

Improved batch code lower bounds

Ray Li

Department of Computer Science
Stanford University
rayyli@cs.stanford.edu

Mary Wootters

Departments of Computer Science and Electrical Engineering
Stanford University
marykw@stanford.edu

Abstract—Batch codes are a useful notion of locality for error correcting codes, originally introduced in the context of distributed storage and cryptography. Many constructions of batch codes have been given, but few lower bound (limitation) results are known, leaving gaps between the best known constructions and best known lower bounds. Towards determining the optimal redundancy of batch codes, we prove a new lower bound on the redundancy of batch codes. Specifically, we study (primitive, multiset) *linear* batch codes that systematically encode n information symbols, with the requirement that any multiset of k symbol requests can be obtained in disjoint ways. We show that such batch codes need $\Omega(\sqrt{nk})$ symbols of redundancy, improving on the previous best lower bounds of $\Omega(\sqrt{n} + k)$ at all $k = n^\varepsilon$ with $\varepsilon \in (0, 1)$. Our proof follows from analyzing the dimension of the order- $O(k)$ tensor of the batch code’s dual code.

Index Terms—Batch Codes, Lower bound, Locality, Tensor product, Distributed Storage

I. INTRODUCTION

In this work, we study *batch codes*, a notion of locality for error correcting codes, and show stronger limitations on batch codes for almost all parameter regimes.

Batch codes were introduced in the context of load-balancing in distributed storage and private information retrieval in cryptography [1]. Informally, a (primitive, multiset) k -batch code is a error correcting code $C : \Sigma^n \rightarrow \Sigma^N$, mapping n information symbols to N codeword symbols, such that every multiset of k information symbols can be recovered from k pairwise disjoint recovering sets. In constructing batch codes, we would like the *locality* parameter k to be as large as possible. On the other hand, we would like to minimize the *redundancy* $N - n$ of our code, representing the number of redundant bits in our encoding. In this work, we prove new limitations on the quantitative tradeoff between the locality and the redundancy of batch codes. Formally, a batch code is defined as follows.

Definition I.1. Let $C : \Sigma^n \rightarrow \Sigma^N$ be a code that maps x_1, \dots, x_n to c_1, \dots, c_N . The code C is a k -batch code if, for every multiset of indices $\{i_1, \dots, i_k\} \subset [n]$, there exist k mutually disjoint sets $R_1, \dots, R_k \subset [N]$ and functions g_1, \dots, g_k such that for all information symbols $x_1, \dots, x_n \in \Sigma$ and codewords $(c_1, \dots, c_N) = C(x_1, \dots, x_n) \in \Sigma^N$ and for all $j \in [k]$, we have $g_j(c_{R_j}) = x_{i_j}$, where c_{R_j} is c restricted

to coordinates in R_j . We say C is a *linear* batch code if Σ is a finite field, the functions g_i are all linear, and C is a linear map. We say C has a *systematic encoding* if $c_i = x_i$ for $i = 1, \dots, n$.

Remark I.2 (Primitive Multiset Batch Codes). We note that our definition of a batch code here is a refinement of the standard, more general notion of batch codes, introduced by [1]. In that definition of a batch code, n symbols are encoded into “buckets” of symbols such that the total size of all buckets is N , and each batch of k symbols can be decoded by reading at most one symbol from each bucket. A batch code is a *multiset* batch code if (a) the k symbols can form a multiset and (b) the k symbols can be decoded by querying k pairwise disjoint sets of buckets. When each bucket can store a single symbol, the multiset batch code is said to be *primitive*. Because we only focus on primitive multiset batch codes in this work, we drop the adjectives “primitive” and “multiset” throughout and simply use “batch code,” as per Definition I.1.

The main goal in constructing batch codes is to determine the minimum redundancy $r(n, k)$ of a k -batch code encoding n symbols. Many works [1]–[9], have constructed batch codes, giving good upper bounds on $r(n, k)$. Figure I and Section I-A give a summary of the known constructions.

