Electronic Colloquium on Computational Complexity, Report No. 156 (2026)

Improved Subexponential Upper Bounds for 3-Restricted Matching Vector Families

Sidhant Saraogi*

Abstract

Matching Vector families (MVFs) are defined by two ordered lists of vectors in $\mathbb{Z}_m^n$ whose inner products satisfy specific residue patterns modulo an integer m. Most famously, restricted MVFs are used to construct the best-known constant-query Locally Decodable codes (LDCs).

We prove an upper bound of $2^{O\left(\sqrt{n \log n \log m}\right)}$ on the size of 3-restricted MVFs in $\mathbb{Z}_m^n$ for $m \leq \sqrt{n}$, substantially improving on the previous best bound of $2^{O(n/\log n)}$ by Bhowmick, Dvir and Lovett (STOC'13, SICOMP'14). Our proof relies on a new polynomial-method argument that controls collisions in sumsets of matching vectors.

1 Introduction

A Matching Vector (MV) family is defined by two ordered lists of vectors $\mathcal{U} = \{u_1, \ldots, u_K\}$ and $\mathcal{V} = \{v_1, \ldots, v_K\}$ in $\mathbb{Z}_m^n$ for some $m, n \geq 1$ such that the inner products of pairs of vectors from $\mathcal{U}$ and $\mathcal{V}$ satisfy certain residue patterns modulo m. More precisely:

Definition 1 (Matching Vector Family)**.** Let $S \subseteq \mathbb{Z}_m \setminus \{0\}$. An S-matching vector family is a set of vectors $\mathcal{U} = \{u_1, \ldots, u_K\}$, $\mathcal{V} = \{v_1, \ldots, v_K\}$, $\mathcal{U}, \mathcal{V} \subseteq \mathbb{Z}_m^n$ such that:

- $\forall i \in [K], \langle u_i, v_i \rangle = 0$.
- $\forall i \neq j,\ i, j \in [K],\ \langle u_i, v_j \rangle \in S$.

Such an MV family is referred to as a $|S| + 1$-restricted MV family. When $S = \mathbb{Z}_m \setminus \{0\}$, we simply refer to it as an MV family. We refer to K as the size of the MV family.

Famously, MV families were used to construct the best-known constant-query Locally Decodable codes (LDCs) [Yek08, Efr09, DGY11]. MVFs have since found a wider array of applications including to private information retrieval [Yek08, Efr09, GKS25, ABL26], conditional disclosure of secrets [ABL26], circuit complexity, and catalytic computation [HPR26]. They are also connected to the study of Ramsey graphs [Gro00] and polynomials weakly representing the OR function [Gro95].

As is common, we denote by $\mathbf{MV}(m, n)$ the size of the largest possible MV family in $\mathbb{Z}_m^n$ and by $\mathbf{MV}(m, n, r)$ the size of the largest possible r-restricted MV family in $\mathbb{Z}_m^n$.

*Georgetown University. Email: ss4456@georgetown.edu.

ISSN 1433-8092

When m has several distinct prime factors, the situation is much richer. For suitable constant moduli m with t distinct prime factors, Grolmusz [Gro00] constructed MVFs whose off-diagonal inner products lie in a canonical set of $2^t - 1$ nonzero residues. In our terminology, these are 2^t-restricted MVFs, and they satisfy

$$\mathbf{MV}(m, n, 2^t) \geq \exp\left(\Omega\left(\frac{(\log n)^t}{(\log\log n)^{t-1}}\right)\right).$$

This remains the best-known construction of MVFs. Based on [BDL14] and the recently proved Polynomial Freiman-Rusza Theorem [GGMT26], we know that the best upper bound for MVFs is

$$\mathbf{MV}(m, n) \leq (c_m)^{n/\log n}$$

for some constant c_m depending only on m.

Locally Decodable Codes (LDCs). The most surprising application of MVFs is to the construction of LDCs, codes that allow for extremely efficient decoding of individual message bits.

