact that a's minimal polynomial over $\mathbb{F}_Q$ has degree d. Also, $f(X)$ must be monic since $P(0) = 1$. Consequently, we must have:

$$f(X) = \prod_{i=0}^{d-1} (X - a^{Q^i}).$$

Now, for any $t \in \mathbb{F}_Q^\times$, we have:

$$\begin{aligned}
(1+at)^p &= (1+at)^{(Q^d-1)/(Q-1)} \\
&= (1+at)^{1+Q+Q^2+\ldots+Q^{d-1}} \\
&= \prod_{i=0}^{d-1} (1+at)^{Q^i} \\
&= \prod_{i=0}^{d-1} (1+a^{Q^i}t) && \text{since } t \in \mathbb{F}_Q \Rightarrow t^{Q^i} = t \\
&= (-t)^d \prod_{i=0}^{d-1} (-1/t - a^{Q^i}) \\
&= (-t)^d f(-1/t) \\
&= (-t)^d \cdot (-1/t)^d \cdot P(t) \\
&= P(t).
\end{aligned}$$

So for any $t \in T$, we have $(1+at)^p = P(t) = 1 \Rightarrow 1 + at \in \mu_p(\mathbb{F}_{Q^d})$, as desired. □

This approach also explains Efremenko's example $m = 7 \times 73$ over $\mathbb{F}_{2^9}$: taking $Q = 8$ and $d = 3$ so that $(8^3 - 1)/(8 - 1) = 73$ is prime, we can obtain the 73-row of the root-of-unity grid from the above argument over the extension $\mathbb{F}_{8^3}/\mathbb{F}_8$. The 7-row is simpler to construct, by using the observation that $\mathbb{F}_8^\times = \mu_7(\mathbb{F}_{2^9})$. Thus, after choosing a suitable two-point set T inside $\mathbb{F}_8$, the two rows together form a root-of-unity grid over $\mathbb{F}_{2^9}$, and hence give a 3-sparse decoding polynomial for $m = 7 \times 73$.

Corollary 4.8 (Rederivation of Efremenko's decoding polynomial [Efr09])**.** *Let $m = 7 \times 73$. Then there exists an S_m^*-decoding polynomial over $\mathbb{F}_{2^9}$ of sparsity 3.*

Proof. Let $p_1 = 7$ and $p_2 = 73$. By Lemma 4.2, it suffices to show that there is a root-of-unity grid over $\mathbb{F}_{2^9}$. Take $\theta \in \mathbb{F}_{2^3}$ such that $\theta^3 + \theta + 1 = 0$. For our root-of-unity grid, we will take $T = \{1, \theta, 0\}$. First, we will take $a_1 = \theta$. We have $1 + a_1 T = \{1 + \theta, 1 + \theta^2, 1\}$. These are all nonzero elements of $\mathbb{F}_8$, so this choice of a_1 works precisely because $\mathbb{F}_8^\times = \mu_7(\mathbb{F}_{2^9})$.

To choose a_2, we will use Lemma 4.7 with $Q = 8$ and $d = 3$. We will take:

$$P(X) = 1 + X(X-1)(X-\theta)$$

$$
\begin{aligned}
&= 1 + X(X+1)(X+\theta) && \text{since we are in characteristic 2}\\
\Rightarrow f(X) &= X^3\,(1 + (-1/X)(-1/X-1)(-1/X-\theta))\\
&= X^3 + (X+1)(X\theta+1)\\
&= X^3 + \theta X^2 + (1+\theta)X + 1.
\end{aligned}
$$

By definition of P, we have $P(0) = P(1) = P(\theta) = 1$. A direct check verifies that $f(X)$ has no roots over $\mathbb{F}_8$ and is therefore irreducible. Thus Lemma 4.7 is applicable so there exists $a_2 \in \mathbb{F}_{2^9}^\times$ such that $1 + a_2\{1, \theta\} \subseteq \mu_{73}(\mathbb{F}_{2^9})$. Noting that $1 + a_2 \cdot 0 = 1 \in \mu_{73}(\mathbb{F}_{2^9})$, we are done. □

4.4 Comparison with [CFL+13]

Here, we briefly explain why none of our results appear to capture the results by [CFL+13] in the $k = 2$ case. Their theorem says that if $m = p_1 p_2$ and $m + 1 = 2^d$ for a prime d, then there is an S_m^*-decoding polynomial of sparsity 3 over $\mathbb{F}_{2^d}$. By the equivalence we have shown, this is the same as saying that there is a root-of-unity grid over $\mathbb{F}_{2^d}$ with respect to m.

Firstly the rough-repunit/Weil-bound approach in Section 4.2 requires each p_i to be a very large divisor of $q^{d_i} - 1$ (at the very least, much larger than $q^{d_i/2}$). In the setting of [CFL+13], however, the two primes p_1, p_2 are the two prime factors of the same number $2^d - 1$. It is clearly not possible for both p_1 and p_2 to be larger than $2^{d/2}$, so the rough-repunit approach does not explain why this configuration gives rise to a root-of-unity grid.

The norm-kernel construction of Section 4.3 also does not recover the theorem of [CFL+13] in general. We saw that it explains Efremenko's example $m = 7 \cdot 73$ because $73 = (8^3 - 1)/(8 - 1)$ is prime (in addition to some other conditions). By contrast, in the theorem of [CFL+13], neither p_1 nor p_2 need arise as a prime of the form $(Q^e - 1)/(Q - 1)$ for a convenient intermediate field $\mathbb{F}_Q$. Thus the norm-kernel argument gives a conceptual explanation of Efremenko's original example, but not of the full family constructed in [CFL+13].

4.5 Stitching Together Simultaneous Root-of-Unity Grids

In some parameter regimes, in particular, when the number of servers is superconstant, the optimal thing to do will be to construct many simultaneous root-of-unity grids over the same characteristic and stitch them together into a sparse decoding polynomial.

Lemma 4.9 (Stitching together many root-of-unity grids to get a sparse decoding polynomial)**.** *Let $t \geq 0$, $r > 0$ and $2 \leq k_1 \leq \ldots \leq k_r$ be integers. Suppose there exists a prime q, integers $\{d_{i,j}\}_{i\in[r],j\in[k_i]} \cup \{d'_i\}_{i\in[t]}$, and pairwise distinct primes $\{p_{i,j}\}_{i\in[r],j\in[k_i]} \cup \{p'_i\}_{i\in[t]}$ such that the following conditions hold:*

- *For every $i \in [r]$, there is a $(\mathbb{F}_{q^{D_i}}, m_i)$-root-of-unity grid, where*

$$
\begin{aligned}
m_i &= \prod_{j\in[k_i]} p_{i,j}, \text{ and}\\
D_i &= \operatorname{lcm}(d_{i,1}, \ldots, d_{i,k_i}).
\end{aligned}
$$

- *For every $i \in [t]$, $p'_i | q^{d'_i} - 1$.*

Then there exists an S_m^*-decoding polynomial over $\mathbb{F}_{q^D}$ with sparsity $\leq 2^t(k_1+1)\ldots(k_r+1)$, where

$$
m := \left(\prod_{i=1}^{t} p'_i\right) \cdot \left(\prod_{i=1}^{r} m_i\right), \text{ and}
$$

$$D := \operatorname{lcm}(d'_1, \dots, d'_t, D_1, \dots, D_r).$$

Proof. We have all of the following:

- There is an $S^*_{p'_i}$-decoding polynomial over $\mathbb{F}_{q^{d'_i}}$ of sparsity 2 for all $i \in [t]$. This is by Claim 3.1.
- For every $i \in [r]$, there is an $S^*_{m_i}$-decoding polynomial over $\mathbb{F}_{q^{D_i}}$ of sparsity $k_i + 1$. This is by Lemma 4.2 applied to the given root-of-unity grids.

We now compose these decoding polynomials together using Theorem 3.6 to obtain the conclusion. □

We also note that when $r = 0$, one can instantiate the above construction without requiring any conjectures. This is exactly Corollary 3.7.

5 Number-Theoretic Conjectures for Rough Repunits

In this section, we state and justify two number-theoretic conjectures that imply the existence of root-of-unity grids via the rough repunit-approach. The two conjectures are for two different regimes: a constant number of servers, for which a single root-of-unity grid suffices (Section 5.1) and superconstant number of servers, for which the optimal construction of the decoding polynomials (see Section 6 for the optimization in the superconstant server regime) will be to stitch together several simultaneous root-of-unity grids occurring over the same characteristic (Section 5.2).

5.1 Constant # of Servers Regime

We begin by recalling Conjecture 1.1 that we rely on for the constant-server case, as well as our main theorem:

Conjecture 5.1 (Generalized repunit rough numbers, restated from Conjecture 1.1)**.** *For every positive integer $k > 1$, there exists a prime $q \geq k + 1$ for which there exist k distinct positive integers $d_1, \dots, d_k \geq 6$ such that $q^{d_i} - 1$ is $1/(3k)$-rough for all $i \in [k]$.*

Theorem 5.1 (Sparse decoding polynomials from Conjecture 5.1 (restated from Theorem 1.2))**.** *Assuming Conjecture 5.1, for any constant k there exists m that is a product of k distinct primes, a finite field $\mathbb{F}$ with $m \mid |\mathbb{F}| - 1$, and an S^*_m-decoding polynomial of sparsity $k + 1$.*

Proof. If $k = 1$, this is immediate from Claim 3.1. So assume from now that $k > 1$. Assuming Conjecture 5.1, there exist $q, d_1, \dots, d_k$ satisfying the preconditions of Corollary 4.6. The corollary gives us a finite field $\mathbb{F}$ and modulus m that is a product of k distinct primes such that there is an $(\mathbb{F}, m)$-root-of-unity grid. Then Lemma 4.2 tells us that there exists an S^*_m-decoding polynomial over $\mathbb{F}$ of sparsity $k + 1$, as desired. □

Now that we have proven Theorem 1.2, we turn our attention to justifying the plausibility of Conjecture 5.1.

Justification for Conjecture 5.1. As noted in the introduction, it would suffice to ensure that $d_i \geq 3k$ and $(q^{d_i}-1)/(q-1)$ is prime. The generalized repunit conjecture [Bou09, Dub93], a generalization of the Mersenne prime conjecture, states that there exist infinitely many base-q repunits for prime q. The conjecture we need is only weaker; we only need $(q^{d_i}-1)/(q-1)$ to be a "rough" number, or in other words, have a large prime divisor.

There are two different ways to justify that the above conjecture ought to be true:

- *Fix q then pick the d_i's.* As we have discussed, assuming the generalized repunit conjecture is enough to show that this approach will work. To justify why the generalized repunit conjecture might be true, since d is an odd prime there is no "obvious reason" preventing $\frac{q^d-1}{q-1}$ from being prime. So assuming $\frac{q^d-1}{q-1}$ is as likely to be prime as a random integer of the same size, the probability that a particular d works for us is $\sim \frac{1}{d \log q}$. Summing over all primes $d \geq 3k$, the number of good primes d is heuristically:
$$\approx \sum_{d \geq 3k \text{ prime}} \frac{1}{d \log q} = \frac{1}{\log q} \sum_{d \geq 3k \text{ prime}} \frac{1}{d},$$
which diverges to ∞ by Mertens's theorem.

- *Fix the d_i's then pick q.* An alternative strategy would be to pick arbitrary distinct primes $d_1, \ldots, d_k \geq 3k$, then conjecture that a good value of q should exist. Indeed, this is implied by Schinzel's hypothesis H [SS58, BH62] applied to the $k+1$ polynomials $X, \frac{X^{d_1}-1}{X-1}, \ldots, \frac{X^{d_k}-1}{X-1}$. Informally, this hypothesis states that for any tuple of irreducible polynomials, they should simultaneously evaluate to primes infinitely often unless this claim can be refuted by looking at the polynomials' evaluations modulo some fixed prime p.