On the other hand, few limitations are known on the optimal locality versus redundancy tradeoff $r(n, k)$. An easy lower bound shows that a k -batch code has minimum Hamming distance k , and thus must have redundancy at least k . The only nontrivial lower bound on the redundancy of (linear) batch codes is given by [10], who showed that a k -batch code has redundancy at least $\Omega(\sqrt{n})$ when $k \geq 3$, and this is tight up to a logarithmic factor for constant k [4]. A priori, given these two lower bounds, and the fact that there exist n^ε batch codes with redundancy n^δ for every $\varepsilon \in (0, 1)$ and for $\delta = \delta(\varepsilon) < 1$ [8], it seemed possible that the optimal δ could match the best known lower bounds at $\max(1/2, \varepsilon)$. In our work, we refute this possibility for all $\varepsilon \in (0, 1)$, showing that the redundancy must be $\Omega(\sqrt{nk})$ (under two reasonable assumptions—a systematic encoding and linearity—that also appear in prior lower bounds [10]). We note that Theorem I.3 holds for any values $3 \leq k \leq n$, not just values in the regime $k = n^\varepsilon$ for $\varepsilon \in (0, 1)$.

Theorem I.3. *There exists an absolute constant $c > 0$ such that, for all positive integers $3 \leq k \leq n$, a linear k -batch code*

RL is supported by NSF Grants DGE-1656518, CCF-1814629. MW is partially supported by NSF Grant CCF-1844628 and by a Sloan Research Fellowship.




Fig. 1. $r(n, k)$, the minimum possible redundancy of a k -batch code encoding n information symbols. Blue dashed segments indicate code constructions. Red solid segments indicate lower bounds for linear batch codes. Our lower bound additionally assumes systematic encoding.

systematically encoding n symbols must have redundancy at least $c\sqrt{nk}$.

Though our lower bound assumes linearity and a systematic encoding, we point out that the only other nontrivial lower bounds for batch codes [10] also assume linearity. Furthermore, [10] proves the $\Omega(\sqrt{n})$ lower bound in two ways, with one proof additionally assuming systematic encoding, and with the other proof handling non-systematic encoding but additionally assuming that the code is binary (our lower bound works for linear batch codes over any field). We also point out that many existing constructions of batch codes are linear and have systematic encoding.

A. Related work

a) *Prior constructions of batch codes:* We would like to understand $r(n, k)$, the minimum possible redundancy of a k -batch code encoding n information symbols. The following upper bounds on $r(n, k)$ are known, with the best known ones illustrated in Figure I, along with the known lower bounds.

- $r(n, k) \leq O(k^4)$, for $k = n^\varepsilon$ with $1/5 < \varepsilon \leq 7/32$ [2]
- $r(n, n^{1/4}) \leq n^{7/8}$ [2]
- $r(n, k) \leq \tilde{O}(\sqrt{n})$ for any fixed k [4].
- $r(n, n^\varepsilon) \leq O(n^{2/3+5\varepsilon/3})$ for $\varepsilon < 1/2$ [5].
- $r(n, n^\varepsilon) \leq O(n^{5/6+\varepsilon/3})$ for $\varepsilon < 1/2$ [5].
- $r(n, n^\varepsilon) \leq \tilde{O}(n^{(3\varepsilon+1)/2})$ for $0 < \varepsilon < 1/3$ [6].
- $r(n, n^\varepsilon) \leq \tilde{O}(n^{\log_4(3)+(2-\log_2(3))\varepsilon})$ for $0 < \varepsilon < 1/2$ [7].
- $r(n, n^\varepsilon) \leq O(n^\delta)$ for $0 \leq \varepsilon < 1$, where $\delta = \delta(\varepsilon) < 1$ [8]

b) *Other notions of locality:* Batch codes are related to several other notions of locality. Two closely related notions are private information retrieval (PIR) codes [5], [11] and codes with the disjoint repair group property (DRGP) [12]–[15], which are relaxations of batch codes that only require k disjoint repair groups for a single information symbol

(codeword symbol in the case of DRGP). In the case that the locality parameter k is linear in the block length, PIR codes and DRGP codes are in fact equivalent to constant query locally decodable codes (LDCs) and locally correctable codes (LCCs), respectively. Other related notions include Locally Repairable Codes [16], LRCs with availability [17], batch codes with availability [18] where the repair group sizes are also bounded, switch codes [19]–[21], which are a special case of batch codes, and combinatorial batch codes [22], which are a special case of (non-primitive) multiset batch codes. In these settings, determining the optimal locality versus redundancy tradeoff is an interesting question. We hope our techniques could be useful for proving lower bounds for some of these other notions of locality. For more details about some of these other notions, we refer the reader to the survey [23].