Definition 2 (Locally Decodable Codes)**.** A q-ary code $C : \mathbb{F}_q^K \to \mathbb{F}_q^N$ is said to be (r, δ, ε)-locally decodable if there exists a randomized decoding algorithm $\mathcal{A}$ such that:

1. For all $\mathbf{x} \in \mathbb{F}_q^K$, $i \in [K]$, and all vectors $\mathbf{y} \in \mathbb{F}_q^N$ such that
$$d(C(\mathbf{x}), \mathbf{y}) \leq \delta N,$$
we have
$$\Pr[\mathcal{A}^{\mathbf{y}}(i) = \mathbf{x}(i)] \geq 1 - \varepsilon,$$
where the probability is over the random coin tosses of $\mathcal{A}$.
2. $\mathcal{A}$ makes at most r queries to $\mathbf{y}$.

[Yek08] showed that an r-restricted MV family of size K in $\mathbb{Z}_m^n$ can be used to construct an r-query locally decodable code with message length K and codeword length m^n. MVFs are central to the construction of matching vector codes. Building on this framework, a line of works [Yek08, Efr09, DGY11] constructed the best-known constant-query LDCs. In particular, the [Efr09] construction gives 3-query LDCs of codeword length

$$N = \exp\left(\exp\left(O\left(\sqrt{\log K \log\log K}\right)\right)\right).$$

On the other hand, despite a long line of works [KT00, Kd04, GKST06, AGKM23, JM25, BHKL25], we only know that

$$N \geq \tilde{\Omega}(K^{r/(r-2)})$$

for r-query LDCs. As a result, the study of lower bounds for Matching Vector codes has gained importance. Based on the [BDL14] upper bound, constant-query MV codes must satisfy

$$N \geq K^{\Omega(\log\log K)}.$$

for constant m.

1.1 Our Results

In this work, we show a substantially stronger subexponential upper bound for 3-restricted MVFs

Theorem 3. *For every $n \geq 2$ and every integer m satisfying $2 \leq m \leq \sqrt{n}$,*

$$\mathbf{MV}(m, n, 3) \leq 2^{O\left(\sqrt{n \log m \log n}\right)}.$$

In particular, for constant m,

$$\mathbf{MV}(m, n, 3) \leq 2^{O\left(\sqrt{n \log n}\right)}.$$

As a corollary, we can prove a strong lower bound on certain 3-query matching vector codes.

Corollary 4. *Every 3-query matching vector code obtained from a 3-restricted MVF $\mathcal{U}, \mathcal{V} \subseteq \mathbb{Z}_m^n$ with $2 \leq m \leq \sqrt{n}$, message length K, and codeword length $N = m^n$, satisfies*

$$N \geq \exp\left(\Omega\left(\frac{(\log K)^2}{\log \log K}\right)\right).$$

1.2 Techniques

Our work continues the study of additive structure in $\mathcal{U}$ and $\mathcal{V}$. In [BDL14], this structure is explored through the lens of approximate duality [ZBS11]. In a different direction, [ADL+25] suggested studying collisions in sumsets of $\mathcal{U}$: many distinct subsets having the same sum should impose strong structure on the matching-vector family. We refer to the collection of all s-subsets whose corresponding vectors have the same sum as a *collision class*. Our proof develops this viewpoint directly by showing that no collision class can be too large.

The key step is to show that a collision class cannot contain a large moonflower[1] [LMW26]. The polynomial method enters by starting from the low-rank matrix of inner-product residues between $\mathcal{U}$ and $\mathcal{V}$. From this matrix, we construct a new matrix that encodes the inner-product structure between the equal sums and suitable vectors in $\mathcal{V}$. The presence of a moonflower inside the collision class forces this new matrix to have the form $J - I$, and hence to have large rank. At the same time, its construction from the original low-rank residue matrix via a multilinear polynomial gives a strong upper bound on its rank. Comparing these two bounds limits the size of the moonflower.

Finally, let us fix a collision class and choose a maximal moonflower inside it. Because the moonflower is maximal, the union of its sets must intersect every other set in the class. Otherwise, we could add a disjoint set and obtain a larger moonflower. Thus every representation of the fixed s-sum contains at least one element from this relatively small union. Once such an element is fixed and removed, the remaining $s - 1$ elements form a representation of a corresponding $(s-1)$-sum. This bounds the size of an s-sum collision class in terms of the size of the maximal moonflower and the largest $(s-1)$-sum collision class. This recurrence gives a uniform bound on the size of every s-sum collision class. Combining this with the fact that there are only m^n possible sums yields the desired upper bound on K.

