




Decoding Balanced Linear Codes with Preprocessing

Andrej Bogdanov  



School of Electrical Engineering and Computer Science, University of Ottawa, Canada

Rohit Chatterjee  

Department of Computer Science, National University of Singapore, Singapore

Yunqi Li  

Department of Computer Science, National University of Singapore, Singapore

Prashant Nalini Vasudevan  

Department of Computer Science, National University of Singapore, Singapore

Abstract

Prange’s information set algorithm is a well-known decoding algorithm for linear codes. It decodes corrupted codewords of most \mathbb{F}_2 -linear codes C of message length n up to relative error rate $O(\log n/n)$ in $\text{poly}(n)$ time. We show that the error rate can be improved to $O((\log n)^2/n)$, provided: (1) the decoder has access to a polynomial-length advice string that depends on C only, and (2) C is $n^{-\Omega(1)}$ -balanced.

As a consequence we improve the error tolerance in decoding random linear codes if inefficient preprocessing of the code is allowed. This reveals potential vulnerabilities in cryptographic applications of Learning Noisy Parities with low noise rate.

Our main technical result is that the Hamming weight of Hw , where the rows of H are a random sample of *short dual* codewords, measures the proximity of a received word w to the code in the regime of interest. Given such H as advice, our algorithm corrects errors by locally minimizing this measure. We show that for most codes, the error rate tolerated by our decoder is asymptotically optimal among all algorithms whose decision is based on thresholding Hw for an arbitrary polynomial-size advice matrix H .

2012 ACM Subject Classification Mathematics of computing → Coding theory

Keywords and phrases Linear codes, nearest codeword problem, learning parity with noise

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.23

Related Version *Full Version*: <https://eprint.iacr.org/2025/1932>

Funding Work supported by an NSERC Discovery Grant, and by the National Research Foundation, Singapore, under its NRF Fellowship programme, award no. NRF-NRFF14-2022-0010.

Acknowledgements We thank the anonymous ITCS 2026 reviewers for pointing out the connection to statistical decoding algorithms and other insights.

1 Introduction

Decoding corrupted codewords is a challenging algorithmic task. Its complexity is far from being understood. The choice of the code is crucial in this context. Crafted codes like Reed-Solomon can be decoded optimally, all the way up to half the minimum distance. In contrast, for random linear codes no known algorithm can substantially outperform brute force. What is the algorithmic complexity of decoding a generic code of given rate and distance? This question is captured by the *nearest codeword problem* (NCP).

NCP takes two inputs: A linear code C of message length n and blocklength $m \geq n$, and a target word w . (We restrict attention to codes over the binary alphabet.) In the *decision* version, the goal is to distinguish whether w is close to C or far from C . In the *estimation*



© Andrej Bogdanov, Rohit Chatterjee, Yunqi Li, and Prashant Nalini Vasudevan; licensed under Creative Commons License CC-BY 4.0

17th Innovations in Theoretical Computer Science Conference (ITCS 2026).

Editor: Shubhangi Saraf; Article No. 23; pp. 23:1–23:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

version, the goal is to calculate the distance from w to C up to some approximation factor. The *search* version asks for the codeword closest to w under the promise that w is sufficiently close to C .

Of the three variants, the search one is the hardest. Decoding up to relative distance η yields an estimation algorithm with approximation ratio at most $1/\eta$. It also distinguishes those words that are η -close to C from all others, provided that η is within the decoding radius.

The complexity of decoding

The randomized “information set decoding” algorithm of Prange [43] is essentially still the best asymptotically efficient one for NCP.¹ It finds the codeword closest to w in expected time $\exp(O(\eta n))$ under the promise that w is within relative distance η of C . Setting η to $O((\log m)/n)$ yields a polynomial-time algorithm with approximation ratio $n/c \log m$ for any constant c .

No asymptotic improvement to Prange’s algorithm is known even in the average-case, i.e., for random linear codes C . In the vanishing rate regime $m \gg n$, which will be of our main interest, most such codes are $O(\sqrt{n/m})$ -balanced: all codewords have Hamming weight in the range $1/2 \pm \Theta(\sqrt{n/m})$, which is close to the best possible [47, 36, 4, 40, 32].