B. Preliminaries

a) *Basic notation:* We use $[N]$ to denote the set $\{1, \dots, N\}$. For a vector $c \in \mathbb{F}^N$, let $\text{supp}(c) \stackrel{\text{def}}{=} \{i \in [N] : c_i \neq 0\}$. For two vector spaces V, W over \mathbb{F} , we use $V \leq W$ to denote V is a subspace of W .

By abuse of notation, we identify a code $C : \mathbb{F}^n \rightarrow \mathbb{F}^N$ with a systematic encoding by its image $C \leq \mathbb{F}^N$, which is a subspace of \mathbb{F}^N . In the rest of this paper, all codes have systematic encoding and are represented this way.

b) *Dual codes:* The dual code $C^\perp \leq \mathbb{F}^N$ of a linear code $C \leq \mathbb{F}^N$ is the subspace of all codewords orthogonal to every codeword of C . One can easily check that $\dim C^\perp = N - \dim C$, so in particular the dimension of C^\perp is equal to the redundancy of C . We call elements of C^\perp *dual codewords (of C)*. We use the following standard fact of dual codes, and include a proof for completeness.

Lemma I.4. *Suppose $C \leq \mathbb{F}^N$ is a linear code with a systematic encoding, and there exists an index $i \in [N]$, a set $R \subset [N] \setminus \{i\}$, and a linear function g , such that $g(c|_R) = c_i$ for all codewords $c \in C$. Then there exists a nonzero dual codeword c^\perp such that $\text{supp}(c^\perp) \subset R \cup \{i\}$ and $i \in \text{supp}(c^\perp)$.*

Proof. Suppose g is the function $g(c|_R) = \sum_{j \in R} \alpha_j c_j$. Then $-c_i + \sum_{j \in R} \alpha_j c_j = 0$ for all codewords c . Thus, the vector c^\perp with $c_j^\perp = -1$ if $j = i$, $c_j^\perp = \alpha_j$ if $j \in R$, and $c_j^\perp = 0$ otherwise, is a dual codeword. Vector c^\perp also satisfies $\text{supp}(c^\perp) \subset R \cup \{i\}$ and $i \in \text{supp}(c^\perp)$ by construction, as desired. \square

c) *Tensor products:* We let e_i denote the standard basis vector in \mathbb{F}^N , so that $(e_i)_i = 1$ and $(e_i)_j = 0$ for all $j \neq i$. For $v^{(1)}, \dots, v^{(s)} \in \mathbb{F}^N$, we define the tensor product $v^{(1)} \otimes \dots \otimes v^{(s)} \in \mathbb{F}^{N^s}$ to be the vector indexed by tuples $(i_1, \dots, i_s) \in [N]^s$ with

$$(v^{(1)} \otimes \dots \otimes v^{(s)})_{(i_1, \dots, i_s)} = \prod_{j=1}^s v_{i_j}^{(j)}.$$

A tensor of this form is called a *simple tensor*; more generally a *tensor* is any linear combination of simple tensors. We note that the set of simple tensors

$$\left\{ \bigotimes_{j=1}^s e_{i_j} : (i_1, \dots, i_s) \in [N]^s \right\}$$

forms a basis for \mathbb{F}^{N^s} . We refer to this as the standard tensor basis for \mathbb{F}^{N^s} . Accordingly, every vector in \mathbb{F}^{N^s} can be written as a linear combination of the standard tensor basis, which we call the *standard basis representation*. Because these tensors form a basis, we have the following useful fact.