[1] A family of sets is a moonflower if each set contains an element that belongs to none of the others. Pairwise disjoint sets would also suffice for our argument, but moonflowers capture the precise combinatorial structure that we need.

1.3 Concurrent Work

Independently, Aggarwal and Obremski [AO26] obtain the same upper bound for 3-restricted MVFs.

Acknowledgements

The author would like to thank Alexander Golovnev for helpful discussions during the writeup of this work.

AI Usage

The main ideas behind this work were obtained while exploring multiple directions to understand the intuition developed in [ADL+25] in extended conversations with ChatGPT 5.5. Specifically, the main idea behind the proof of Lemma 11 was developed by ChatGPT 5.5 and refined into its current form by the author. Subsequent conversations were not able to provide any further insight into improving the bound for 3-restricted MVFs or for r-restricted MVFs more generally. While multiple LLMs were used in reviewing and editing the manuscript, all mathematical claims, proofs, and errors are the sole responsibility of the author.

2 Preliminaries

2.1 Matching Vector Families

We will utilize some known lemmas on restricted MV families. Succinctly, we can turn any 3-restricted MVF modulo m into a *typical* 3-restricted MVF whose inner products modulo m satisfy some nice properties and m itself is constrained in the following fashion.

Lemma 5 ([BDL14, ADL+25])**.** *Let $(\mathcal{U}, \mathcal{V})$ be a S-matching vector family in $\mathbb{Z}_m^n$ and for any proper divisor $m'|m$, $(\mathcal{U}, \mathcal{V})$ is not a matching vector family modulo m'. Then, m has at most $|S|$ prime factors.*

Let us now define a typical 3-restricted MVF

Definition 6 (Typical 3-restricted MVF)**.** We say a 3-restricted MVF $\mathcal{U}, \mathcal{V}$ modulo m is *typical* if $m = p^a q^b$, where p, q are distinct primes and $a, b \geq 1$. Furthermore, the residues $S = \{\alpha, \beta\} \subseteq \mathbb{Z}_m \setminus \{0\}$ satisfy the following residue structure

$$\alpha \equiv 0 \pmod{p^a q^{b-1}}, \qquad \alpha \not\equiv 0 \pmod{q^b},$$

and

$$\beta \equiv 0 \pmod{p^{a-1} q^b}, \qquad \beta \not\equiv 0 \pmod{p^a}.$$

Finally, the following lemma combined with the previous one allows us to convert a 3-restricted MVF into a typical one without any loss of parameters.

Lemma 7 ([ADL+25])**.** *Let $\mathcal{U}, \mathcal{V}$ be a 3-restricted MVF modulo m such that for any proper divisor $m' \mid m$, $\mathcal{U}, \mathcal{V}$ is not a 3-restricted MVF modulo m' and that m is not a prime power. Then $\mathcal{U}, \mathcal{V}$ is a typical 3-restricted MVF.*

2.2 Moonflowers

We will utilize the concept of moonflowers in our lower bound. Moonflowers were introduced by Lovett, Meka, and Wang in recent work inspired by sunflowers, with applications to code sparsification. We do not use any results from that work in our proof, but moonflowers are the right object to capture our lower bound structure.

Definition 8 (Moonflower)**.** A family of sets $A_1, \ldots, A_L$ is an L-moonflower if for every $h \in [L]$, there is an element

$$i_h \in A_h$$

such that

$$i_h \notin A_t \qquad \text{for every } t \neq h.$$

We call i_h a unique element of A_h.

2.3 Linear Algebra

In this section, we prove two auxiliary linear-algebraic lemmas. Our first lemma allows us to pass from matrices over the prime-power ring $\mathbb{Z}_{q^b}$ to matrices over the field $\mathbb{F}_q$, while preserving the rank bound needed in our main polynomial-method argument. For $a \in q^{b-1}\mathbb{Z}_{q^b}$, let $\tau_q(a) \in \mathbb{F}_q$ denote the coefficient of q^{b-1}, that is,

$$\tau_q(q^{b-1}c) = c \bmod q.$$

This is well-defined.