  We take this opportunity to justify the applicability of Schinzel's hypothesis H: first, note that since the d_i's are prime these polynomials are all irreducible. Secondly, we need to rule out the existence of a prime p that divides one of these polynomials for every evaluation point x. We can take x to be a generator of the multiplicative group $\mathbb{F}_p^\times$, then we have $p \nmid x$ and for some d_i we have $p \mid x^{d_i} - 1 \Rightarrow p - 1 \mid d_i$, which forces $p - 1 \in \{1, d_i\}$. The latter case is impossible since p, d_i are both primes with $d_i \geq 6$, so they cannot be consecutive. So this leaves us with the former case of $p = 2$. However, it is easy to see that when x is odd all these polynomials are odd (since the d_i's are odd), so this case cannot happen either.

  So if Schinzel's hypothesis H were true, it would follow that there are infinitely many positive integers q such that $q, \frac{q^{d_1}-1}{q-1}, \ldots, \frac{q^{d_k}-1}{q-1}$ are all prime, as claimed.

5.2 Superconstant # of Servers Regime

To optimize the constructions of S-decoding polynomials for the superconstant server regime, it turns out to be optimal to stitch together many root-of-unity grids over the same characteristic. To carry out this optimization and estimate costs of the final PIR scheme, we need to estimate the parameters (such as m, q, field size) of the decoding polynomial obtained via this stitching. The following conjecture does this.

Conjecture 5.2. *Let $t \geq 0, r > 0$ and $2 \leq k_1 \leq \ldots \leq k_r$ be integers. Then there exists a prime q, integers $\{d_{i,j}\}_{i \in [r], j \in [k_i]} \cup \{d'_i\}_{i \in [t]}$, and pairwise distinct primes $\{p_{i,j}\}_{i \in [r], j \in [k_i]} \cup \{p'_i\}_{i \in [t]}$ such that the following conditions hold:*

- *For every $i \in [r]$, there is a $(\mathbb{F}_{q^{D_i}}, m_i)$-root-of-unity grid, where*

$$m_i = \prod_{j \in [k_i]} p_{i,j}, \text{ and}$$
$$D_i = \operatorname{lcm}(d_{i,1}, \ldots, d_{i,k_i}).$$

- *For every $i \in [t]$, $p'_i | q^{d'_i} - 1$.*

Moreover, we can upper bound the prime q, modulus $m := \left(\prod_{i=1}^{r} \prod_{j=1}^{k_i} p_{i,j}\right) \cdot \left(\prod_{i \in [t]} p'_i\right)$ and size of the field $\mathbb{F}_{q^D}$, where $D := \operatorname{lcm}(D_1, \ldots, D_r, d'_1, \ldots, d'_t)$ as follows:

$$q \leq 2k_r,$$
$$m \leq \exp\left\{ O\left(t \log(t+K) + (\log k_r) \cdot \sum_{i=1}^{r} k_i^2 K_i \log K_i \right) \right\}, \text{ and}$$
$$\log |\mathbb{F}_{q^D}| \leq (\log k_r) \cdot \left(\prod_{i=1}^{r} O(k_i K_i \log K_i)^{k_i} \right) \cdot \exp(O(t \log(t+K))),$$

where $K_i := \sum_{j=i}^{r} k_j$ and $K = K_1 = \sum_{j=1}^{r} k_j$.

Remark. *We note that although we use the full generality of Conjecture 5.2 to find the optimal choice of decoding polynomials, our final PIR constructions require the conjecture to hold only for some choices of $t, r, k_1, \ldots, k_r$. In particular, we only ever consider cases where $t = 0$ and all the $k_1 = \ldots = k_r$.*

Our strategy for finding r small root-of-unity grids, given $k_1, \ldots, k_r$ will be to first pick a characteristic q, and then search for *prime* degrees $d_{i,j}$ such that $q^{d_{i,j}} - 1$ is $(1/(3k_i))$-rough, for all $i \in [r], j \in [k_i]$—we call such degrees *usable*. We note that conversion of rough repunits $q^{d_{i,j}} - 1$ to root-of-unity grids does not require the $d_{i,j}$'s to be prime; in the following conjecture we insist on prime degrees because it is easier to reason about their density. Having picked enough usable degrees $d_{i,j}$, we set $p_{i,j} = P^+(q^{d_{i,j}} - 1)$, and the rest of the root-of-unity grid follows from these choices. Then, estimating the expected size of m, q, D boils down to estimating the density of usable primes.

Conjecture 5.3 (Density of Usable Degrees)**.** *For every $\varepsilon \in (0, 1/6]$, prime q and integer $\Delta \geq 6$, such that $q^{5\Delta/12} > \Delta^3 \log q$.*

$$\Pr_{\substack{\text{prime } d \\ d \leftarrow [\Delta, 2\Delta]}} [q^d - 1 \text{ is } \varepsilon\text{-rough}] \geq \log \frac{1}{1-\varepsilon} + O\left(\frac{1}{\Delta}\right).$$

The purpose of the above conjecture is to give us many root-of-unity grids over the same characteristic that can be stitched together into one large decoding polynomial. We capture this with the below lemma:

Theorem 5.2. *Assume Conjecture 5.3 holds. Then, Conjecture 5.2 holds, and therefore for every choice of $t \geq 0, r > 0$ and $2 \leq k_1 \leq \ldots \leq k_r$, there exists an S^*_m-decoding polynomial over $\mathbb{F}_{q^D}$ with sparsity $\leq 2^t (k_1 + 1) \cdots (k_r + 1)$, where*

$$q \leq 2k_r,$$
$$m \leq \exp\left\{ O\left(t \log(t+K) + (\log k_r) \cdot \sum_{i=1}^{r} k_i^2 K_i \log K_i \right) \right\}, \text{ and}$$

$$\log |\mathbb{F}_{q^D}| \le (\log k_r) \cdot \left(\prod_{i=1}^{r} O(k_i K_i \log K_i)^{k_i} \right) \cdot \exp(O(t \log(t+K))),$$

where $K_i := \sum_{j=i}^{r} k_j$ *and* $K = K_1 = \sum_{j=1}^{r} k_j$.

Proof. Fix any prime $q > k_r$. By Lemma 3.23, there exists prime $k_r \le q \le 2k_r$. Call prime $d \ge 6$ *usable for level* i if $q^d - 1$ is ε_i rough, where $\varepsilon_i = 1/(3k_i)$. For each i, we will pick k_i usable primes $d_{i,1}, \dots, d_{i,k_i}$ such that $\{d_{i,j}\}_{i,j}$ are pairwise distinct. Therefore, by Corollary 4.6, there exists a $(\mathbb{F}_{q^{D_i}}, m_i)$-root-of-unity grid, where $D = \operatorname{lcm}(d_{i,1}, \dots, d_{i,k_i})$ and $m_i = p_{i,1} \cdots p_{i,k_i}$, where $p_{i,j} = P^+(q^{d_{i,j}} - 1)$.

By Conjecture 5.3, we have that for random prime $d \leftarrow [\Delta_i, 2\Delta_i]$, large enough[6] $\Delta_i \ge 6, k_i, \log\log q$, we have that d is usable for level i with probability around

$$\log \frac{1}{1-\varepsilon_i} + O\left(\frac{1}{\Delta_i}\right) = \varepsilon_i + O(\varepsilon_i^2) + O\left(\frac{1}{\Delta_i}\right) = \Omega\left(\frac{1}{k_i}\right) + O\left(\frac{1}{k_i^2}\right).$$

To get ℓ usable primes greater than Δ_i, we need to search over roughly $O(k_i \ell)$ primes greater than Δ_i. To make sure that the interval $[\Delta_i, 2\Delta_i]$ contains enough primes, we would like that $\pi(2\Delta_i) - \pi(\Delta_i) \ge \Omega(k_i \ell)$ for all $i \in [r]$. By the prime number theorem,

$$\pi(2\Delta_i) - \pi(\Delta_i) = \Theta\left(\frac{2\Delta_i}{\log 2\Delta_i} - \frac{\Delta_i}{\log \Delta_i}\right),$$

so taking $\Delta_i = O(k_i \ell \log(k_i \ell))$ suffices, and we expect to find $d_{i,1}, \dots, d_{i,k_i} \le O(k_i \ell \log(k_i \ell))$.

Since $\varepsilon_1 \ge \dots \ge \varepsilon_r$, and ε-roughness is a stronger condition for smaller ε, we have every degree usable at level i must be usable for all levels $i' \le i$. Therefore to ensure that all the $d_{i,j}$'s are distinct, let us start picking from $i = r$ and move downwards. We first pick k_r distinct prime degrees to be usable at level r, so we expect $d_{r,1}, \dots, d_{r,k_r} \le O(k_r^2 \cdot \log(k_r^2))$. To pick distinct usable degrees at level i, we want to avoid all $K_{i+1} = K_i - k_i$ degrees picked so far, and pick k_i additional degrees usable at level i, so we expect to find $d_{i,1}, \dots, d_{i,k_i} \le O(k_i K_i \cdot \log(k_i K_i))$. Since $\log(k_i K_i) \le 2 \log K_i$ for all $i \in [r]$, we get the slightly simpler bound

$$d_{i,1}, \dots, d_{i,k_i} \le O(k_i K_i \log(K_i)).$$

After having picked all the primes $p_{i,j}$ for $i \in [r], j \in [k_i]$, we look for t additional small primes p'_i for $i \in [t]$. Keep in mind that all the primes $\{p_{i,j}\}_{i \in [r], j \in [k_i]} \cup \{p'_i\}_{i \in [t]}$ must be pairwise distinct. To avoid picking the same prime twice, we will simply look for the smallest t primes that do not already occur in the set $\{p_{i,j}\}_{i \in [r], j \in [k_i]} \cup \{q\}$ of size $K + 1$. By the prime number theorem, we expect $p'_t \le O((t+K)\log(t+K))$. Now all that remains to do is to find d'_i, $i \in [t]$ such that $p'_i | q^{d'_i} - 1$ for all $i \in [t]$. This can be done by picking $d'_i = \operatorname{ord}_{p'_i}(q) \le p'_i$ for all $i \in [t]$.

Finally, we can bound the modulus m:

$$\begin{aligned} m &= \left(\prod_{i=1}^{t} p'_i \right) \cdot \left(\prod_{i=1}^{r} \prod_{i=1}^{k_i} p_{i,j} \right) \\ &= O((t+K)\log(t+K))^t \cdot q^{\sum_{i=1}^{r} k_i \cdot O(k_i K_i \log(K_i))} \\ &= \exp\left\{ O\left(t \log(t+K) + (\log q) \cdot \sum_{i=1}^{r} k_i^2 K_i \log(K_i) \right) \right\} \end{aligned}$$

[6] Note that the precondition for Conjecture 5.3 that $q^{5\Delta/12} \gg \Delta^3 \log q$ is satsified as long as $\Delta \ge \Omega(\log\log q)$.

$$\leq \exp\left\{ O\left(t\log(t+K) + (\log k_r)\cdot \sum_{i=1}^{r} k_i^2 K_i \log(K_i) \right) \right\}$$

and the logarithm of the field size

$$\begin{aligned}
\log|\mathbb{F}_{q^D}| &= (\log q)\cdot D \\
&= (\log q)\cdot \operatorname{lcm}(D_1,\dots D_r, d_1',\dots,d_t') \\
&\leq (\log q)\cdot \left(\prod_{i=1}^{t} d_i'\right)\cdot\left(\prod_{i=1}^{r}\prod_{i=1}^{k_i} d_{i,j}\right) \\
&\leq O(\log k_r)\cdot O((t+K)\log(t+K))^t \cdot \left(\prod_{i=1}^{r} O(k_i K_i \log(K_i))^{k_i}\right) \\
&\leq O(\log k_r)\cdot \exp(O(t\log(t+K)))\cdot \left(\prod_{i=1}^{r} O(k_i K_i \log(K_i))^{k_i}\right).
\end{aligned}$$

□

5.3 Heuristics Justifying Conjecture 5.3

We present two heuristics (Heuristics 5.4 and 5.5) that each individually support Conjecture 5.3. First, we introduce some useful notation.