Lemma I.5. *Let $(e^{(1)}, w^{(1)}), \dots, (e^{(D)}, w^{(D)})$ be pairs of tensors in \mathbb{F}^{N^s} such that, for all $i = 1, \dots, D$, tensor $e^{(i)}$ is in the standard tensor basis and the standard basis representation of $w^{(i)}$ contains $e^{(i)}$ and none of $e^{(i+1)}, \dots, e^{(D)}$. Then $w^{(1)}, \dots, w^{(D)}$ are linearly independent.*

Proof. Suppose for contradiction $\sum_{i=1}^D \alpha_i w^{(i)} = 0$ is a nonzero linear combination of $w^{(1)}, \dots, w^{(D)}$. Let j be the largest index such that $\alpha_j \neq 0$. By definition of $w^{(j)}$, we may write $w^{(j)} = \beta e^{(j)} + w'$ for some $\beta \neq 0$ and some $w' \in \mathbb{F}^{N^s}$ not containing $e^{(j)}$ in its standard basis representation. Then the coefficient of $e^{(j)}$ in the standard basis representation of $\sum_{i=1}^D \alpha_i w^{(i)} = (\sum_{i < j} \alpha_i w^{(i)} + \alpha_j w') + \alpha_j \beta e^{(j)}$ is exactly $\alpha_j \beta$, since $w^{(1)}, \dots, w^{(j-1)}$ and w' do not contain $e^{(j)}$ in their standard basis representation, so $\alpha_j \beta = 0$. This contradicts that $\alpha_j, \beta \neq 0$. \square

For a subspace $V \leq \mathbb{F}^N$ and integer s , we let $V^{\otimes s}$ denote the subspace of \mathbb{F}^{N^s} spanned by simple tensors $v^{(1)} \otimes \dots \otimes v^{(s)}$ for $v^{(1)}, \dots, v^{(s)} \in V$. A standard fact says that $\dim V^{\otimes s} = (\dim V)^s$, for all subspaces V and positive integers s .

II. PROOF OF THE MAIN THEOREM

A. Sketch of the proof

Though our proof is short, we provide a brief summary to highlight the main ideas. For intuition in this sketch, we strengthen the definition of batch codes to require that the repair groups R_1, \dots, R_k are not only pairwise disjoint but also that they are each disjoint from $\{i_1, \dots, i_k\}$. In the main proof, we remove this assumption.

Let C be a k -batch code. Let $t = \lfloor k/2 \rfloor$. Since $2t \leq k$, C is also a $2t$ -batch code by definition. Let $V = C^\perp$. Call a simple tensor $e_{i'_1} \otimes \dots \otimes e_{i'_{2t}} \in \mathbb{F}^{N^{2t}}$ *good* if the multiset $\{i'_1, \dots, i'_{2t}\} \subset [n]$ contains exactly t distinct elements, each appearing exactly twice. Given such a multiset $\{i'_1, i'_2, \dots, i'_{2t}\}$ for a good tensor, the definition of a $2t$ -batch code together with Lemma I.4 guarantee the existence of $2t$ dual codewords, one for each i'_j , so that the support of the corresponding codeword is $R_j \cup \{i'_j\}$, and where (by assumption for this sketch) the R_j are all disjoint from each other and $\{i'_1, \dots, i'_{2t}\}$. By tensoring these dual codewords, we obtain a tensor $w_{i'_1, \dots, i'_{2t}} \in V^{\otimes 2t}$ whose standard basis representation contains the good simple tensor $\bigotimes_{j=1}^{2t} e_{i'_j}$ and, crucially, no other good simple tensors.

Because these tensors $w_{i'_1, \dots, i'_{2t}}$ have only one good simple tensor each, unique to them, they must be linearly independent by Lemma I.5. Thus, the dimension of $V^{\otimes 2t}$ is lower bounded by the number of tensors $w_{i'_1, \dots, i'_{2t}}$, which is at least the number of good simple tensors, which we can count to be $\binom{n}{t} \binom{2t}{2, 2, \dots, 2} \geq \Omega(nt)^t$. As $\dim(V^{\otimes 2t}) = (\dim V)^{2t}$, we conclude that $\dim V \geq (\Omega(nt)^t)^{1/2t} = \Omega(\sqrt{nt}) = \Omega(\sqrt{nk})$, as desired.

B. Full proof

Proof of Theorem I.3. Let $t = \lfloor k/3 \rfloor$. Since C is a k -batch code, C is also a $3t$ -batch code. Define vector spaces $V \stackrel{\text{def}}{=} C^\perp$ and $W \stackrel{\text{def}}{=} V^{\otimes 2t}$. In the standard tensor basis of $\mathbb{F}^{N^{2t}}$, call a simple tensor $e_{i'_1} \otimes \dots \otimes e_{i'_{2t}}$ *good* if the multiset $\{i'_1, \dots, i'_{2t}\}$ contains exactly t distinct elements, each appearing exactly twice.