On the hardness side, there is no polynomial-time algorithm that approximates the distance to the nearest codeword within an arbitrary constant factor unless P equals NP, or within a $O(2^{\log^{1-\epsilon} n})$ factor unless NP has quasi-polynomial-time algorithms [9, 14]. In addition, there is no sub-exponential-time approximation within *some* constant factor under the Exponential-Time Hypothesis [16].

Does this hardness stem from the choice of the code C or of the corrupted codeword w ? The nearest codeword problem *with pre-processing* is useful for studying their relative contribution. In this problem, there is a pre-processing phase, in which the algorithm sees only the code C and produces a bounded advice string H that depends on C . There is no restriction on the complexity of computing H . In the online phase, the algorithm reads w and produces its answer. The estimation variant of NCP with pre-processing is known to be NP-hard to approximate to within a fixed constant factor [31, 44], and SETH-hard to solve exactly in $2^{(1-\epsilon)n}$ time [45].

Our results

Our main result (Theorem 7) is that pre-processing improves upon Prange’s algorithm for highly balanced codes. Under the promise that C is β -balanced, the Nearest Codeword at relative distance η can be found in time $\text{poly}(m) \cdot \exp[O(\eta n / \log(1/\max\{\beta, \eta\}))]$ using advice of comparable length.

When the code is $n^{-\Omega(1)}$ -balanced, we obtain a polynomial-time algorithm with preprocessing for decoding near-codewords at relative distance $(\log n)^2/n$. In contrast, Prange’s algorithm would take quasipolynomial time to decode at this distance.

The balance assumption is satisfied by most linear codes of rate $n^{-\Omega(1)}$ (as n grows), as well as by many explicit constructions [41, 6, 5, 11, 46]. Is it necessary? In Theorem 15 we show that it is for a certain class of algorithms with preprocessing.

¹ It was rediscovered by Berman and Karpinski [12] as an approximation algorithm, and partially derandomized by Alon, Panigrahy, and Yekhanin [7].

Specifically, the advice in Algorithm 2 that solves the decision version of NCP is a matrix H that depends on the code. The algorithm decides proximity to the code based on the Hamming weight of Hw (modulo 2). This weight is small if w is close to the code and large if it is far.

Our Corollary 16 shows that for most codes C , any algorithm whose decision is the threshold of Hw for *some* matrix H cannot separate words that are $\omega((\log^2 n)/n)$ -close to the code from those that are $(1/2 - 3\sqrt{n/m})$ -far, provided H has size polynomial in n . Thus even under almost extremal assumptions on the balance of C , the decoding distance cannot be improved.

Theorem 15 is restrictive in assuming that proximity to the code is decided by a threshold of parities. In Corollary 17 we show that the same conclusion extends to a polynomial threshold function of parities. While further generalizations are an interesting open question, we cannot expect to fully eliminate the complexity assumption on the distinguisher in an unconditional lower bound. For instance, the possibility that every code can be decoded by a polynomial-size *threshold of thresholds* of parities is consistent with the current sorry state of computational complexity theory. Polynomial-size thresholds of thresholds of parities could conceivably compute all of NP.

Ideas and techniques

The algorithm

Our advice H consists of independent samples from a carefully chosen distribution D on length- m binary strings, conditioned on each sample being a dual codeword in C^\perp . We denote (the density of) this conditional distribution by D_C .

A sample of D is a sum $h = u_1 + \dots + u_\ell$ of independent standard basis (one-hot) vectors in \mathbb{F}_2^m chosen with replacement. When ℓ is even, the support of D_C is guaranteed to be nonempty as it contains the all-zero string, so D_C is well-defined.

The distribution D_C favors light dual codewords, with the weight controlled by the parameter ℓ . When ℓ is slightly larger than the distance of C^\perp , we expect D_C to be concentrated on dual codewords of weight close to the minimum distance of C^\perp .