Lemma 9. *Let $R = \mathbb{Z}_{q^b}$ where q is a prime and $b \geq 1$. Let*

$$X \in R^{M \times n}, \qquad Y \in R^{n \times N}.$$

Suppose every entry of XY lies in $q^{b-1}R$. Define $L \in \mathbb{F}_q^{M \times N}$ by

$$L_{ij} = \tau_q((XY)_{ij}).$$

Then

$$\operatorname{rank}_{\mathbb{F}_q}(L) \leq n.$$

Proof. Take the Smith normal form of

$$X = ADB$$

over R, where A, B are invertible and D is diagonal with at most n nonzero rows. Since XY has entries in $q^{b-1}R$, so does

$$DBY = A^{-1}XY.$$

Let

$$L'_{ij} = \tau_q((DBY)_{ij}).$$

The matrix L' has at most n nonzero rows, hence

$$\operatorname{rank}_{\mathbb{F}_q}(L') \leq n.$$

Since $XY = A(DBY)$ and $\tau_q(ra) = (r \bmod q)\tau_q(a)$ for $a \in q^{b-1}R$, we have

$$L = \overline{A}L',$$

where $\overline{A} := A \bmod q$ entrywise. Therefore

$$\operatorname{rank}_{\mathbb{F}_q}(L) \le \operatorname{rank}_{\mathbb{F}_q}(L') \le n.$$

□

Our second lemma bounds the rank of a matrix obtained by applying a multilinear polynomial entrywise to several low-rank matrices.

Lemma 10. *Let $B_1, \dots, B_d$ be $M \times N$ matrices over a field $\mathbb{F}$, and let*

$$\operatorname{rank}_{\mathbb{F}}(B_h) \le n \qquad \text{for every } h \in [d].$$

Let $f \in \mathbb{F}[z_1, \dots, z_d]$ be multilinear, and define P entry-wise by

$$P_{ij} = f((B_1)_{ij}, \dots, (B_d)_{ij}).$$

Then

$$\operatorname{rank}_{\mathbb{F}}(P) \le (n+1)^d.$$

Proof. Write

$$f(z_1, \dots, z_d) = \sum_{T \subseteq [d]} c_T \prod_{h \in T} z_h.$$

For each $T \subseteq [d]$, define B_T entrywise by

$$(B_T)_{ij} = \prod_{h \in T} (B_h)_{ij}.$$

Thus B_T is the Hadamard product of the matrices B_h with $h \in T$. By the Hadamard-product rank inequality,

$$\operatorname{rank}_{\mathbb{F}}(B_T) \le \prod_{h \in T} \operatorname{rank}_{\mathbb{F}}(B_h) \le n^{|T|}.$$

For $T = \emptyset$, B_T is the all-ones matrix whose rank is 1. Since

$$P = \sum_{T \subseteq [d]} c_T B_T,$$

rank subadditivity gives

$$\operatorname{rank}_{\mathbb{F}}(P) \le \sum_{T \subseteq [d]} n^{|T|} = (n+1)^d.$$

□

3 Proofs

Throughout this section, let $\mathcal{U} = \{u_1, \ldots, u_K\}$ and $\mathcal{V} = \{v_1, \ldots, v_K\}$ be a typical 3-restricted MVF modulo $m = p^a q^b$. For $A \subseteq [K]$, write

$$u_A := \sum_{i \in A} u_i.$$

For $w \in \mathbb{Z}_m^n$ and $d \geq 1$, define

$$C_w^{(d)} = \left\{ A \in \binom{[K]}{d} : u_A = w \right\}.$$

Our main lemma states that for each $w \in \mathbb{Z}_m^n$, there cannot be a large moonflower in $C_w^{(d)}$, the family of d-subsets whose corresponding matching vectors in $\mathcal{U}$ sum to w.