Definition 5.3. *Fix $\varepsilon \in (0, 1/6]$. Throughout this section, when ε is clear from the context, we use the following notation:*

$$\tau_{q,d} := q^{d(1-\varepsilon)} \geq (q^d-1)^{(1-\varepsilon)}, \text{ and}$$
$$N_{q,d} := \begin{cases} \frac{q^d-1}{d(q-1)} & \text{if } d|q-1, \\ \frac{q^d-1}{q-1} & \text{otherwise.} \end{cases}$$

Claim 5.1. *For any $\varepsilon \in (0, 1/6]$ and primes q, d, we have that $P^+(q^d-1) \geq \tau_{q,d}$ if and only if $P^+(N_{q,d}) \geq \tau_{q,d}$.*

Proof. First note that $N_{q,d}$ is an integer: $q-1|q^d-1$ for all $d \geq 1$, and $d|q-1$, then $q = 1 \bmod d$

$$\frac{q^d-1}{q-1} = 1 + q + \cdots + q^{d-1} = 0 \bmod d.$$

We will show that every prime divisor of $(q^d-1)/N_{q,d}$ is strictly smaller than $\tau_{q,d}$. Since $\varepsilon \leq 1/6$, $q, d \geq 2$, we have that $\tau_{q,d} = q^{d(1-\varepsilon)} \geq q^{5d/6} > q-1$. Also $\tau_{q,d} \geq q^{5d/6} \geq 2^{5d/6} > d$ for all $d \geq 2$. It then follows that q^d-1 has a prime factor at least $\tau_{q,d}$ if and only if $N_{q,d}$ does. □

5.3.1 Naive Dickman Random Factorization Heuristics

Dickman characterized the probability that a random number drawn from an interval is smooth (Lemma 3.28). As a first heuristic, we model the repunit integers q^d-1 as random integers of size q^d, and apply the Dickman heuristic to these numbers.

Heuristic 5.4. *[Dickman heuristic for $q^d - 1$] For every prime q, $\alpha > 0$, and integer Δ,*

$$\Pr_{\substack{\text{prime } d \\ d \leftarrow [\Delta, 2\Delta]}} [P^+(q^d - 1) \leq q^{\alpha d}] = \rho\left(\frac{1}{\alpha}\right) + O\left(\frac{1}{\Delta}\right).$$

Lemma 5.4. *Under Heuristic 5.4, Conjecture 5.3 holds.*

Proof. Since $\tau_{q,d} := q^{(1-\varepsilon)d} \geq (q^d - 1)^{1-\varepsilon}$, we have that that $x \geq \tau_{q,d} \implies x \geq (q^d - 1)^{1-\varepsilon}$, so applying $x = q^d - 1$ this means that

$$\begin{aligned}
\Pr_d[q^d - 1 \text{ is } \varepsilon\text{-rough}] &\geq \Pr_d[P^+(q^d - 1) \geq \tau_{q,d}] \\
&= \Pr_d[P^+(q^d - 1) \geq q^{d(1-\varepsilon)}] && \text{Notation 5.3} \\
&= 1 - \rho\left(\frac{1}{1-\varepsilon}\right) + \left(\frac{1}{\Delta}\right) && \text{Heuristic 5.4} \\
&= \log\frac{1}{1-\varepsilon} + \left(\frac{1}{\Delta}\right) && \text{Lemma 3.28, since } 5/6 \leq \alpha \leq 1.
\end{aligned}$$

□

Justification of Heuristic 5.4. For a random (not necessarily prime) integer $d \leftarrow [\Delta, 2\Delta]$, modeling $q^d - 1$ as a random integer of size q^d ignores the cyclotomic structure of such numbers:

$$q^d - 1 = \prod_{d'|d} \Phi_{d'}(q),$$

where $\Phi_{d'}(q)$ is the value of the d'th cyclotomic polynomial, evaluated at q (see Fact 3.22). For example, if d is even, every prime factor of $q^d - 1$ must divide at least one of $q^{d/2} - 1$, $q^{d/2} + 1$, ruling out any hope for $q^d - 1$ to be rough. In general, for non-prime d, and in the extreme case where d has a lot of (small) prime factors, one would expect to run into similar issues. This is why we assume d is prime in our heuristic; Claim 5.1 shows that we are then effectively treating $N_{q,d}$ as a random number that satisfies the Dickman heuristic.

5.3.2 Heuristics Accounting for the Cyclotomic Structure of $q^d - 1$

In order to sidestep the cyclotomic structure of $q^d - 1$ and use the Dickman heuristic for random integers, the previous section and Heuristic 5.4 restricted attention to prime d, effectively treating $N_{q,d}$ as a random integer that satisfies the Dickman heuristic, but this still ignores the fact that $N_{q,d}$ itself has structure that could make it look very different from a random number.

Concretely, by Lemma 3.26 and Claim 5.1, we know that for rough $q^d - 1$, $p = P^+(N_{q,d})$ if and only if $\mathrm{ord}_p(q) = d$, which implies that $d|p - 1$. In other words, $p = 1 \bmod d$. Therefore, the prime divisors of $N_{q,d}$ come from a sparse arithmetic progression of primes of density $1/\varphi(d)$. In this section, instead of heuristically treating $N_{q,d}$ as a random number, we use heuristically estimate the density of large primes p such that $p|N_{q,d}$.

Heuristic 5.5. *Let $\varepsilon \in (0, 1/6]$ and let q be a prime and let Δ be an integer*

$$\mathop{\mathbb{E}}_{d} \sum_{\substack{\text{primes} \\ p \in [\tau_{q,d}, N_{q,d}]}} \mathbb{1}[p | N_{q,d}] = \mathop{\mathbb{E}}_{d}\left[\sum_{\substack{\text{primes} \\ p \in [\tau_{q,d}, N_{q,d}] \\ p = 1 \bmod d}} \frac{d-1}{p-1} \right] + O\left(\frac{1}{\Delta}\right),$$

where the expectation is taken over random prime $d \leftarrow [\Delta, 2\Delta]$.

Justification for Heuristic 5.5. For $p \geq \tau_{q,d}$, we know from the proof of Lemma 3.27 that $p | N_{q,d}$ if and only if $\mathrm{ord}_p(q) = d$. This can happen only if $d | p-1$, i.e., $p = 1 \bmod d$. Since $\mathbb{F}_p^\times$ is cyclic, the number of elements if order d is $\varphi(d) = d-1$. If q was sampled uniformly at random from $\mathbb{F}_p^\times$, conditioned on $p = 1 \bmod d$, the probability that $\mathrm{ord}_p(q) = d$ is $(d-1)/(p-1)$. Summing this density and averaging over d yields the expression in the heuristic. In our setting, q is a fixed prime. Heuristic 5.5 asserts that this quantity nevertheless gives a good estimate in the fixed-q case.

Lemma 5.5. *Assume that the GRH for all L-functions. Under Heuristic 5.5, Conjecture 5.3 holds.*

Proof. Throughout this proof, we are sampling d to be a random prime in the interval $[\Delta, 2\Delta]$. First, we claim that $q^d - 1$ is ε-rough if and only $P^+(N_{q,d}) \geq \tau_{q,d}$.

Since $\tau_{q,d} \geq (q^d-1)^{1-\varepsilon}$, we have that $P^+(q^d-1) \geq \tau_{q,d}$ implies that $q^d - 1$ is ε-rough.

$$\begin{aligned}
\Pr_d[q^d - 1 \text{ is } \varepsilon\text{-rough}] &\geq \Pr_d[P^+(q^d-1) \geq \tau_{q,d}] \\
&= \Pr_d[P^+(N_{q,d}) \geq \tau_{q,d}] && \text{Claim 5.1} \\
&= \mathop{\mathbb{E}}_d \#\{\text{prime } p | N_{q,d} \,:\, p \geq \tau_{q,d}\},
\end{aligned}$$

where the last equality holds because $\#\{\text{prime } p | N_{q,d} \,:\, p \geq \tau_{q,d}\} \leq 1$. To see, this, suppose for contradiction there were two distinct primes $p_1, p_2 | N_{q,d}$ such that $p_1, p_2 \geq \tau_{q,d}$, we would have that

$$q^d \geq N_{q,d} \geq p_1 p_2 \geq \tau_{q,d}^2 = q^{2d(1-\varepsilon)} \geq q^{5d/3},$$

a contradiction. We can rewrite $\#\{\text{prime } p | N_{q,d} \,:\, p \geq \tau_{q,d}\}$ as a sum of an indicator over primes p in the interval

$$\begin{aligned}
\Pr_d[q^d - 1 \text{ is } \varepsilon\text{-rough}] &\geq \mathop{\mathbb{E}}_d \#\{\text{prime } p | N_{q,d} \,:\, p \geq \tau_{q,d}\} \\
&= \mathop{\mathbb{E}}_d \sum_{\substack{\text{primes} \\ p \in [\tau_{q,d}, N_{q,d}]}} \mathbb{1}[p | N_{q,d}] \\
&= \mathop{\mathbb{E}}_d \sum_{\substack{\text{primes} \\ p \in [\tau_{q,d}, N_{q,d}] \\ p = 1 \bmod d}} \frac{d-1}{p-1} + O\left(\frac{1}{\Delta}\right) && \text{Heuristic 5.5} \\
&= \mathop{\mathbb{E}}_d \sum_{\substack{\text{primes} \\ p \in [\tau_{q,d}, N_{q,d}] \\ p = 1 \bmod d}} \frac{d-1}{p}\left(1 + \frac{1}{(p-1)}\right) + O\left(\frac{1}{\Delta}\right)
\end{aligned}$$

$$= \mathop{\mathbb{E}}_{d} \sum_{\substack{\text{primes} \\ p \in [\tau_{q,d}, N_{q,d}] \\ p = 1 \bmod d}} \frac{d-1}{p} \left(1 + O\left(\frac{1}{\tau_{q,d}}\right)\right) + O\left(\frac{1}{\Delta}\right).$$

By Lemma 3.29, we can lower bound this quantity, and get that

$$\begin{aligned}
&\Pr_{d}[q^d - 1 \text{ is } \varepsilon\text{-rough}] \\
&\quad \geq \mathop{\mathbb{E}}_{d}\left[(d-1) \cdot \left(\frac{1}{\varphi(d)}(\log\log N_{q,d} - \log\log \tau_{q,d}) + O\left(\frac{\log \tau_{q,d}}{\sqrt{\tau_{q,d}}}\right)\right) \cdot \left(1 + O\left(\frac{1}{\tau_{q,d}}\right)\right)\right] + O\left(\frac{1}{\Delta}\right) \\
&\quad = \mathop{\mathbb{E}}_{d}\left[\log \frac{\log N_{q,d}}{\log \tau_{q,d}} + O\left(\frac{d \log \tau_{q,d}}{\sqrt{\tau_{q,d}}} + \frac{d}{\tau_{q,d}}\right)\right] + O\left(\frac{1}{\Delta}\right) \\
&\quad = \log \frac{1}{1-\varepsilon} + O\left(\frac{1}{\Delta}\right) + O\left(\frac{d \log \tau_{q,d}}{\sqrt{\tau_{q,d}}}\right) + O\left(\frac{1}{\Delta}\right),
\end{aligned}$$

where the second equality holds because d is an odd prime, so $\varphi(d) = d - 1$, and the third equality holds because of the following calculation:

- If $d \nmid q - 1$,

$$\begin{aligned}
\frac{\log N_{q,d}}{\log \tau_{q,d}} = \frac{\log \frac{q^d(1-1/q^d)}{q-1}}{\log q^{d(1-\varepsilon)}} &= \frac{d \log q + \log(1 - 1/q^d) - \log(q-1)}{d(1-\varepsilon)\log q} \geq \frac{d \log q + O(1) + O(\log q)}{d(1-\varepsilon)\log q} \\
&= \frac{1}{1-\varepsilon}\left(1 + O\left(\frac{1}{d}\right)\right)
\end{aligned}$$