For every t -tuple $(i_1, \dots, i_t) \in [n]^t$ with $i_1 > \dots > i_t$, consider the t -multiset of symbols $\cup_{j=1}^t \{x_{i_j}, x_{i_j}, x_{i_j}\}$. Since C is a $3t$ -batch code, these have recovery sets $R_{j,1}, R_{j,2}, R_{j,3}$ for $j = 1, \dots, t$ that are pairwise disjoint. For any j , at least two of the $R_{j,1}, R_{j,2}, R_{j,3}$ do not contain i_j because they are pairwise disjoint. Without loss of generality, $R_{j,1}$ and $R_{j,2}$ do not contain i_j . Thus, for each j , applying Lemma I.4 to $R_{j,1}$ and $R_{j,2}$, there exist two dual codewords $c^{(j,1)}, c^{(j,2)}$ such that $\text{supp}(c^{(j,1)}) \cap \text{supp}(c^{(j,2)}) = \{i_j\}$. Define $\mathcal{D}_{i_1, \dots, i_t} \stackrel{\text{def}}{=} (c^{(j,\ell)})_{1 \leq j \leq t, \ell=1,2}$.

Call a bijection $\pi : [2t] \rightarrow [t] \times [2]$ a *good map*. For a good map π , let $\pi_1 : [2t] \rightarrow [t]$ denote the first coordinate of π and let $\pi_2 : [2t] \rightarrow [2]$ denote the second coordinate of π . For each good map π and each i_1, \dots, i_t , let $e_{i_1, \dots, i_t, \pi}$ denote the simple tensor

$$e_{i_1, \dots, i_t, \pi} \stackrel{\text{def}}{=} \bigotimes_{j=1}^{2t} e_{i_{\pi_1(j)}}. \quad (1)$$

Clearly $e_{i_1, \dots, i_t, \pi}$ is a good simple tensor, since π is a bijection and thus each $\pi_1(j)$ appears twice.

For each good map π and each i_1, \dots, i_t , if $\mathcal{D}_{i_1, \dots, i_t} = (c^{(j,\ell)})_{1 \leq j \leq t, \ell=1,2}$, then let $w_{i_1, \dots, i_t, \pi}$ denote the tensor

$$w_{i_1, \dots, i_t, \pi} \stackrel{\text{def}}{=} \bigotimes_{j=1}^{2t} c^{(\pi_1(j), \pi_2(j))}. \quad (2)$$

As $c^{(\pi_1(j), \pi_2(j))}$ are all in V , each $w_{i_1, \dots, i_t, \pi}$ is in $W = V^{\otimes 2t}$.

Claim II.1. *Each $w_{i_1, \dots, i_t, \pi}$, written in the standard basis of $\mathbb{F}^{N^{2t}}$ has at most 3^t good simple tensors.*

Proof. Note that $\text{supp}(c^{(j,\ell)}) \subseteq \{i_j\} \cup R_{j,\ell}$, and furthermore that $R_{j,\ell}$ are pairwise disjoint. Thus, each of i_1, \dots, i_t appears in at most three supports $\text{supp}(c^{(j,\ell)})$; that is, i_j can appear in $\text{supp}(c^{(j,1)})$, $\text{supp}(c^{(j,2)})$, and then in $\text{supp}(c^{(j',\ell)})$ for at most one other (j', ℓ) , so that $i_j \in R_{j',\ell}$. Further, any element $i \in [N] \setminus \{i_1, \dots, i_t\}$ appears in at most one set $R_{j,\ell}$ and hence in at most one support $\text{supp}(c^{(j,\ell)})$. Thus, any simple tensor $e_{i'_1} \otimes \dots \otimes e_{i'_{2t}}$ appearing in the standard basis representation of $w_{i_1, \dots, i_t, \pi}$ has at most three of any of e_{i_1}, \dots, e_{i_t} and at

most one any other e_i in the product. Thus, in any good simple tensor $e_{i'_1} \otimes \cdots \otimes e_{i'_{2t}}$ in the standard basis representation of $w_{i_1, \dots, i_t, \pi}$, there must be two of each of e_{i_1}, \dots, e_{i_t} in the product. For each e_{i_j} , there are at most three of the $2t$ positions where it can appear in the standard basis representation: in the (two) positions j' such that $\pi(j')_1 = j$, or the (at most one) position j' such that $R_{\pi(j')_1, \pi(j')_2}$ contains $\{i_j\}$. Thus, there are at most $\binom{3}{2} = 3$ choices for the positions of e_{i_j} for each j , so there are at most 3^t good tensors in w_{i_1, \dots, i_t} . \square