In Lemma 9 we show that for a suitable choice of ℓ the Fourier transform

$$\hat{D}_C(w) = \mathbb{E}_{h \sim D_C} [(-1)^{\langle h, w \rangle}] = \mathbb{E}_{h \sim D} [(-1)^{\langle h, w \rangle} \mid h \in C^\perp]$$

distinguishes close from far codewords: $\hat{D}_C(w)$ is bounded away from zero when w is close to the code and extremely close to zero when it is at relative distance $1/2 - n^{-\Omega(1)}$ from all codewords and their bit-wise negations.

Our Algorithm 2 decides proximity based on an empirical estimate of $\hat{D}_C(w)$. By a large deviation bound, the efficiently computable function

$$\tilde{D}_H(w) = \mathbb{E}_{\text{row } h \text{ of } H} [(-1)^{\langle h, w \rangle}]$$

uniformly approximates $\hat{D}_C(w)$ for at least one choice (in fact, for most choices) of H . This H is the advice provided to the algorithm.

Our search algorithm (Algorithm 3) is based on one additional idea. For sufficiently light w , $\hat{D}_C(w)$ is monotonically decreasing. Thus the 1-entries in w are precisely those whose flip results in an increase of $\hat{D}_C(w)$, or of its proxy $\tilde{D}_H(w)$. As $\hat{D}_C(w)$ is periodic modulo C , the errors in w can be corrected by flipping those bits w_i that increase the value of $\tilde{D}_H(w)$.

Lower bound

To prove Theorem 15 we show that for any sufficiently short advice H , the function \tilde{D}_H cannot distinguish *random* codewords $w = Cx + e$ corrupted by random noise e exceeding the noise rate supported by our algorithm from truly random strings w in expectation. The former are typically close to the code, while the latter are very far. Thus it is impossible for a threshold distinguisher to tell them apart.

As a consequence, our algorithms cannot be improved by optimizing the advice H , and in particular by a different choice of distribution D . The key property of D used in our analysis is that the Fourier coefficients $\hat{D}(w)$ are noticeable when w is light (its relative Hamming weight is $O((\log n)^2/n)$) and vanishing when w is heavy (see (6)). In words, D is noticeably biased under all light linear tests but almost unbiased under all heavy linear tests.

Lemma 19 states that this is optimal: If D is noticeably biased under all light tests of relative Hamming weight $\omega((\log n)^2/n)$ then it must be somewhat biased under a heavy test of weight $(1 \pm n^{-\Omega(1)})/2$. This consequence may be of interest in pseudorandomness, specifically the study of small-biased sets. Lemma 19 says that one cannot ignore the contribution of light tests, which are generally easy to fool.

Lemma 19 is a consequence of Theorem 15 and our proof of Proposition 6. We find this argument somewhat unsatisfying as it detours into coding theory and algorithms to establish a statement about distributions. In Section 4.3 we provide an alternative direct analytical proof of it.

Parallels and tangents to decoding in lattices

The closest vector problem for lattices is the analogue of the nearest codeword problem for codes. Aharonov and Regev [2] gave an efficient algorithm with preprocessing that approximates the distance to any lattice of dimension n up to a $\sqrt{n}/\log n$ approximation factor. Liu, Lyubashevsky, and Micciancio [37] gave a decoding algorithm with preprocessing for lattices under the promise that the distance between the target vector w and the lattice is within $O(\sqrt{\log n/n})$ times the length of its shortest vector.

There are strong parallels between their algorithms and ours. In the lattice setting, the analogue of our distribution D is the multivariate standard normal, H is a matrix of samples from D conditioned on membership in the dual lattice, and the proximity decision is made based on the norm of the vector Hw .

One important difference is that the algorithm of Aharonov and Regev works without any assumption on the length of the dual shortest vector, which is the analogue of dual distance for lattices. As we show in Theorem 15, a bounded dual distance assumption is necessary in the code setting.² For codes, the length of the advice must be exponential in the product of the relative proximity parameter η and the dual distance d .

Interestingly, Aharonov and Regev's $\sqrt{n}/\log n$ approximation ratio for the closest vector problem with preprocessing is matched by Goldreich and Goldwasser's statistical zero-knowledge proof [33] for the same problem. Analogously, our $n/(\log n)^2$ approximation ratio for the nearest codeword problem for balanced codes with preprocessing is matched by the statistical zero-knowledge proof of Brakerski et al. [23] for the same problem. As in the

² This still leaves open whether the balance assumption used in our algorithms can be replaced by a weaker one like bounded rate or bounded dual distance. Balance implies bounded rate. Bounded rate implies bounded dual distance. The reverse implications do not hold.

lattice setting, the two algorithms are quite different. (At a technical level, the sampling of short combinations of codewords in the code and its dual, respectively, and the Fourier analysis of this process is a commonality.)