- If $d | q - 1$, we have that $d \leq q - 1$, and so $\log d \leq \log q$, and an almost identical calculation as above shows that

$$\frac{\log N_{q,d}}{\log \tau_{q,d}} \geq \frac{1}{1-\varepsilon}\left(1 + O\left(\frac{1}{d}\right)\right).$$

In both cases, since $\log N_{q,d} / \log \tau_{q,d} = \frac{1}{1-\varepsilon}\left(1 + O\left(\frac{1}{d}\right)\right)$, we have that

$$\log \frac{\log N_{q,d}}{\log \tau_{q,d}} = \log \frac{1}{1-\varepsilon} + \log\left(1 + O\left(\frac{1}{d}\right)\right) = \log \frac{1}{1-\varepsilon} + O\left(\frac{1}{d}\right) = \log \frac{1}{1-\varepsilon} + O\left(\frac{1}{\Delta}\right).$$

Finally, observing that $2d(\log \tau_{q,d})/\sqrt{\tau_{q,d}} \leq O(d(\log \tau_{q,d})q^{-5d/12}) \ll O(1/\Delta)$ whenever $q^{5\Delta/12} > \Delta^3 \log q$ finishes the proof. □

6 PIR from Our Decoding Polynomials

We begin by quickly completing the proof of Corollary 1.4 (PIR in the constant-server regime):

Proof of Corollary 1.4. Let $k = s - 1$. By Theorem 1.2, there exists m that is a product of k distinct primes and a finite field $\mathbb{F}$ with $m \mid |\mathbb{F}| - 1$ such that there is an S_m^*-decoding polynomial over $\mathbb{F}$ of sparsity $k + 1 = s$. The conclusion is now immediate by plugging into Theorem 1.1. □

Now we turn to the case of superconstant s. Here, Theorem 1.1 does not apply out of the box so we will use the following workflow to compute bounds on the PIR communication complexity:

- We invoke our root-of-unity grid framework to obtain $q, m, \mathbb{F}$ with bounds on these quantities such that there exists an S_m^*-decoding polynomial over $\mathbb{F}$ of sparsity $\leq s$.

  This part splits into three regimes (chosen according to what will give the best end guarantee for the communication complexity):

  - Section 6.1: when $s \leq O((\log\log n)^{1/3}/(\log\log\log n)^{2/3})$, we will use Conjecture 5.2 and Lemma 4.9 with $t = 0$, $r = 1$, and $k_1 = s - 1$. In other words, we build our decoding polynomial from just one root-of-unity grid. This is intuitive as this is what we do in the constant-server case.
  - Section 6.2: when s ranges from $\Omega((\log\log n)^{1/3}/(\log\log\log n)^{2/3})$ to $\exp(o(\sqrt{\log\log n/\log\log\log n}))$, we will use multiple root-of-unity grids of the same size k_i, but this size will be smaller than k. (This still goes via Conjecture 5.2 and Lemma 4.9.)
  - Section 6.3: finally, when $s \geq \exp(\Omega(\sqrt{\log\log n/\log\log\log n}))$, our techniques combined with Conjecture 5.2 do not provide any advantage over using the standard construction of decoding polynomials. For completeness, we have restated the relevant decoding polynomial result in Corollary 3.7 and derived the resulting communication guarantee in Theorem 3.9.

- Once we have bounds on $q, m, |\mathbb{F}|$, we can feed these into Corollary 3.4 [BBR94, Gro00, DGY11] to bound the dimension d of a matching vector family modulo qm of size $\geq n$.
- Finally, we can plug these into Theorem 3.8 to obtain a bound on the PIR communication complexity CC.

6.1 Small Superconstant-Server Regime

Theorem 6.1 (PIR in the few-server regime)**.** *Assume Conjecture 5.2, and suppose n, s are such that $s = \omega(1)$ and:*

$$s \leq O((\log\log n)^{1/3}/(\log\log\log n)^{2/3}).$$

Then there exists an s-server PIR for a database of size n achieving total communication:

$$\mathsf{CC} = \exp\left(\exp\left(O\left(\frac{\log\log n}{s}\right)\right)\right).$$

Proof. We begin by setting parameters and calculating the resulting bounds on $m, |\mathbb{F}|$. We invoke Lemma 4.9 with $t = 0$, $r = 1$, and $k_1 = s - 1$. The resulting decoding polynomial has sparsity $\leq 2^t(k_1 + 1) = s$. The bounds from the lemma tell us that:

$$\begin{aligned} q &\leq O(s), \\ m &\leq \exp\left(O\left((\log q) \cdot (s-1)^3 \log(s-1)\right)\right) \\ &= \exp(O(s^3(\log s)^2)), \text{ and} \\ \log|\mathbb{F}| &\leq (\log q) \cdot O((s-1)^2 \log(s-1))^{s-1} \\ &\leq O(1)^s \cdot s^{2s} \cdot (\log s)^s. \end{aligned}$$

In Corollary 3.4, we will take the modulus to be $qm \leq \exp(O(s^3(\log s)^2))$ and the number of primes ℓ is exactly s (accounting for the extra prime we get from q as well).

Claim 6.1. *We have:*

$$\exp(\Omega(s^2)) \le \exp((s-1)\log\log n - O(s^3(\log s)^2)) \le \eta \le \exp(s\log\log n).$$

Proof. We have:

$$\begin{aligned}\eta &= \frac{(\log n)^{\ell-1}}{qm} \\ &\ge \frac{(\log n)^{s-1}}{\exp(O(s^3(\log s)^2))} \\ &= \exp\left((s-1)\log\log n - O(s^3(\log s)^2)\right),\end{aligned}$$

and $\eta \le (\log n)^{s-1} \le \exp(s\log\log n)$. The leftmost inequality follows from the fact that $s\log\log n \ge \tilde{\Omega}(s^4)$. □

Claim 6.2. *We have:*

$$d \le \exp\left(\exp\left(O\left(\frac{\log\log n}{s}\right)\right)\right).$$

Proof. By Claim 6.1, we will use the second branch of Corollary 3.4 and compute:

$$\begin{aligned}d &= \exp\left(O\left(\log n \cdot \frac{(\log \eta)^{1-1/\ell}}{\ell \eta^{1/\ell}}\right)\right) \\ &\le \exp\left(O\left(\log n \cdot \frac{(s\log\log n)^{1-1/s}}{\ell \eta^{1/\ell}}\right)\right) \\ &\le \exp\left(O\left(\log n \cdot \frac{(s\log\log n)^{1-1/s}}{s\cdot\exp((1-1/s)\log\log n - O(s^2(\log s)^2))}\right)\right) \\ &= \exp\left(2^{O(s^2(\log s)^2)} \cdot (\log n)^{1/s}(\log\log n)^{1-1/s}\right) \\ &= \exp\left(\exp\left(O(s^2(\log s)^2) + \frac{\log\log n}{s} + \left(1-\frac{1}{s}\right)\log\log\log n\right)\right) \\ &= \exp\left(\exp\left(O\left(\frac{\log\log n}{s}\right)\right)\right).\end{aligned}$$

In the final step, we are using the fact that $s \le O((\log\log n)^{1/3}/(\log\log\log n)^{2/3})$ to absorb the $O(s^2(\log s)^2)$ term. □

Now to complete the proof and bound CC, by Theorem 3.8 it remains to verify that:

$$s\log|\mathbb{F}| \le \exp\left(\exp\left(O\left(\frac{\log\log n}{s}\right)\right)\right).$$

From the earlier computations, we have:

$$\begin{aligned}s\log|\mathbb{F}| &\le O(1)^s \cdot s^{2s} \cdot (\log s)^s \\ &\le \exp(O(s\log s)) \\ &= \exp(\exp(O(\log s))),\end{aligned}$$

and now the conclusion follows by noting that $s\log s \le \tilde{O}((\log\log n)^{1/3}) \le O(\log\log n)$. □

6.2 Intermediate-Server Regime

Theorem 6.2 (PIR in the intermediate-server regime)**.** *Assume Conjecture 5.2, and suppose s is such that:*

$$\Omega\left(\frac{(\log\log n)^{1/3}}{(\log\log\log n)^{2/3}}\right) \leq s \leq \exp\left(o\left(\sqrt{\frac{\log\log n}{\log\log\log n}}\right)\right).$$

Then there exists an s-server PIR for a database of size n achieving total communication:

$$\exp\left(\exp\left(O\left(\frac{(\log\log n)^{2/3}(\log\log\log n)^{1/3}}{(\log s)^{1/3}} \cdot \left(\log\left(\frac{\log\log n}{(\log\log\log n)\cdot(\log s)^2}\right)\right)^{2/3}\right)\right)\right).$$

Proof. We begin by setting parameters and calculating the resulting bounds on $m, |\mathbb{F}|$. Define

$$x = \frac{\log\log n}{(\log\log\log n)\cdot(\log s)^2} = \omega(1),$$

and $a = \Theta((x\log x)^{1/3})$. We will invoke Lemma 4.9 with $t = 0$, $r = \lfloor \frac{\log s}{\log(a+1)} \rfloor$, and $k_1 = \ldots = k_r = a$. The resulting decoding polynomial has sparsity $2^t \prod_{i=1}^{r}(k_i+1) = (a+1)^r \leq s$. As an additional sanity check, we verify that $a + 1 \leq s$; indeed, we have:

$$\begin{aligned}
&a + 1 \leq s \\
\Leftarrow\ &(x\log x)^{1/3} \leq O(s) \\
\Leftrightarrow\ &x \leq O(s^3/\log s) \\
\Leftrightarrow\ &\frac{\log\log n}{\log\log\log n} \leq O(s^3 \log s) \\
\Leftrightarrow\ &\Omega\left(\frac{(\log\log n)^{1/3}}{(\log\log\log n)^{2/3}}\right) \leq s,
\end{aligned}$$

which we are assuming to be the case. Continuing on, we have $\ell = ar + 1$ and $K = ar$. For any $i \in [r]$, we have $K_i = a(r - i + 1)$.

Claim 6.3. *We have:*

$$\begin{aligned}
q &\leq 2a \\
m &\leq \exp(O(a^3 r^2 \log a \log\log\log n)) \\
\log|\mathbb{F}| &\leq \exp(O(ar\log\log\log n)).
\end{aligned}$$

Proof. By the bounds stated in Lemma 4.9, we have:

$$\begin{aligned}
&q \leq 2a, \\
&m \leq \exp\left(O\left((\log q)\cdot\sum_{i=1}^{r} k_i^2 K_i \log K_i\right)\right) \\
&\quad = \exp\left(O\left((\log q)\cdot\sum_{i=1}^{r} a^2 \cdot a(r-i+1)\cdot(\log(a(r-i+1)))\right)\right)
\end{aligned}$$

$$= \exp\left(O\left(a^3 \log a \cdot \left(\sum_{i=1}^{r}(r-i+1)\cdot(\log(a(r-i+1)))\right)\right)\right)$$
$$= \exp\left(O\left(a^3 \log a \cdot \left(r^2 \log a + r^2 \log r\right)\right)\right), \text{ and}$$
$$\log|\mathbb{F}| \leq (\log q)\cdot\left(\prod_{i=1}^{r} O(k_i K_i \log K_i)^{k_i}\right)$$
$$\leq (\log q)\cdot\prod_{i=1}^{r} O(a^2(r-i+1)\log(a(r-i+1)))^a$$
$$= (\log q)\cdot\exp\left(O\left(a\sum_{i=1}^{r}(O(1)+\log a+\log(r-i+1)+\log\log(a(r-i+1)))\right)\right)$$
$$= \exp\left(O\left(ar(\log a+\log r)\right)\right).$$

To complete the proof of the claim, it suffices to show that $\log a + \log r \leq O(\log\log\log n)$. Indeed, we have:

$$\begin{aligned}\log a &= \Theta(\log x)\\ &\leq O(\log\log\log n), \text{ and}\\ \log r &\leq \log\log s\\ &\leq O(\log\log\log n).\end{aligned}$$

□

Claim 6.4. *We have:*

$$\exp(\Omega(\ell^2)) \leq \exp(ar\log\log n - O(a^3r^2\log a\cdot\log\log\log n)) \leq \eta \leq \exp(ar\log\log n),$$

and $ar \leq o(\log\log n)$.