Lemma 11 (Equal-sum moonflower bound)**.** *Let $\mathcal{U}, \mathcal{V}$ be a typical 3-restricted MVF modulo $m = p^a q^b$. For every $w \in \mathbb{Z}_m^n$ and every $d \geq 1$, if*

$$A_1, \ldots, A_L \in C_w^{(d)}$$

is a moonflower, then

$$L \leq d\left((n+1)^d + 1\right).$$

Proof. Fix $w \in \mathbb{Z}_m^n$, $d \geq 1$, and a moonflower

$$A_1, \ldots, A_L \in C_w^{(d)}.$$

The proof has three main components. We begin with the full matrix of inner products

$$M_{ij} := \langle u_j, v_i \rangle \in \mathbb{Z}_m.$$

The diagonal entries of M are 0, while every off-diagonal entry lies in $S = \{\alpha, \beta\}$.

Moving from $\mathbb{Z}_m$ to $\mathbb{F}_q$. First, we reduce all inner products modulo q^b. Since

$$\alpha \in q^{b-1}\mathbb{Z}_{q^b}, \qquad \beta \equiv 0 \pmod{q^b},$$

and the diagonal entries are 0, all entries lie in $q^{b-1}\mathbb{Z}_{q^b}$. Define

$$Q_{ij} = \tau_q(\langle u_j, v_i \rangle).$$

By Lemma 9, we know that

$$\operatorname{rank}_{\mathbb{F}_q}(Q) \leq n.$$

Moreover, since $\alpha \not\equiv 0 \pmod{q^b}$, there is a nonzero $\gamma \in \mathbb{F}_q$ such that

$$Q_{ij} = \gamma \cdot \mathbf{1}[\langle u_j, v_i \rangle = \alpha].$$

Let $\Pi = \gamma^{-1} Q$, then we have $\operatorname{rank}_{\mathbb{F}_q}(\Pi) = \operatorname{rank}_{\mathbb{F}_q}(Q) \leq n$. Note that $\Pi_{ij} = 1[\langle u_j, v_i \rangle = \alpha]$. Π is a low-rank indicator matrix for the residue α.

From elements to sums. Now, for each $h \in [L]$, choose a unique element

$$i_h \in A_h, \qquad i_h \notin A_t \quad \text{for all } t \neq h.$$

We will use Π to build a new matrix indexed by the sets $A_1, \ldots, A_L$, capturing the interaction between the sums u_{A_t} and the unique element v_{i_h}. Define

$$\ell_h = |\{j \in A_h : \langle u_j, v_{i_h} \rangle = \alpha\}|.$$

Since $i_h \in A_h$ gives the diagonal inner product 0, we have $\ell_h \in \{0, \ldots, d-1\}$. By the pigeonhole principle, some value ℓ occurs for at least L/d of the sets. Restricting to those sets and relabeling, we obtain a sub-moonflower $A_1, \ldots, A_{L'}$ with $L' \geq L/d$ and $\ell_h = \ell$ throughout.

For $h \neq t$, let

$$r_{ht} = |\{j \in A_t : \langle u_j, v_{i_h} \rangle = \alpha\}|.$$

Since $i_h \notin A_t$, all terms in $\langle u_{A_t}, v_{i_h} \rangle$ are off-diagonal. Since $u_{A_h} = u_{A_t}$, taking their respective inner products with v_{i_h} gives

$$r_{ht}\alpha + (d - r_{ht})\beta \equiv \ell\alpha + (d - 1 - \ell)\beta \pmod{m}.$$

Thus r_{ht} must lie in the set

$$T_\ell = \{r \in \{0, \ldots, d\} : r\alpha + (d-r)\beta \equiv \ell\alpha + (d-1-\ell)\beta \pmod{m}\}.$$

Crucially, $\ell \notin T_\ell$, since $r = \ell$ would force $\beta \equiv 0 \pmod{m}$. Therefore, along the diagonal block pair (h, h) the row indexed by i_h sees exactly ℓ occurrences of the residue α, while along every off-diagonal block pair (h, t) it sees a number of occurrences belonging to T_ℓ, a set that avoids ℓ.