One intriguing open question is whether Aharonov and Regev’s \sqrt{n} -approximate non-interactive (coNP) refutation for closest vector in lattices has an analogue for balanced codes.

Implications on learning noisy parities and cryptography

Learning Noisy Parities (LPN) [17] is an average-case variant of NCP. The code C is random, and the decision problem is to distinguish a randomly corrupted random codeword from a uniformly random string. The hardness of LPN has found several uses in cryptography, including public-key encryption [3], collision-resistant hashing [8, 23, 49], and more [22, 15, 1]. LPN-based schemes are attractive for their computational simplicity and plausible post-quantum security.

Some of these constructions [23, 22, 15, 49, 1] assume security of LPN at noise rate $O((\log n)^2/n)$. Our Algorithm 3 does not render them insecure owing to its inefficient preprocessing phase. However, it highlights potential concerns in settings where the code may be used as a public parameter, like a hash key as in [1] or a common random string as in [22]. Such schemes may invoke security of some LPN instances where the code C is available *long term*, and not generated ephemerally by running some algorithm in the scheme. This would allow an adversary to mount a longer duration attack to calculate the advice H and then apply Algorithm 3 to break the scheme. Our results also rule out non-uniform security: If the adversary is only bounded in size but may otherwise depend arbitrarily on C then noise of rate $O((\log n)^2/n)$ is insecure.

This non-uniform security model is in particular captured by the *linear test framework* of Boyle et al. [21, 27, 20]. They quantify the insecurity of a code C by the maximum bias

$$\max_{h \in \mathbb{F}_2^m} \mathbb{E} \left[(-1)^{\langle h, w \rangle} \right],$$

where w is a randomly corrupted random codeword of C . Couteau et al. [27, Lemma 6] observe that the insecurity becomes negligible when the dual distance of C is sufficiently large. The same logic underlies the proof of our lower bound Theorem 15.

More on decoding algorithms

Prange’s algorithm has been applied to cryptanalyze code-based schemes. There are several concrete improvements in the high-distance regime. None of them are asymptotic in the exponent. Concrete progress, some of it heuristic, has relied on two techniques: Information set decoding and statistical decoding.

Information set decoding algorithms

The expected running time of Prange’s algorithm on worst-case inputs (close to half the minimum distance) is asymptotically dominated by 2^{cn} with $c \approx 0.058$. In applications, decoding at lower distances is reduced to decoding near half the minimum distance on a code of lower dimension.

There has been considerable effort in improving this constant [38, 10, 39, 18]. Both and May [19] obtain $c \approx 0.047$. Bernstein [13] obtains a quadratic *quantum* speedup of Prange’s algorithm, and Kachigar and Tillich [35] are able to get a quantum speedup for the

23:6 Decoding Balanced Linear Codes with Preprocessing

approaches in [38, 10]. Ducas, Esser, Etinski and Kirshanova [29] provide further lower-order improvements motivated by cryptographically relevant concrete parameter settings building on the approach of [38].

Statistical decoding algorithms

The idea of using inner products with light dual codewords to estimate distance is due to Al Jabri [34]. It is the basis of statistical decoding algorithms for codes [34, 42, 25, 26] and lattices [30]. At low rates these currently outperform information set decoding algorithms.

The best analyses of these statistical decoding algorithms have so far been heuristic, relying on unproven but experimentally verified conjectures. It is unclear whether they remain valid in our regime of very low noise in very low-rate codes.

Some of these algorithms, or steps thereof, assume properties of random linear codes and random errors, even though the actual instances they operate on may not be random. This is one context in which the relation between worst-case and average-case instances of the NCP is of interest.