Proof. First we lower bound η:

$$\begin{aligned}\eta &= \frac{(\log n)^{\ell-1}}{qm}\\ &\geq \frac{(\log n)^{ar}}{\exp(O(a^3r^2\log a\cdot\log\log\log n))}\\ &= \exp(ar\log\log n - O(a^3r^2\log a\cdot\log\log\log n)).\end{aligned}$$

The upper bound on η is straightforward: $\eta \leq (\log n)^{\ell-1} = \exp(ar\log\log n)$. To establish the leftmost inequality, note that $\ell^2 = a^2r^2$, so it suffices to show that:

$$\begin{aligned}&ar\log\log n \geq \omega(a^3r^2\log a(\log a+\log r))\\ \Leftrightarrow\ &\log\log n \geq \omega(a^2r\log a(\log a+\log r))\\ \Leftrightarrow\ &\log\log n \geq \omega(a^2\log s(\log a+\log r)).\end{aligned}$$

To check this, note that:

$$a^2\log s(\log a+\log r) \leq O\left(x^{2/3}(\log x)^{2/3}\cdot\log s\cdot(\log x+\log\log s)\right)$$

$$\leq O\left(\frac{(\log\log n)^{2/3}}{(\log s)^{1/3}} \cdot (\log\log\log n + \log\log s)\right)$$
$$\leq o(\log\log n),$$

as claimed. Lastly, we have already shown that $a^2r^2 \leq o(ar\log\log n) \Rightarrow ar \leq o(\log\log n)$. □

Claim 6.5. *We have:*

$$d \leq \exp\left(\exp\left(O\left(\frac{(\log\log n)^{2/3}(\log\log\log n)^{1/3}}{(\log s)^{1/3}} \cdot \left(\log\left(\frac{\log\log n}{(\log\log\log n)\cdot(\log s)^2}\right)\right)^{2/3}\right)\right)\right).$$

Proof. By Claim 6.4, we can use the second branch of Corollary 3.4 and compute:

$$\begin{aligned}
d &= \exp\left(O\left(\log n \cdot \frac{(\log\eta)^{1-1/\ell}}{\ell\eta^{1/\ell}}\right)\right) \\
&\leq \exp\left(O\left(\log n \cdot \frac{(ar\log\log n)^{1-1/\ell}}{\ell\eta^{1/\ell}}\right)\right) \\
&\leq \exp\left(O\left(\log n \cdot \frac{(ar\log\log n)^{1-1/\ell}}{ar\exp\left((1-\frac{1}{ar+1})\log\log n - O(a^2r\log a(\log\log\log n))\right)}\right)\right) \\
&= \exp\left(\exp\left(O\left(\frac{\log\log n}{ar} + \log\log\log n + a^2r\log a(\log\log\log n)\right)\right)\right).
\end{aligned}$$

We address each of the three terms in the parentheses above separately. First, we have:

$$\begin{aligned}
a^2r\log a(\log\log\log n) &\leq O\left(x^{2/3}(\log x)^{2/3} \cdot \frac{\log s}{\log x} \cdot \log x \cdot (\log\log\log n)\right) \\
&= O\left(x^{2/3}(\log x)^{2/3} \cdot \log s \cdot (\log\log\log n)\right) \\
&= O\left(\frac{(\log\log n)^{2/3}}{(\log\log\log n)^{2/3}(\log s)^{4/3}}(\log x)^{2/3} \cdot \log s \cdot \log\log\log n\right) \\
&= O\left(\frac{(\log\log n)^{2/3}(\log\log\log n)^{1/3}}{(\log s)^{1/3}}(\log x)^{2/3}\right),
\end{aligned}$$

which is equivalent to the desired bound by definition of x. From this, we can also infer the same upper bound on $\log\log\log n$. As for the first of the three terms, we have:

$$\begin{aligned}
\frac{\log\log n}{ar} &= \Theta\left(\frac{\log\log n}{(x\log x)^{1/3}(\log s/\log x)}\right) && \text{by definition of } a, r \\
&= \Theta\left(\frac{\log\log n}{x^{1/3}\log s} \cdot (\log x)^{2/3}\right) \\
&= \Theta\left(\frac{\log\log n}{\log s} \cdot \frac{(\log\log\log n)^{1/3}\cdot(\log s)^{2/3}}{(\log\log n)^{1/3}}(\log x)^{2/3}\right) \\
&= \Theta\left(\frac{(\log\log n)^{2/3}(\log\log\log n)^{1/3}}{(\log s)^{1/3}} \cdot (\log x)^{2/3}\right).
\end{aligned}$$

This completes the proof of this claim. □

Now to complete the proof and bound CC, by Theorem 3.8 it remains to verify that:

$$s \log |\mathbb{F}| \leq \exp\left(\exp\left(O\left(\frac{(\log\log n)^{2/3}(\log\log\log n)^{1/3}}{(\log s)^{1/3}}(\log x)^{2/3}\right)\right)\right).$$

First we show that s satisfies this bound:

$$\begin{aligned}
\log s &\leq o(\sqrt{\log\log n}) \\
\Rightarrow (\log s)^{1/3} &\leq o((\log\log n)^{1/6}) \\
\Rightarrow (\log s)^{1/3}\log\log s &\leq o((\log\log n)^{2/3}) \\
\Rightarrow \log\log s &\leq \frac{(\log\log n)^{2/3}}{(\log s)^{1/3}} \\
&\leq O\left(\frac{(\log\log n)^{2/3}(\log\log\log n)^{1/3}}{(\log s)^{1/3}}(\log x)^{2/3}\right) \\
\Rightarrow s &\leq \exp\left(\exp\left(O\left(\frac{(\log\log n)^{2/3}(\log\log\log n)^{1/3}}{(\log s)^{1/3}}(\log x)^{2/3}\right)\right)\right).
\end{aligned}$$

Similarly, we have:

$$\begin{aligned}
\log|\mathbb{F}| &\leq \exp(O(ar\log\log\log n)) \\
&\leq \exp(O(\log\log n\cdot\log\log\log n)) \\
&\leq \exp(\exp(O(\log\log\log n))) \\
&\leq \exp\left(\exp\left(O\left((\log\log n)^{1/2}\right)\right)\right) \\
&\leq \exp\left(\exp\left(O\left(\frac{(\log\log n)^{2/3}}{(\log s)^{1/3}}\right)\right)\right) && \text{by the above bound on } \log s \\
&\leq \exp\left(\exp\left(O\left(\frac{(\log\log n)^{2/3}(\log\log\log n)^{1/3}}{(\log s)^{1/3}}(\log x)^{2/3}\right)\right)\right).
\end{aligned}$$

□

6.3 Many-Server Regime

It now remains to deal with the regime where $s \geq \exp(\Omega(\sqrt{\log\log n/\log\log\log n}))$. In this regime, it appears that our techniques together with Conjecture 5.2 do not offer any advantage at all over prior work [Efr09]. Put another way, the communication-minimizing setting of parameters in Conjecture 5.2 would be to set r to 0—or equivalently, just use the 2^ℓ-sparse decoding polynomial modulo a product of ℓ primes given by Corollary 3.7. This is exactly the state-of-the-art approach prior to our work [Efr09], which we summarized in Theorem 3.9. For completeness, we re-derive its communication bound here, making no claim to originality.

Proof of Theorem 3.9. We begin by setting parameters and calculating the resulting bounds on $m, |\mathbb{F}|$. To do this, we invoke Corollary 3.7 with $\lfloor \log_2 s \rfloor$ primes (note that the corollary is applicable as $s \geq 32 \Rightarrow \lfloor \log_2 s \rfloor \geq 5$) to obtain $m, \mathbb{F}$ of characteristic $q = 2$ such that there is an S^*_m-decoding polynomial over $\mathbb{F}$ of sparsity $\leq 2^{\lfloor \log_2 s \rfloor} \leq s$ and we have the bounds:

$$m \leq (\log_2 s \log\log_2 s)^{\log_2 s}$$

$$\log|\mathbb{F}| \leq \exp\left(O(\log s \log\log s)\right).$$

Then for the purposes of Corollary 3.4, we have $\ell = \lfloor \log_2 s \rfloor + 1$.

Claim 6.6. *We have:*

$$\exp(\Omega(\ell^2)) \leq \frac{(\log n)^{\lfloor \log_2 s \rfloor}}{2(\ell \log \ell)^\ell} \leq \eta \leq (\log n)^{\lfloor \log_2 s \rfloor}.$$

Proof. First we lower bound η:

$$\begin{aligned}\eta &= \frac{(\log n)^{\ell-1}}{2m} \\ &\geq \frac{(\log n)^{\lfloor \log_2 s \rfloor}}{2(\ell \log \ell)^\ell}.\end{aligned}$$

The upper bound on η is immediate from $\eta \leq (\log n)^{\ell-1}$. Finally, to justify the leftmost inequality, we want to show that:

$$\begin{aligned}\Omega(\ell^2) &\leq \log\left(\frac{(\log n)^{\lfloor \log_2 s \rfloor}}{2(\ell \log \ell)^\ell}\right) \\ &= \lfloor \log_2 s \rfloor \log\log n - \log 2 - \ell \log \ell - \ell \log\log \ell \\ \Leftrightarrow \Omega(\ell^2) &\leq \lfloor \log_2 s \rfloor \log\log n.\end{aligned}$$

This follows from noting that $\ell^2 \leq O((\log s)^2) \leq O(\log s \cdot \log\log n)$, since we are assuming that $s \leq \log n$. □

Claim 6.7. *We have:*

$$d \leq \exp\left(O\left((\log n)^{\frac{1}{\lfloor \log_2 s \rfloor + 1}} \cdot (\log\log n)^{1 - \frac{1}{\lfloor \log_2 s \rfloor + 1}} \cdot \log s \cdot \log\log s\right)\right).$$

Proof. By Claim 6.6, we use the second branch of Corollary 3.4 to compute:

$$\begin{aligned}d &= \exp\left(O\left(\log n \cdot \frac{(\log \eta)^{1-1/\ell}}{\ell \eta^{1/\ell}}\right)\right) \\ &\leq \exp\left(O\left(\log n \cdot \frac{((\ell-1)\log\log n)^{1-1/\ell}}{\ell \eta^{1/\ell}}\right)\right) \\ &\leq \exp\left(O\left(\log n \cdot \frac{((\ell-1)\log\log n)^{1-1/\ell}}{\ell (\log n)^{1-1/\ell}/(2^{1/\ell}\ell \log \ell)}\right)\right) \\ &= \exp\left(O\left(\log n \cdot \frac{(\log\log n)^{1-1/\ell}}{(\log n)^{1-1/\ell}/(\ell \log \ell)}\right)\right) \\ &= \exp\left(O\left((\log n)^{1/\ell} \cdot (\log\log n)^{1-1/\ell} \cdot \ell \log \ell\right)\right) \\ &= \exp\left(O\left((\log n)^{\frac{1}{\lfloor \log_2 s \rfloor + 1}} \cdot (\log\log n)^{1 - \frac{1}{\lfloor \log_2 s \rfloor + 1}} \cdot \log s \cdot \log\log s\right)\right).\end{aligned}$$

□

Now to complete the proof and bound CC, by Theorem 3.8 it remains to verify that:

$$s \log|\mathbb{F}| \leq \exp\left(O\left((\log n)^{\frac{1}{\lfloor \log_2 s \rfloor + 1}} \cdot (\log\log n)^{1 - \frac{1}{\lfloor \log_2 s \rfloor + 1}} \cdot \log s \cdot \log\log s\right)\right).$$

The s term on the LHS is clearly absorbed by the $\log s$ on the RHS, and we know that $\log |\mathbb{F}| \leq \exp(O(\log s \log\log s))$ so this term will also be absorbed.