Claim II.2. Each $w_{i_1, \dots, i_t, \pi}$, written in the standard basis of $\mathbb{F}^{N^{2t}}$ contains the simple tensor $e_{i_1, \dots, i_t, \pi}$.

Proof. For each $j = 1, \dots, 2t$, the dual codeword $c^{(\pi_1(j), \pi_2(j))}$ has a nonzero coefficient in coordinate $i_{\pi_1(j)}$, so $w_{i_1, \dots, i_t, \pi} = \bigotimes_{j=1}^{2t} c^{(\pi_1(j), \pi_2(j))}$ has a nonzero coefficient in coordinate $(i_{\pi_1(1)}, i_{\pi_1(2)}, \dots, i_{\pi_1(2t)})$. Hence, the standard basis representation of $w_{i_1, \dots, i_t, \pi}$ contains the simple tensor $\bigotimes_{j=1}^{2t} e_{i_{\pi_1(j)}} = e_{i_1, \dots, i_t, \pi}$. \square

Now consider the set E of all good simple tensors $e_{i_1, \dots, i_t, \pi}$ with $i_1 > \cdots > i_t$ and π a good map. Note that there are exactly $\binom{n}{t} \binom{2t}{2, 2, \dots, 2}$ elements of E , as we may first choose $i_1 > \cdots > i_t$ and then choose π_1 . (Notice that the definition of $e_{i_1, \dots, i_t, \pi}$ does not depend on π_2). Choose a sequence $(e^{(1)}, w^{(1)}), (e^{(2)}, w^{(2)}), \dots \in E \times W$ as follows: given $(e^{(1)}, w^{(1)}), \dots, (e^{(r)}, w^{(r)})$, choose $e^{(r+1)} = e_{i_1, \dots, i_t, \pi}$ to be a good simple tensor not in the standard basis representation of any of $w^{(1)}, \dots, w^{(r)}$, and let $w^{(r+1)} \stackrel{\text{def}}{=} w_{i_1, \dots, i_t, \pi}$, which contains $e^{(r+1)}$ by Claim II.2. By Claim II.1, since each $w^{(r)}$ has at most 3^t good simple tensors in the standard basis representation, this process can be continued for at least $D \stackrel{\text{def}}{=} |E|/3^t$ steps. Furthermore, we guarantee that $w^{(r)}$ contains $e^{(r)}$ and none of $e^{(r+1)}, \dots, e^{(D)}$ for any $r = 1, \dots, D$. Thus, by Lemma I.5, the tensors $w^{(1)}, \dots, w^{(D)}$ are linearly independent. Hence, we have

$$\begin{aligned} (\dim V)^{2t} &= \dim W \geq D = |E|/3^t = \binom{n}{t} \cdot \frac{(2t)!}{2^t} \cdot \frac{1}{3^t} \\ &\geq \binom{n}{t}^t \cdot \frac{(2t)^{2t}/e^{2t}}{2^t \cdot 3^t} = \frac{n^{2t}}{(3e^2/2)^t}. \end{aligned} \quad (3)$$

Hence, the redundancy is at least

$$\dim V \geq \frac{1}{\sqrt{3e^2/2}} \cdot \sqrt{nt} \geq 0.1\sqrt{nk}, \quad (4)$$

as desired. Here, we used that $t = \lfloor k/3 \rfloor \geq k/5$. \square

REFERENCES

- [1] Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai, "Batch codes and their applications," in *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pp. 262–271, 2004.
- [2] A. S. Rawat, Z. Song, A. G. Dimakis, and A. Gál, "Batch codes through dense graphs without short cycles," *IEEE Transactions on Information Theory*, vol. 62, no. 4, pp. 1592–1604, 2016.
- [3] H. Lipmaa and V. Skachek, "Linear batch codes," in *Coding theory and applications*, pp. 245–253, Springer, 2015.
- [4] A. Vardy and E. Yaakobi, "Constructions of batch codes with near-optimal redundancy," in *2016 IEEE International Symposium on Information Theory (ISIT)*, pp. 1197–1201, 2016.