Worst-case to average-case reductions for balanced NCP

For balanced linear codes, the hardness of NCP reduces to that of LPN [23], albeit with extreme parameters. Yu and Zhang [48] somewhat improve and generalize this result to another restricted kind of codes they call *independent* (balanced) codes. Debris-Alazard and Resch [28] give a reduction from worst case NCP on balanced codes (with inverse polynomial noise rate) to average case NCP (with noise rate inverse polynomially close to half).

2 Concepts and Notation

2.1 Notations

Linear codes

A linear code \mathcal{C} is a collection of vectors $\{C \cdot x : x \in \mathbb{F}_2^n\}$ with C being the generator matrix in $\mathbb{F}_2^{m \times n}$. By default, we assume C has full rank and $m > n$. Consider any linear code \mathcal{C} with a generator matrix C , there exists a parity-check matrix $C^\perp \in \mathbb{F}_2^{m \times (m-n)}$, such that $(C^\perp)^T C = 0^{(m-n) \times n}$. C^\perp also generates the dual code of \mathcal{C} , denoted by $\mathcal{C}^\perp = \{h \in \mathbb{F}_2^m : \langle h, v \rangle = 0, \forall v \in \mathcal{C}\}$.

Weight and distance

We denote the Hamming weight of a vector $v \in \mathbb{F}_2^m$ by $\text{wt}(v)$. For any code $\mathcal{C} \subseteq \mathbb{F}_2^m$ and vector $w \in \mathbb{F}_2^m$, let $\text{dist}(\mathcal{C}, w) = \min_{v \in \mathcal{C}} \text{wt}(v + w)$, indicate the minimum distance of w to a code \mathcal{C} . With these notations, we are ready to define the notion of *balanced codeword* and to characterize the distance of a vector from a linear code.

► **Definition 1** (Balanced codeword). For $\beta \in [0, 1]$, a length- m vector w is called β -balanced if $\text{wt}(w) \in \frac{1}{2}(1 \pm \beta)m$. Correspondingly, a code \mathcal{C} is said to be β -balanced if every non-zero codeword is β -balanced.

► **Definition 2.** For $\eta, \beta \in [0, 1]$, w is called η -close to a code \mathcal{C} if $\text{dist}(\mathcal{C}, w) \leq \eta m$; w is called β -separated from the code \mathcal{C} , if $(w + v)$ is β -balanced for every $v \in \mathcal{C}$.

Distributions

We identify discrete distributions D with their probability mass functions, i.e., $D(x)$ is the probability that x is the outcome when sampling from D .

2.2 Nearest Codeword Problem

The search version of *nearest codeword problem* is defined as follows:

► **Definition 3** (Nearest Codeword Problem). For $m, n \in \mathbb{N}$, $0 < \eta < 1$ with $C \in \mathbb{F}_2^{m \times n}$, $w \in \mathbb{F}_2^m$, given input (C, w) with the promise that $\text{dist}(C, w) < \eta m$, the (search) nearest codeword problem NCP_η is to find s , such that $s \in \arg \min_{x \in \mathbb{F}_2^n} \text{dist}(C \cdot x, w)$.

Beyond this general definition, we introduce a variant, called *balanced nearest codeword problem*, in which the code is restricted to meet the balance property.

► **Definition 4** (Balanced Nearest Codeword Problem). For $m, n \in \mathbb{N}$, $0 < \beta, \eta < 1/6$, consider $C \in \mathbb{F}_2^{m \times n}$, $w \in \mathbb{F}_2^m$, given input (C, w) with the promise that C is a β -balanced code and $\text{dist}(C, w) < \eta m$, the (search) balanced nearest codeword problem $\text{BNCP}_{\beta, \eta}$ is to find $\arg \min_{x \in \mathbb{F}_2^n} \text{dist}(C \cdot x, w)$.

We note that when $\beta, \eta < 1/6$, with the promise that w is η -close to a β -balanced code C , there must be a unique close codeword. The *decisional balanced nearest codeword problem* is given below. For simplicity, we use a single parameter β to capture both the balance of the code and the separation of the NO instances.