Finally, we justify the coarser bound of:

$$\exp\left(O\left((\log n)^{\frac{1}{\lfloor \log_2 s \rfloor + 1}} \cdot (\log\log n)^{1-\frac{1}{\lfloor \log_2 s \rfloor + 1}} \cdot \log s \cdot \log\log s\right)\right) \leq \exp\left(\exp\left(O\left(\log\log\log n + \frac{\log\log n}{\log s}\right)\right)\right).$$

The $\frac{\log\log n}{\log s}$ term from the RHS absorbs the $(\log n)^{\frac{1}{\lfloor \log_2 s \rfloor + 1}}$ from the LHS. The $\log\log\log n$ term from the RHS absorbs the $(\log\log n)^{1-\frac{1}{\lfloor \log_2 s \rfloor + 1}}$ term on the LHS. Finally, we have $\log s \cdot \log\log s \leq \exp(O(\log\log s)) \leq \exp(O(\log\log\log n))$ so the $\log\log\log n$ also absorbs the remaining terms on the LHS. □

7 Concrete Empirical Results for Small s

While we believe Conjecture 1.1 and therefore Theorem 1.2 and Corollary 1.4 to be highly plausible, we still believe it is important to complement this with explicit examples of PIR protocols achieving the guarantees we claim for finite values of s, to the extent possible. A similar empirical study was conducted by [CFL+13] to support their conjectural construction of 3-sparse decoding polynomials modulo products of two primes. In this section, we present the empirical results that we obtained to this end.

For each prime characteristic q, we aim to find $q-1$ primes d_i satisfying the preconditions of Lemma 4.5, taking p_i to be the largest prime divisor of $q^{d_i}-1$ and ensuring that the p_i's are distinct. If we find k such primes, then we know there exists a $(k+1)$-term decoding polynomial with respect to some product of k primes. Moreover, it turns out that there also exists a $(k'+1)$-term decoding polynomial with respect to some product of k' primes for any $k' \leq k$. To see this "monotonicity" property, simply note that a root-of-unity grid for k primes can be turned into a root-of-unity grid for k' primes by dropping some rows and columns. (In fact, monotonicity was also shown in more generality for share conversion schemes by Beimel and Lasri [BL26, Theorem 1.4].)

Bearing this in mind, we search over primes d satisfying the inequality of Lemma 4.5 when we take p to be the largest prime divisor of q^d-1. Note that the precondition $q^{\gcd(d_1,\ldots,d_k)} \geq k+1$ is automatically ensured by $q \geq k+1$. Our search has two distinct components which we detail next.

Before that, we remark that an important distinction is that some of the numbers we are considering are so large that it is not easy to verify their primality with certainty. Such numbers are tested with "probable prime" tests. Therefore some of our constructions rely on some probable primes actually being prime. We will explicitly highlight those cases.

Base q repunit primes. First, in Table 2 we tabulate for various characteristics q degrees d such that the base q repunit $(q^d-1)/(q-1)$ is a (probable) prime. For such a d, we can take $p=(q^d-1)/(q-1)$ and $h=q-1$, so the third condition in Lemma 4.5 becomes:

$$\begin{aligned} &\frac{q^d-k}{(q-1)^k} > (k-1)\sqrt{q^d}(1-1/(q-1)^k)+1 \\ \Leftrightarrow\ & q^d-k > (k-1)q^{d/2}((q-1)^k-1)+(q-1)^k, \end{aligned}$$

which holds for any $d \geq 2k+2$. In particular, any $d \geq 2q$ will be usable for every $k \leq q-1$. Values of $d < 2q$ will also be usable for us in some cases; any such cases are included in Table 3.

We include all characteristics ≤ 19 and after that we include characteristics for which many such d are known. This table was built using the information at https://pzktupel.de/Primetables/TableRepunitGenS.php. For completeness, we also provide OEIS references for these sequences when known.

Characteristic q	**OEIS ID**	**# of (Probable) Primes (Clipped at $q-1$)**	**Degrees d**
3	A028491	2	7, 13
5	A004061	4	11, 13, 47, 127
7	A004063	6	131, 149, 1699, 14221, *35201, 126037*
11	A005808	10	73, 139, 907, 1907, 2029, 4801, 5153, 10867, 20161, *293831*
13	A016054	11	137, 283, 883, 991, 1021, 1193 3671, *18743, 31751, 101089, 1503503*
17	A006034	8	47, 71, 419, 4799, *35149, 54919 74509, 1990523*
19	A006035	9	47, 59, 61, 107, 337, 1061 *9511, 22051, 209359*
37	A128003	10	181, 251, 463, 521, *7321, 36473 48157, 87421, 168527, 249341*
43	A240765	6	*6277, 26777, 27299, 40031, 44773, 194119*
67		7	367, 1487, 3347, 4451, *10391, 13411, 167449*
71		7	157, 1583, *31079, 55079, 72043, 150697, 172259*

Table 2: For each characteristic q, we list here known prime degrees $d \geq 2q$ such that $(q^d - 1)/(q - 1)$ is either a prime or a probable prime. Probable primes are italicized. The $d \geq 2q$ requirement is to ensure that these primes are usable for every value $k \leq q - 1$; some values $d < 2q$ may still be usable for particular values of k. We only list up to $q - 1$ primes in each row, since this is all we need when $q \geq k + 1$.

Characteristic q	**Largest Achievable $k \le Q-1$**	**Degrees d in Addition to Table 2**
3	2	None required
5	4	
7	6	
11	10	
13	12	101
17	16	131, 229, 359, 647, 701, 2269, 2459, 2879
19	18	311, 347, 907, 1571, 1759, 2539, *7019, 12919, 18947*
37	23	71, 971, 1451, 1559, 1979, *2371, 2683* 3371, *5431, 6173, 9521, 14197, 55661*
43	17	181, 311, 457, 653, 1021, 1283 *2081, 8317, 22543, 22639, 33623*
67	20	727, 1279, 1361, 1567, 2081, *2693, 3539* *3719, 3907, 6373, 8893, 19609, 56171*
71	15	41, 269, 349, 431, 569, 613, 1231, 2711

Table 3: For each characteristic q, we tabulate the largest $k \le Q-1$ for which we found k (probable) primes satisfying the conditions in Lemma 4.5. We searched over degrees $d \le 60000$ and took the largest (probable) prime factor of $(q^d-1)/(q-1)$ in each case. These prime factors (when known) were computed using queries to FactorDB. Once again, probable primes are italicized.

Base q repunit rough composites. For characteristics $q \le 11$, we can already attain a root-of-unity grid with $k = q-1$ using just the primes already discussed. For $q > 11$, we augment our search by including cases where $(q^d-1)/(q-1)$ is a rough composite integer. By Corollary 4.6, this is still sufficient for our purposes. To this end, we exhaustively search over $d \le 60000$ and see whether taking p to be the largest prime divisor of $(q^d-1)/(q-1)$ will satisfy the third condition of Lemma 4.5. We obtain these prime divisors (when known) using queries to FactorDB. We tabulate the results (in addition to repunit primes with $d < 2q$ that are still usable for these cases) in Table 3.

Summary. We now summarize with a proof of Corollary 1.5:

Proof of Corollary 1.5. On the unconditional side (so without relying on probable primes), the largest k we obtain for Theorem 1.2 is $k = 14$ for $q = 19$. Feeding this into the proof of Corollary 1.4 gives the claimed PIR for up to 15 servers. The relevant degrees d are as follows:

- d such that $(19^d-1)/18$ is prime: $31, 47, 59, 61, 107, 337, 1061$.
- d such that $(19^d-1)/18$ is a rough composite: $83, 311, 347, 601, 907, 953, 1571$. ($d = 1759$ and 2539 are also available alternatives, although these do not allow us to obtain a tuple for $k \ge 15$ because $d = 31, 83$ are only usable for $k \le 14$.)

  The cases of $d = 83, 601, 953$ are not shown in Table 3 because they are not usable for the $k = 18$ result tabulated there, although they are usable for the $k = 14$ result we are showing here.

On the conditional side (assuming that integers passing probable prime tests are actually prime), the largest k we obtain for Theorem 1.2 is $k = 23$ for $q = 37$. Feeding this into the proof of Corollary 1.4 gives the

claimed PIR for up to 24 servers. This relies on 6 probable primes from Table 2 and 7 probable primes from Table 3. For clarity, we list these below:

- The 6 probable primes from Table 2 are $\frac{37^d-1}{36}$ for $d = 7321, 36473, 48157, 87421, 168527, 249341$.
- The 7 probable primes from Table 3 are:
  - $(37^{2371} - 1)/3413169294372$;
  - $(37^{2683} - 1)/307412007186684769953576371198892$;
  - $(37^{5431} - 1)/76456531836$;
  - $(37^{6173} - 1)/3931185602686020572988$;
  - $(37^{9521} - 1)/1419551395838858218572$;
  - $(37^{14197} - 1)/12198192238493172$; and
  - $(37^{55661} - 1)/4007628$.

□

8 Conclusion

For any constant k, our work—assuming Conjecture 1.1—completely resolves the question of the minimal sparsity achievable by an S_m^*-decoding polynomial for m a product of k distinct primes, by showing that sparsity $k+1$ is achievable. This is the best one can hope for due to lower bounds shown by [GK26, BL26]. Via the results by [GKS25], our result implies that for any constant $s \geq 2$ there exists an s-server PIR protocol that on an n-bit database requires communication $\exp(O((\log n)^{1/s}(\log\log n)^{1-1/s}))$, which is strictly better than the previous state of the art for any $s \geq 4$. As additional contributions, we have run experiments to remove the need for any reliance on conjectures for $s \leq 15$, and also taken a first step towards reducing the communication of PIR when the number of servers grows with n.

We next provide a summary of some related work, before concluding with some future directions.

8.1 Related Work

Decoding polynomials. The central object of study in our paper is S-decoding polynomials. These were first introduced by Efremenko [Efr09], where he introduced this notion as a powerful abstraction for PIR. He also provided an example of a decoding polynomial attaining nontrivial sparsity: for $m = 511 = 7 \times 73$, he showed that there is an S_m^*-decoding polynomial of sparsity 3. This was followed up by work by [CFL+13], who demonstrated that under a plausible number-theoretic conjecture, there is an S_m^*-decoding polynomial of sparsity 3 for infinitely many $m = p_1 p_2$. More recently, work by Ghasemi and Kopparty [GK26] shows lower bounds against the sparsity of S-decoding polynomials, in settings including but not limited to the S_m^*-decoding setting that is most relevant to this work.

Secret share conversion. A closely related object to decoding polynomials is *share conversion*. This is an elegant abstraction introduced by [BIKO12] that generalizes S-decoding polynomials and remains applicable to PIR. In share conversion, a secret $a \in S \subseteq \mathbb{Z}_m$ is converted into shares $\mathsf{sh}_1, \ldots, \mathsf{sh}_T$ that are distributed to T parties such that each individual share perfectly hides a. The defining requirement is there are local functions $C_1, \ldots, C_T$ with outputs in $\mathbb{F}$ such that $\sum_{i\in[T]} C_i(\mathsf{sh}_i) = 0$ if and only if $a \neq 0$.

To provide some intuition, we show why an S-decoding polynomial directly gives share conversion, as was shown by [BIKO12]: let the decoding polynomial be $P(X) = \sum_{i\in[T]} \alpha_i X^{\beta_i}$. Then to share a, we will sample $r \leftarrow \mathbb{Z}_m$ and set $\mathsf{sh}_i = r + a\beta_i$. The ith party's local computation is $C_i(\mathsf{sh}) = \alpha_i \gamma^{\mathsf{sh}}$. Now we have:

$$\sum_{i\in[T]} C_i(\mathsf{sh}_i) = \sum_{i\in[T]} \alpha_i \gamma^{\mathsf{sh}_i} = \sum_{i\in[T]} \alpha_i \gamma^{r+a\beta_i} = \gamma^r \cdot P(\gamma^a),$$

which is indeed 0 for $a \in S\backslash\{0\}$ and nonzero for $a = 0$. As mentioned earlier, share conversion also captures constructions that cannot be captured by S-decoding polynomials. In fact, the alternative presentation by [ABL25] of the PIR construction of [GKS25] adds the generalization that decoding polynomials can be replaced by share conversion from $\mathbb{Z}_m$ to a field $\mathbb{F}$ of characteristic not dividing m. A natural question this raises is whether the results of [GKS25] can be improved by using a share conversion with few parties in a way that decoding polynomials would not capture. However, recent follow-up work by Beimel and Lasri [BL26] answers this negatively; in the setting where the share conversion is from $\mathbb{Z}_m$ to $\mathbb{F}$ and the characteristic of $\mathbb{F}$ does not divide m, share conversion and decoding polynomials are equivalent objects.