- [5] H. Asi and E. Yaakobi, "Nearly optimal constructions of pir and batch codes," *IEEE Transactions on Information Theory*, vol. 65, no. 2, pp. 947–964, 2019.
- [6] N. Polyanski and I. Vorobyev, "Constructions of batch codes via finite geometry," in *2019 IEEE International Symposium on Information Theory (ISIT)*, pp. 360–364, 2019.
- [7] R. Polyanskaya, N. Polyanski, and I. Vorobyev, "Binary batch codes with improved redundancy," *IEEE Transactions on Information Theory*, vol. 66, no. 12, pp. 7360–7370, 2020.
- [8] L. Holzbaur, R. Polyanskaya, N. Polyanski, and I. Vorobyev, "Lifted reed-solomon codes with application to batch codes," *IEEE International Symposium on Information Theory (ISIT)*, pp. 634–639, 2020.
- [9] L. Holzbaur, R. Polyanskaya, N. Polyanski, I. Vorobyev, and E. Yaakobi, "Lifted multiplicity codes," *arXiv preprint arXiv:2008.04717*, 2020.
- [10] S. Rao, A. Vardy, and M. Wootters, "Lower bounds on the redundancy of linear codes with disjoint repair groups," in *2022 IEEE International Symposium on Information Theory (ISIT)*, 2022.
- [11] A. Fazeli, A. Vardy, and E. Yaakobi, "Codes for distributed pir with low storage overhead," in *2015 IEEE International Symposium on Information Theory (ISIT)*, pp. 2852–2856, 2015.
- [12] V. Guruswami and C. Wang, "Linear-algebraic list decoding for variants of reed-solomon codes," *IEEE Transactions on Information Theory*, vol. 59, no. 6, pp. 3257–3268, 2013.
- [13] S. L. Frank-Fischer, V. Guruswami, and M. Wootters, "Locality via partially lifted codes," *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 2017.
- [14] R. Li and M. Wootters, "Lifted Multiplicity Codes and the Disjoint Repair Group Property," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)* (D. Achlioptas and L. A. Végh, eds.), vol. 145 of *Leibniz International Proceedings in Informatics (LIPIcs)*, (Dagstuhl, Germany), pp. 38:1–38:18, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019.
- [15] J. Hastings, A. Kanne, R. Li, and M. Wootters, "Wedge-lifted codes," *IEEE International Symposium on Information Theory (ISIT)*, 2021.
- [16] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin, "On the locality of codeword symbols," *IEEE Transactions on Information theory*, vol. 58, no. 11, pp. 6925–6934, 2012.
- [17] A. S. Rawat, D. S. Papailiopoulos, A. G. Dimakis, and S. Vishwanath, "Locality and availability in distributed storage," in *2014 IEEE International Symposium on Information Theory*, pp. 681–685, 2014.
- [18] H. Zhang and V. Skachek, "Bounds for batch codes with restricted query size," in *2016 IEEE International Symposium on Information Theory (ISIT)*, pp. 1192–1196, 2016.
- [19] Z. Wang, O. Shaked, Y. Cassuto, and J. Bruck, "Codes for network switches," in *2013 IEEE International Symposium on Information Theory*, pp. 1057–1061, 2013.
- [20] Z. Wang, H. M. Kiah, and Y. Cassuto, "Optimal binary switch codes with small query size," in *2015 IEEE International Symposium on Information Theory (ISIT)*, pp. 636–640, 2015.
- [21] Y. M. Chee, F. Gao, S. T. H. Teo, and H. Zhang, "Combinatorial systematic switch codes," in *2015 IEEE International Symposium on Information Theory (ISIT)*, pp. 241–245, 2015.
- [22] M. B. Paterson, D. R. Stinson, and R. Wei, "Combinatorial batch codes," *Advances in Mathematics of Communications*, vol. 3, no. 1, p. 13, 2009.
- [23] V. Skachek, "Batch and PIR codes and their connections to locally repairable codes," in *Network Coding and Subspace Designs*, pp. 427–442, Springer, 2018.