► **Definition 5** (Decisional Balanced Nearest Codeword Problem). For $m, n \in \mathbb{N}$, $0 < \beta, \eta < 1/6$, consider $C \in \mathbb{F}_2^{m \times n}$, $w \in \mathbb{F}_2^m$, given input (C, w) with the promise that C is a β -balanced code, the decisional balanced nearest codeword problem $\text{DBNCP}_{\beta, \eta}$ is to decide between the following two cases:

$$\begin{aligned} \text{YES} &= \{(C, w) : w \text{ is } \eta\text{-close to the code } C\}; \\ \text{NO} &= \{(C, w) : w \text{ is } \beta\text{-separated from the code } C\}. \end{aligned}$$

2.3 Random Linear Codes

A random linear code is a code specified by a uniformly random choice of generator matrix $C \in \mathbb{F}_2^{m \times n}$. By the Gilbert-Varshamov bound, a random linear code is $3\sqrt{n/m}$ -balanced except with probability $2^{-\Omega(n)}$.

We reproduce this argument. As each nonzero codeword of a random linear code is random, by Hoeffding's inequality it is β -balanced except with probability $2 \exp(-\beta^2 m/2)$. By a union bound, the probability that there exists an unbalanced codeword is then at most

$$(2^n - 1) \cdot 2e^{-\beta^2 m/2} < 2^{n - \beta^2 m/2 + 1}, \quad (1)$$

which is $2^{-\Omega(n)}$ when $\beta = 3\sqrt{n/m}$.

3 Algorithm with Preprocessing

Our main result is Theorem 7, the search algorithm with preprocessing for the balanced codeword problem. As the algorithm for the decisional problem DBNCP is easier to describe and contains the main idea, we describe it first and prove its correctness in Proposition 6.

23:8 Decoding Balanced Linear Codes with Preprocessing

Both of these algorithms involve preprocessing. An algorithm with preprocessing for NCP is one that, in addition to the instance (C, w) , is also given an advice string that is a function of C (but not of w). This advice is not required to be efficiently computable.

In the following statements n and m are the message length and blocklength of the code, respectively.

► **Proposition 6.** *Assuming $0 < \beta, \eta < 1/6$, there is an algorithm with preprocessing for $\text{DBNCP}_{\beta, \eta}$ with both advice size and running time $m^2 \cdot \exp[O(\eta n / \log(1/\beta))]$.*

► **Theorem 7.** *Assuming $1/m \leq \beta, \eta < 1/8$, there is an algorithm with preprocessing for $\text{BNCP}_{\beta, \eta}$ with both advice size and running time $(m^4 \log^2(1/\alpha)/n^2) \cdot \exp[O(\eta n / \log(1/\alpha))]$, where $\alpha = \beta + 2\eta$.*

In particular, when $\beta = n^{-\Omega(1)}$ and $\eta = O(\log^2 n/n)$, Algorithm 3 (which proves Theorem 7) runs in polynomial time.

► **Corollary 8.** *When $\beta = n^{-\Omega(1)}$, $\eta = O(\log^2 n/n)$, there is an algorithm with preprocessing for $\text{BNCP}_{\beta, \eta}$ with advice size and running time polynomial in n .*

Since the decisional variant DBNCP reduces to BNCP , Theorem 7 is effectively more general than Proposition 6. There is a gap in complexities in the regime $\eta \gg \beta$. This gap owes to our usage of β to represent both the balance and separation parameters in DBNCP .

In Section 3.3 we present an alternative proof of Theorem 7 (with slightly worse advice size) by reduction to Proposition 6. While the resulting algorithm is less natural, this argument explains the deterioration from $\log 1/\beta$ in Proposition 6 to $\log 1/(\beta + 2\eta)$ in Theorem 7.

3.1 Decision Algorithm

In this section, we present the decision algorithm and prove Proposition 6. The algorithm is as follows, where **Pre** is the preprocessing stage, and **Decide** is the algorithm itself. The algorithms are defined by parameters that we will set later in this section, based on the values of n , m , β , and η .

■ **Algorithm 1** Preprocessing with parameters (ℓ, N) .