Locally-decodable codes. An essentially equivalent object to PIR is locally-decodable codes (LDCs), introduced by Katz and Trevisan [KT00, Tre04]. A PIR protocol with T servers, queries of length $\ell_q := \ell_q(n)$, and answers of length $\ell_a := \ell_a(n)$, implies a smooth T-query LDC mapping n bits into a codeword of length 2^{ℓ_q} over an alphabet of size 2^{ℓ_a}. We refer the reader to Yekhanin's comprehensive survey [Yek12] for a discussion of state-of-the-art locally-decodable codes across different regimes. The cases of Reed-Muller codes and matching-vector codes are also captured by the phase diagram in Figure 2.

Doubly-efficient PIR. The only PIR metrics that our work studies are the number of servers and the communication complexity. One could also ask for a protocol where all parties run in sublinear time $n^{o(1)}$ per query. Such a protocol is known as doubly-efficient PIR, and necessitates working in an alternate model such as PIR-with-preprocessing [BIM00]. The problem of doubly-efficient PIR is very important and has seen exciting advances in recent years [LMW23]; however, these protocols are ultimately incomparable to those we consider in this work. They are computationally more efficient, but at the expense of either requiring more communication (albeit still sublinear in n) or only being secure against computationally-bounded adversaries.

8.2 Open Questions

Our work nails down the best communication-server tradeoff achievable for a constant number of servers, while staying within the "matching vectors + decoding polynomials" framework [GKS25] and adhering to the matching vector construction by [BBR94, Gro00, DGY11]. This leaves the following outstanding gaps:

Superconstant servers. Within the same framework and matching vector construction, what is the best communication-server tradeoff achievable in regimes where s is superconstant in n? As we will outline in Section 2.3, our work makes progress on answering this question but we suspect improvements should be possible. We see at least a couple of candidate routes towards addressing this:

- Find families of primes $p_1, \ldots, p_k$ that occur more frequently and admit decoding polynomials of sparsity $k + 1$. As discussed in Section 2.2, we seem to be far from a complete understanding of exactly when these decoding polynomials exist so this direction is promising; our work only provides evidence for their existence for infinitely many $m, \mathbb{F}$.

- More modestly, find families of $p_1, \ldots, p_k$ that occur more frequently, at the expense of having decoding polynomials that do not quite hit sparsity $k+1$ but attain some intermediate regime of sparsity that is still useful. Our work focuses on the sparsity exactly $k+1$ case; it may be beneficial when considering a superconstant number of servers to relax this.

Unconditional constructions. Our main theorem rests on a number-theoretic conjecture. While we have provided heuristic justification that this conjecture is plausible, it would be preferable to have unconditional statements. Heuristics of a similar flavor have appeared in previous constructions of matching-vector PIR [Yek08, CFL+13]; we expect proving these heuristics to be extremely difficult since they are related, for example, to the infinitude of Mersenne primes which is wide open. However, one could hope for a different construction from ours that does not rest on Mersenne-like conjectures.

Better constructions of matching vector families. The past decades have seen advances in finding nontrivially sparse S_m^*-decoding polynomials, yet tantalizingly the best matching vector family construction has stood since the work of Grolmusz [Gro00] in 2000! Existing lower bounds [BDL14] are far from suggesting that this construction is tight. We believe this to be a difficult but underexplored and high-impact open direction.

New frameworks for PIR. Another ambitious direction is to find new ways to build low-communication PIR from algebraic primitives, including—but not limited to—matching vectors. Once again, all constructions in the constant-server and sub-polynomial-communication regime are encapsulated by the unifying framework of [GKS25], but we see no reason why this should be the end of the story.

References

[ABL25] Bar Alon, Amos Beimel, and Or Lasri. Simplified PIR and CDS protocols and improved linear secret-sharing schemes. In Benny Applebaum and Huijia (Rachel) Lin, editors, *Theory of Cryptography - 23rd International Conference, TCC 2025, Aarhus, Denmark, December 1-5, 2025, Proceedings, Part II*, volume 16269 of *Lecture Notes in Computer Science*, pages 365–398. Springer, 2025.

[Amb97] Andris Ambainis. Upper bound on the communication complexity of private information retrieval. In *ICALP*, pages 401–407. Springer, 1997.

[Apo13] Tom M Apostol. *Introduction to analytic number theory*. Springer Science & Business Media, 2013.

[BBR94] David A. Mix Barrington, Richard Beigel, and Steven Rudich. Representing boolean functions as polynomials modulo composite numbers. *Computational Complexity*, 4(4):367–382, December 1994.

[BDL14] Abhishek Bhowmick, Zeev Dvir, and Shachar Lovett. New bounds for matching vector families. *SIAM J. Comput.*, 43(5):1654–1683, 2014.

[BH62] Paul T. Bateman and Roger A. Horn. A heuristic asymptotic formula concerning the distribution of prime numbers. *Mathematics of Computation*, 16(79):363–367, 1962.

[BIKO12] Amos Beimel, Yuval Ishai, Eyal Kushilevitz, and Ilan Orlov. Share conversion and private information retrieval. In *Proceedings of the 27th Conference on Computational Complexity, CCC 2012, Porto, Portugal, June 26-29, 2012*, pages 258–268. IEEE Computer Society, 2012.

[BIKR02] Amos Beimel, Yuval Ishai, Eyal Kushilevitz, and Jean-François Raymond. Breaking the $O(n^{1/(2k-1)})$ barrier for information-theoretic private information retrieval. In *FOCS*, pages 261–270. IEEE Computer Society, 2002.

[BIM00] Amos Beimel, Yuval Ishai, and Tal Malkin. Reducing the servers computation in private information retrieval: PIR with preprocessing. In Mihir Bellare, editor, *Advances in Cryptology - CRYPTO 2000, 20th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2000, Proceedings*, volume 1880 of *Lecture Notes in Computer Science*, pages 55–73. Springer, 2000.

[BL26] Amos Beimel and Or Lasri. Private information retrieval: Share conversions vs. decoding polynomials. *CRYPTO*, 2026.

[Bou09] Paul Bourdelais. A generalized repunit conjecture, 2009. Posting to NMBRTHRY mailing list.

[CFL+13] Yeow Meng Chee, Tao Feng, San Ling, Huaxiong Wang, and Liang Feng Zhang. Query-efficient locally decodable codes of subexponential length. *Computational Complexity*, 22(1):159–189, 2013.

[CKGS98] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *J. ACM*, 1998.

[DG16] Zeev Dvir and Sivakanth Gopi. 2-server PIR with subpolynomial communication. *J. ACM*, 2016.

[DGY11] Zeev Dvir, Parikshit Gopalan, and Sergey Yekhanin. Matching vector codes. *SIAM Journal on Computing*, 40(4):1154–1178, 2011.

[Dic30] Karl Dickman. On the frequency of numbers containing prime factors of a certain relative magnitude. *Arkiv for matematik, astronomi och fysik*, 22(10):A–10, 1930.

[Dub93] Harvey Dubner. Generalized repunit primes. *mathematics of computation*, 61(204):927–930, 1993.

[Efr09] Klim Efremenko. 3-query locally decodable codes of subexponential length. In *STOC*, 2009.

[GK26] Fatemeh Ghasemi and Swastik Kopparty. Fourier sparsity of delta functions and matching vector PIRs. In Shubhangi Saraf, editor, *17th Innovations in Theoretical Computer Science Conference, ITCS 2026, Bocconi University, Milan, Italy, January 27-30, 2026*, volume 362 of *LIPIcs*, pages 68:1–68:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2026.

[GKS25] Fatemeh Ghasemi, Swastik Kopparty, and Madhu Sudan. Improved PIR schemes using matching vectors and derivatives. In *STOC*, 2025.

[Gro00] Vince Grolmusz. Superpolynomial size set-systems with restricted intersections mod 6 and explicit Ramsey graphs. *Comb.*, 20(1):71–86, 2000.

[IS10] Toshiya Itoh and Yasuhiro Suzuki. Improved constructions for query-efficient locally decodable codes of subexponential length. *IEICE Transactions on Information and Systems*, 93(2):263–270, 2010.

[KT00] Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In F. Frances Yao and Eugene M. Luks, editors, *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 80–86. ACM, 2000.

[LMW23] Wei-Kai Lin, Ethan Mook, and Daniel Wichs. Doubly efficient private information retrieval and fully homomorphic RAM computation from ring LWE. In *STOC*, 2023.

[Mor14] Pieter Moree. Integers without large prime factors: from Ramanujan to de Bruijn. *Integers*, 14, 2014.

[MV07] Hugh L Montgomery and Robert C Vaughan. *Multiplicative number theory I: Classical theory*, volume 97. Cambridge university press, 2007.

[Rag07] Prasad Raghavendra. A note on Yekhanin's locally decodable codes. *Electron. Colloquium Comput. Complex.*, TR07, 2007.

[Sch76] Wolfgang M Schmidt. *Equations over finite fields: an elementary approach.* Springer, 1976.

[SS58] Andrzej Schinzel and Wacław Sierpiński. Sur certaines hypotheses concernant les nombres premiers. *Acta Arithmetica*, 4(3):185–208, 1958.

[Tre04] Luca Trevisan. Some applications of coding theory in computational complexity. *CoRR*, cs.CC/0409044, 2004.

[WY05] David Woodruff and Sergey Yekhanin. A geometric approach to information-theoretic private information retrieval. In *CCC*. IEEE, 2005.

[Yek08] Sergey Yekhanin. Towards 3-query locally decodable codes of subexponential length. *J. ACM*, 2008.

[Yek12] Sergey Yekhanin. Locally decodable codes. *Foundations and Trends in Theoretical Computer Science*, 6(3), 2012.

A Proof of Corollary 3.4

In all cases, we will proceed by instantiating Theorem 3.3 for suitably chosen h and w.

Case 1: $\eta \leq 2^{O(\ell^2)}$. Fix a constant $C_1 > 0$ and take $w = \lfloor C_1 \log n \rfloor$. Then we have $mw \sim C_1 \frac{(\log n)^\ell}{\eta}$, or equivalently:

$$(mw)^{1/\ell} \sim C_1^{1/\ell} \frac{\log n}{\eta^{1/\ell}}. \tag{8}$$

So since $\eta \geq 1$, this implies $(mw)^{1/\ell} \leq C_1^{1/\ell} \log n$.

Also, set $C_2 > 0$ and $h \sim C_2 \log n$ such that $\binom{h}{w} \geq n$ and $C_2 > 3\max(C_1, 1)$. Then we have:

$$\sum_{j=0}^{\lfloor (mw)^{1/\ell} \rfloor} \binom{h}{j} = O\left(\binom{h}{(mw)^{1/\ell}}\right)$$

$$= O\left(\binom{C_2 \log n}{C_1^{1/\ell} \log n/\eta^{1/\ell}}\right).$$

In the first step above, we are using the well-known inequality that for $a \leq b/3$ we have $\binom{b}{0} + \ldots + \binom{b}{a} = O\left(\binom{b}{a}\right)$, together with the fact that:

$$(mw)^{1/\ell} \leq C_1^{1/\ell} \log n \leq \max(C_1, 1) \log n < \frac{C_2}{3} \log n \sim \frac{h}{3}.$$

Continuing along, if $\eta \leq 2^{O(\ell)}$, then this is $\text{poly}(n)$ which is indeed consistent with the statement of the corollary. If $\eta \geq 2^{\Omega(\ell)}$, then we have:

$$\begin{aligned}\log \binom{C_2 \log n}{C_1^{1/\ell} \log n/\eta^{1/\ell}} &= \Theta\left(\frac{C_1^{1/\ell} \log n}{\eta^{1/\ell}} \cdot \log \frac{C_2 \eta^{1/\ell}}{C_1^{1/\ell}}\right) \\ &= \Theta\left(\frac{\log n \log(\eta^{1/\ell})}{\eta^{1/\ell}}\right),\end{aligned}$$

so this construction achieves the claimed performance.