Designing the polynomial method matrix. We can encode this with a function $g : \{0,1\}^d \to \mathbb{F}_q$ such that

$$g(z) = 0 \quad \text{if } |z| = \ell, \qquad g(z) = 1 \quad \text{if } |z| \in T_\ell.$$

Let $f \in \mathbb{F}_q[z_1, \ldots, z_d]$ be the unique multilinear polynomial agreeing with g on the Boolean cube. Fix orderings $A_t = \{j_{t,1}, \ldots, j_{t,d}\}$ and define

$$P_{ht} := f(\Pi_{i_h, j_{t,1}}, \ldots, \Pi_{i_h, j_{t,d}}).$$

By construction, the input has Hamming weight ℓ when $h = t$, and Hamming weight in T_ℓ when $h \neq t$. Hence $P = J - I$, so

$$L/d - 1 \leq L' - 1 \leq \operatorname{rank}_{\mathbb{F}_q}(P). \tag{1}$$

Now, for $r \in [d]$, define

$$(B_r)_{ht} := \Pi_{i_h, j_{t,r}}.$$

Each B_r is obtained from Π by selecting rows and columns, possibly with repetitions, so $\operatorname{rank}_{\mathbb{F}_q}(B_r) \leq n$. Since P is obtained by applying f entrywise to $B_1, \ldots, B_d$, Lemma 10 gives

$$\operatorname{rank}_{\mathbb{F}_q}(P) \leq (n+1)^d. \tag{2}$$

Combining Equation (1) and Equation (2) completes our proof. □

Using this lemma, we now complete the proof of the main theorem by an induction on the sumset size.

Proof of Theorem 3. First, we may assume without loss of generality that $(\mathcal{U}, \mathcal{V})$ does not remain a 3-restricted MVF modulo any proper divisor $m' \mid m$. Indeed, if it does, replace m by a divisor-minimal modulus with this property. If this modulus is a prime power, then the known prime-power bound, together with $m \le \sqrt{n}$, already gives the desired result. Otherwise, by Lemmas 5 and 7, we may assume that $(\mathcal{U}, \mathcal{V})$ is typical.

Fix an integer $s \ge 2$. For $1 \le d \le s$, define

$$F_d := \max_{u \in \mathbb{Z}_m^n} |C_u^{(d)}|.$$

We first show that $F_s \le n^{3s^2}$. Let us fix u and let $\mathcal{F} \subseteq C_u^{(s)}$ be a largest moonflower of size $M_s \le 2s(n+1)^s$. Set

$$W := \bigcup_{S \in \mathcal{F}} S.$$

Then $|W| \le sM_s$. Moreover, every $T \in C_u^{(s)}$ intersects W, otherwise $\mathcal{F} \cup \{T\}$ would be a larger moonflower. Hence

$$\left|C_u^{(s)}\right| \le \sum_{i \in W} \left|\{T \in C_u^{(s)} : i \in T\}\right| \le |W| F_{s-1} \le sM_s F_{s-1}.$$

Here the second inequality follows by mapping each set T to a set in $C_{u-u_i}^{(s-1)}$ for some $i \in W \subseteq [K]$. Taking the maximum over u gives

$$F_s \le sM_s F_{s-1}.$$

Plugging in Lemma 11, we obtain that

$$F_s \le 2s^2(n+1)^s F_{s-1}.$$

Since the u_i's are distinct, $F_1 = 1$. Hence

$$F_s \le \prod_{d=2}^{s} 2d^2(n+1)^d \le 2^s (s!)^2 (n+1)^{s(s+1)/2} \le n^{3s^2}.$$

when $s \ge 2$ and $n \ge 4$. The sets $C_w^{(s)}$ partition $\binom{[K]}{s}$ as w ranges over $\mathbb{Z}_m^n$, so

$$\binom{K}{s} = \sum_{w \in \mathbb{Z}_m^n} |C_w^{(s)}| \le m^n F_s \le m^n n^{3s^2}.$$