Input: Generator matrix $C \in \mathbb{F}_2^{m \times n}$

Output: Advice (h_1, \dots, h_N)

Sampling $\text{Samp}(C)$:

 Sample length- m unit vectors u_1, \dots, u_ℓ conditioned on $u_1 + \dots + u_\ell \in C^\perp$
 independently randomly;
 return $u_1 + \dots + u_\ell$;

Preprocessing $\text{Pre}(C)$:

for $i \in \{1, \dots, N\}$ **do**
 | Sample $h_i \leftarrow \text{Samp}(C)$;
 end
 return h_1, \dots, h_N ;

Two important distributions

Fix any n , m , and β as in the theorem statement, and any β -balanced code $C \in \mathbb{F}_2^{m \times n}$. Set the parameter $\ell \in \mathbb{N}$ to some fixed value (to be determined later). We define two distributions sampled as follows over the codespace, which will be crucial to the working of our algorithm:

- Distribution D : sample length- m unit vectors u_1, \dots, u_ℓ uniformly randomly, then output $u_1 + \dots + u_\ell$.
- Distribution D_C : sample $h \leftarrow D$ conditioned on $h \in \mathcal{C}^\perp$.

We will set ℓ to be an even integer (among other restrictions), so that the distribution D_C is well-defined. This is guaranteed by the fact that the zero vector 0^m lies in the dual space of any code and occurs with non-zero probability under the distribution D (so D_C has nonempty support).

In the *preprocessing* stage **Pre**, the algorithm samples vectors (h_1, \dots, h_N) from distribution D_C , for $N \in \mathbb{N}$, which we refer to as advice. We note that since the preprocessing stage is allowed to be inefficient, it suffices to show the existence of good advice applicable to all inputs w (which implies that it can be found inefficiently). We present a uniform method for obtaining such advice, that succeeds with high probability, which more than meets this bar.

The *decision algorithm* **Decide** performs as follows: it takes the vector w and the advice (h_1, \dots, h_N) associated with the code \mathcal{C} as input, and compares the value $\sum_i (-1)^{\langle h_i, w \rangle}$ with a hard-coded threshold t to decide whether the vector is close to or separated from the code \mathcal{C} .

■ **Algorithm 2** Decision algorithm with parameters (N, t) .

Input: Vector $w \in \mathbb{F}_2^m$ and advice (h_1, \dots, h_N)

Output: YES or NO, indicating the vector w is closed to or separated from the code \mathcal{C}

Algorithm **Decide** $(w; h_1, \dots, h_N)$:

```

| if  $\sum_i (-1)^{\langle h_i, w \rangle} > t \cdot N$  then
|   | return YES;
| end
| return NO;

```

To prove correctness we analyze the Fourier transform $\hat{D}_C(w) = \mathbb{E} [(-1)^{\langle h, w \rangle}]$ and show that it is a good measure of distance to the code (Lemma 9).

Fourier coefficient of D_C

We describe the probability mass function of D_C in terms of that of D , which is

$$D_C(h) = \frac{\mathbf{1}[h \in \mathcal{C}^\perp]}{Z_C} \cdot D(h),$$

where Z_C is a normalization factor that satisfies

$$Z_C = \sum_{h \in \mathcal{C}^\perp} D(h). \tag{2}$$

The Fourier coefficient of D is defined by

$$\hat{D}(w) = \mathbb{E}_{h \leftarrow D} [(-1)^{\langle h, w \rangle}] = \sum_{h \in \mathbb{F}_2^m} D(h) \cdot (-1)^{\langle h, w \rangle}.$$

23:10 Decoding Balanced Linear Codes with Preprocessing

For the distribution D_C , its Fourier coefficient \hat{D}_C can be represented by \hat{D} ,

$$\begin{aligned}
\hat{D}_C(w) &= \mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right] \\
&= \sum_{h \in \mathbb{F}_2^m} D_C(h) \cdot (-1)^{\langle h, w \rangle} \\
&= \frac{1}{Z_C} \sum_{h \in \mathbb{F}_2^m} \mathbf{1}[h \in \mathcal{C}^\perp] D(h) \cdot (-1)^{\langle h, w \rangle} \\
&= \frac{1}{Z_C \cdot 2^n} \sum_{v \in \mathcal{C}} \sum_{h \in \mathbb{F}_2^m} D(h) \cdot (-1)^{\langle h, w+v \rangle} \\
&= \frac{1}{Z_C \cdot 2^n} \sum_{v \in \mathcal{C}} \hat{D}(w+v), \tag{3}
\end{aligned}$$