Case 2: $\eta \geq 2^{\Omega(\ell^2)}$**.** Here, we will take $\theta \geq 10$ such that $\theta = \Theta(\log \eta/\ell^2)$ and take $w \sim \frac{\log n}{\theta}$ and $h \sim \frac{\log n}{\theta} \cdot 2^\theta$. In this regime, we have $\binom{h}{w} \geq (h/w)^w = n$. Secondly, we have $mw \sim \frac{(\log n)^\ell}{\eta\theta} \Rightarrow (mw)^{1/\ell} \sim \frac{\log n}{\eta^{1/\ell}\theta^{1/\ell}} < h/3$. Consequently, we once again have:

$$\begin{aligned}\sum_{j=0}^{\lfloor (mw)^{1/\ell} \rfloor} \binom{h}{j} &= O\left(\binom{h}{(mw)^{1/\ell}}\right), \text{ and} \\ \log \binom{h}{(mw)^{1/\ell}} &= \log \binom{\log n \cdot 2^\theta/\theta}{\log n/(\eta\theta)^{1/\ell}} \\ &= \Theta\left(\frac{\log n}{(\eta\theta)^{1/\ell}} \log \frac{2^\theta \eta^{1/\ell}}{\theta^{1-1/\ell}}\right) \\ &= \Theta\left(\frac{\log n}{(\eta \log \eta)^{1/\ell}} \log \frac{2^\theta \eta^{1/\ell}}{\theta^{1-1/\ell}}\right) \\ &= \Theta\left(\frac{\log n}{(\eta \log \eta)^{1/\ell}} \left(\theta + \frac{\log \eta}{\ell}\right)\right) \\ &= \Theta\left(\frac{\log n}{(\eta \log \eta)^{1/\ell}} \left(\frac{\log \eta}{\ell}\right)\right) \\ &= \Theta\left(\frac{(\log n) \cdot (\log \eta)^{1-1/\ell}}{\ell \eta^{1/\ell}}\right),\end{aligned}$$

as claimed.

B Proof of Theorem 3.6

Lemma B.1. *Let $S \subseteq \mathbb{Z}_m$ be an arbitrary subset containing 0 and $\mathbb{F}$ a field such that $m||\mathbb{F}| - 1$. Assume there exists some primitive mth root of unity $\gamma_m \in \mathbb{F}$ and an S-decoding polynomial A with respect to γ_m of sparsity*

T. Then for any primitive mth root of unity $\delta_m \in \mathbb{F}$, there exists an S-decoding polynomial B with respect to δ_m of sparsity T.

Proof. Since γ_m, δ_m are primitive mth roots of unity, there exists an integer u such that $\gcd(u, m) = 1$ and $\delta_m^u = \gamma_m$. So if we define $B(X) := A(X^u)$, we have for any $j \in \mathbb{Z}_m$ that $B(\delta_m^j) = A(\delta_m^{ju}) = A(\gamma_m^j)$. The conclusion follows. □

Proof of Theorem 3.6. Let $M = m_1 m_2 \cdots m_r$, $C = \operatorname{lcm}(c_1, \ldots, c_r)$, and $\mathbb{F} = \mathbb{F}_{q^C}$. For each i, since there is an $S_{m_i}^*$-decoding polynomial modulo m_i over $\mathbb{F}_{q^{c_i}}$, the field $\mathbb{F}_{q^{c_i}}$ contains a primitive m_i-th root of unity. Hence $m_i \mid q^{c_i} - 1 \Rightarrow m_i \mid q^C - 1$. Because the m_i's are pairwise coprime, it follows that their product $M \mid q^C - 1$ as well. Thus $\mathbb{F}^\times$ contains a primitive M-th root of unity. Fix one and call it γ.

For each i, set $M_i = M/m_i$ and let $\beta_i = \gamma^{M_i}$ be a primitive m_i-th root of unity. Let $K_i \subseteq \mathbb{F}_{q^C}$ be the unique[7] subfield that is isomorphic to $\mathbb{F}_{q^{c_i}}$. Since β_i is an m_i-th root of unity and $m_i \mid q^{c_i} - 1$, we have $\beta_i \in K_i$. So by Lemma B.1, there exists an $S_{m_i}^*$-decoding polynomial

$$P_i(X) = \sum_{e \in E_i} c_{i,e} X^e$$

of sparsity d_i over K_i with respect to the root of unity β_i. By reducing exponents modulo m_i, we may assume that all elements of E_i lie in $\mathbb{Z}_{m_i}$ and $|E_i| = d_i$. Note also that for each e we have $c_{i,e} \in K_i \Rightarrow c_{i,e} \in \mathbb{F}$. We may now define the polynomial

$$P(X) = \prod_{i=1}^{r} P_i(X^{M_i}) \in \mathbb{F}[X].$$

We claim that P is an S_M^*-decoding polynomial mod M over $\mathbb{F}$. Indeed, for $j \in S_M^*$ we have $P(\gamma^j) \neq 0 \Leftrightarrow P_i(\beta_i^j) = P_i(\gamma^{jM_i}) \neq 0 \forall i$. It immediately follows that for $j = 0$ we have $P_i(1) \neq 0 \forall i \Rightarrow P(1) \neq 0$. For $j \in S_M^* \backslash \{0\}$, there must exist an index i such that $j \not\equiv 0 \pmod{m_i}$. For this i, we therefore have $P_i(\beta_i^j) = 0 \Rightarrow P(\gamma^j) = 0$. This proves that P is an S_M^*-decoding polynomial.

For sparsity, note that P is a product of r polynomials with sparsities $d_1, \ldots, d_r$. This immediately implies that P itself has sparsity $\leq d_1 \ldots d_r$. Although this is enough for our purposes, we show for completeness that P has sparsity exactly $d_1 \ldots d_r$. To show this, we want to show that there are no collisions in the map $(e_1, \ldots, e_r) \mapsto \sum_{i=1}^r e_i M_i \bmod M$ for $e_i \in \mathbb{Z}_{m_i}$. This map is linear so it suffices to show that $\sum_{i=1}^r e_i M_i \equiv 0 \pmod{M}$ implies that $e_i \equiv 0 \pmod{m_i}$ for all i. For a particular i^*, this can be seen by reducing both sides of the congruence modulo m_{i^*}. This implies that $e_{i^*} M_{i^*} \equiv 0 \pmod{m_{i^*}} \Rightarrow e_{i^*} \equiv 0 \pmod{m_{i^*}}$, as claimed. □

C Proof of Lemma 3.29

In this section we prove Lemma 3.29, assuming GRH for all L-functions. First we define the function

$$\vartheta(x; m, a) = \sum_{\substack{\text{primes } p \leq x \\ p = a \bmod m}} \log p.$$

Before we present the proof, we requirng the following two lemmas from analytic number theory.

Lemma C.1 (Corollary 13.8 of [MV07])**.** *Fix any modulus m. Assuming the GRH for all L-functions modulo m. Let a be an integer such that $\gcd(a, m) = 1$ and $x \geq 2$ be a real number. Then,*

$$\left| \vartheta(x; m, a) - \frac{x}{\varphi(m)} \right| \leq O(\sqrt{x} \cdot (\log x)^2).$$

[7] Uniqueness can be seen by noticing that K_i is exactly the splitting field of the polynomial $X^{q^{c_i}} - X$.

Lemma C.2 (Abel summation formula, Theorem 4.2 of [Apo13]). *Let $0 < x < y$ be real numbers. Let $(b_n)_{n=0}^{\infty}$ be a sequence of complex numbers, and let $B(t) = \sum_{n \le t} b_n$ and let $f : \mathbb{R} \to \mathbb{R}$ be a continuously differentiable function on the interval $[x, y]$. Then,*

$$\sum_{\substack{n \in \mathbb{Z}, \\ x < n \le y}} b_n \cdot f(n) = B(y)f(y) - B(x)f(x) - \int_x^y B(t)f'(t)\,dt.$$

Proof of Lemma 3.29. First, define

$$\begin{aligned} b_n &= \begin{cases} \log n & \text{prime } n,\ n = a \bmod m, \\ 0 & \text{otherwise,} \end{cases} \\ B(t) &= \vartheta(t; m, a), \\ f(t) &= \frac{1}{t \log t}. \end{aligned}$$

Observe that for any $1 < x < y \le \infty$, f is continuously differentiable on $[x, y]$, with

$$f'(t) = -\frac{\log t + 1}{(t \log t)^2} \tag{9}$$

By Lemma C.1, we can write $B(t) = \vartheta(t; m, a) = t/\varphi(m) + E(t)$, where $|E(t)| \le O(\sqrt{t}(\log t)^2)$. Then, applying Lemma C.2, we get

$$\begin{aligned} \sum_{\substack{\text{primes } p \in [x,y] \\ p = a \bmod m}} \frac{1}{p} &= \sum_{x < n \le y} b_n \cdot f(n) \\ &= B(y)f(y) - B(x)f(x) - \int_x^y B(t)f'(t)dt && \text{Lemma C.2} \\ &= \frac{1}{\varphi(m)} \left(yf(y) - xf(x) - \int_x^y tf'(t)dt \right) + \xi, \end{aligned}$$

where the penultimate equality follows from writing $B(t) = t/\varphi(m) + E(t)$, and calling $\xi = E(y)f(y) - E(x)f(x) - \int_x^y E(t)f'(t)dt$, and the last inequality follows from Claim C.1 and the following calculation.

$$\begin{aligned} \frac{1}{\varphi(m)} \left(yf(y) - xf(x) - \int_x^y tf'(t)dt \right) &= \frac{1}{\varphi(m)} \left(\frac{1}{\log y} - \frac{1}{\log x} - \int_x^y tf'(t)dt \right) && \text{definition of } f \\ &= \frac{1}{\varphi(m)} \left(\frac{1}{\log y} - \frac{1}{\log x} - tf(t)|_x^y + \int_x^y f(t)dt \right) && \text{integration by parts} \\ &= \frac{1}{\varphi(m)} \left(\int_x^y f(t)dt \right) && \text{evaluating } tf(t) \text{ at } x, y \\ &= \frac{1}{\varphi(m)} \left(\log\log t|_x^y \right) \\ &= \frac{1}{\varphi(m)} \log\log y - \log\log x. \end{aligned}$$

Claim C.1.

$$|\xi| := \left| E(y)f(y) - E(x)f(x) - \int_x^y E(t)f'(t)dt \right| \leq O\left(\frac{\log x}{\sqrt{x}}\right).$$

Proof. By Lemma C.1, we know that $|E(x)f(x)| \leq O(\sqrt{x}(\log x)^2/(x \log x)) = O(\log x/\sqrt{x})$ and similarly for $|E(y)f(y)| \leq O(\log y/\sqrt{y})$. Now, we bound the integral term:

$$\begin{aligned}
\left| \int_x^y E(t)f'(t)dt \right| &\leq O\left(\int_x^y \sqrt{t}(\log t)^2 \cdot \frac{\log t + 1}{(t \log t)^2} dt \right) && \text{by Lemma C.1 and } t \geq 1 \\
&= O\left(\int_x^y \frac{\log t}{t^{3/2}} dt \right) \\
&= O\left(\frac{-2\log t - 4}{\sqrt{t}} \Big|_x^y \right) \\
&\leq O\left(\frac{\log x}{\sqrt{x}} \right).
\end{aligned}$$

The claim then follows by the triangle inequality. □

□

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il