Equivalently,

$$\log \binom{K}{s} \le n \log m + 3s^2 \log n.$$

Now, set

$$s = \left\lfloor \frac{\log K}{24 \log n} \right\rfloor.$$

If $s < 2$, then $\log K = O(\log n)$. Otherwise, we have

$$\begin{aligned} n \log m &\geq \log \binom{K}{s} - 3s^2 \log n \\ &\geq s \log(K/s) - 3s^2 \log n \\ &\geq c \frac{(\log K)^2}{\log n}, \end{aligned}$$

for some absolute constant $c > 0$, where the last inequality follows by plugging in s. Rearranging gives

$$K \leq 2^{O(\sqrt{n \log m \log n})}.$$

□

References

[ABL26] 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*, pages 365–398, Cham, 2026. Springer Nature Switzerland. 1

[ADL+25] Divesh Aggarwal, Pranjal Dutta, Zeyong Li, Maciej Obremski, and Sidhant Saraogi. Improved Lower Bounds for 3-Query Matching Vector Codes. In Raghu Meka, editor, *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, volume 325 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:19, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. 3, 4

[AGKM23] Omar Alrabiah, Venkatesan Guruswami, Pravesh K Kothari, and Peter Manohar. A near-cubic lower bound for 3-query locally decodable codes from semirandom csp refutation. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1438–1448, 2023. 2

[AO26] Divesh Aggarwal and Maciej Obremski. Subexponential upper bounds for 3-restricted matching vector families. Technical Report TR26-141, Electronic Colloquium on Computational Complexity, 2026. 4

[BDL14] Abhishek Bhowmick, Zeev Dvir, and Shachar Lovett. New bounds for matching vector families. *SIAM Journal on Computing*, 43(5):1654–1683, 2014. 2, 3, 4

[BHKL25] Arpon Basu, Jun-Ting Hsieh, Pravesh K. Kothari, and Andrew D. Lin. Improved lower bounds for all odd-query locally decodable codes. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1262–1285, 2025. 2

[DGY11] Zeev Dvir, Parikshit Gopalan, and Sergey Yekhanin. Matching vector codes. *SIAM Journal on Computing*, 40(4):1154–1178, 2011. 1, 2

[Efr09] Klim Efremenko. 3-query locally decodable codes of subexponential length. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 39–44, 2009. 1, 2

[GGMT26] W. Timothy Gowers, Ben J. Green, Freddie Manners, and Terence Tao. Marton's conjecture in abelian groups with bounded torsion. *Annales de la Faculté des sciences de Toulouse : Mathématiques*, Ser. 6, 35(1):1–33, 2026. 2

[GKS25] Fatemeh Ghasemi, Swastik Kopparty, and Madhu Sudan. Improved pir schemes using matching vectors and derivatives. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, page 1648–1656, New York, NY, USA, 2025. Association for Computing Machinery. 1

[GKST06] Oded Goldreich, Howard Karloff, Leonard J Schulman, and Luca Trevisan. Lower bounds for linear locally decodable codes and private information retrieval. *computational complexity*, 15:263–296, 2006. 2

[Gro95] Vince Grolmusz. On the weak mod m representation of Boolean functions. *Chicago Journal of Theoretical Computer Science*, 1995(2), July 1995. 1

[Gro00] Vince Grolmusz. Superpolynomial size set-systems with restricted intersections mod 6 and explicit ramsey graphs. *Combinatorica*, 20(1):71–86, January 2000. 1, 2

[HPR26] Alexandra Henzinger, Edward Pyne, and Seyoon Ragavan. Catalytic tree evaluation from matching vectors, 2026. 1

[JM25] Oliver Janzer and Peter Manohar. A kq/q-2 lower bound for odd query locally decodable codes from bipartite kikuchi graphs. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 75–85, 2025. 2

[Kd04] Iordanis Kerenidis and Ronald de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. *Journal of Computer and System Sciences*, 3(69):395–420, 2004. 2

[KT00] Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 80–86, 2000. 2

[LMW26] Shachar Lovett, Raghu Meka, and Yimeng Wang. Moonflowers and efficient code sparsification, 2026. 3

[Yek08] Sergey Yekhanin. Towards 3-query locally decodable codes of subexponential length. *Journal of the ACM (JACM)*, 55(1):1–16, 2008. 1, 2

[ZBS11] Noga Zewi and Eli Ben-Sasson. From affine to two-source extractors via approximate duality. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 177–186, New York, NY, USA, 2011. Association for Computing Machinery. 3

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il