where the fourth equality follows the fact that

$$\mathbf{1}[h \in \mathcal{C}^\perp] = \frac{1}{2^n} \sum_{v \in \mathcal{C}} (-1)^{\langle h, v \rangle},$$

since if $h \in \mathcal{C}^\perp$, for all $v \in \mathcal{C}$, $(-1)^{\langle h, v \rangle} = 1$; if $h \notin \mathcal{C}^\perp$, there exists $v' \in \mathcal{C}$, such that $\langle h, v' \rangle = 1$, then

$$\sum_{v \in \mathcal{C}} (-1)^{\langle h, v \rangle} = \frac{1}{2} \sum_{v \in \mathcal{C}} \left((-1)^{\langle h, v \rangle} + (-1)^{\langle h, v+v' \rangle} \right) = 0.$$

Since $\hat{D}_C(0^m) = 1$, the normalization factor is

$$Z_C = \frac{1}{2^n} \sum_{v \in \mathcal{C}} \hat{D}(v). \tag{4}$$

Combining (3) and (4), we obtain that

$$\hat{D}_C(w) = \mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right] = \frac{\sum_{v \in \mathcal{C}} \hat{D}(v+w)}{\sum_{v \in \mathcal{C}} \hat{D}(v)}. \tag{5}$$

Fourier coefficient of D

Recall that the distribution D outputs the XOR of ℓ independent random unit vectors over \mathbb{F}_2^m ,

$$\begin{aligned}
\hat{D}(w) &= \mathbb{E}_{h \leftarrow D} \left[(-1)^{\langle h, w \rangle} \right] = \mathbb{E}_{u_1, \dots, u_\ell} \left[\prod_i (-1)^{\langle u_i, w \rangle} \right] \\
&= \prod_i \mathbb{E}_{u_i} \left[(-1)^{\langle u_i, w \rangle} \right] = \left(1 - 2 \cdot \frac{\text{wt}(w)}{m} \right)^\ell. \tag{6}
\end{aligned}$$

The last equality is obtained by

$$\mathbb{E}_{u_i} \left[(-1)^{\langle u_i, w \rangle} \right] = \frac{m - \text{wt}(w)}{m} \cdot 1 + \frac{\text{wt}(w)}{m} \cdot (-1) = 1 - 2 \cdot \frac{\text{wt}(w)}{m}.$$

We now turn to establishing guarantees for our algorithm. We start by showing that the function $\mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right]$ serves as an “ideal” distinguishing function between inputs η -close to the code, and those β -separated from the code (hence helping us decide instances of the DBNCP problem). Note that this function is not necessarily efficient to compute, but we will later show that this is exactly what our preprocessing step will help us handle.

► **Lemma 9.** For $m, n \in \mathbb{N}$, $0 < \beta, \eta < 1/6$. Given any $\text{DBNCP}_{m,\beta,\eta}$ instance (C, w) , letting $\ell = 2 \cdot \lceil n/\log(1/\beta) \rceil$ be an even integer, we have that:

If w is η -close to the code \mathcal{C} ,

$$\mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right] > \frac{1}{2} (1 - 2\eta)^{2(n/\log(1/\beta)+1)};$$

If w is β -separated from the code \mathcal{C} ,

$$\mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right] < 2^{-n}.$$

Proof of Lemma 9. Consider any $\text{DBNCP}_{m,\beta,\eta}$ instance (C, w) , where C is the generator matrix of a β -balanced code \mathcal{C} , which means that for $v = C \cdot x$ with all non-zero $x \in \mathbb{F}_2^n$, the Hamming weight $\text{wt}(v) \in \frac{1}{2}(1 \pm \beta)m$, thus $0 \leq \hat{D}(v) \leq \beta^\ell$ according to (6). As $\hat{D}(0^m) = 1$, we have

$$1 \leq \sum_{v \in \mathcal{C}} \hat{D}(v) \leq 1 + (2^n - 1) \cdot \beta^\ell. \quad (7)$$

When $\ell = 2 \cdot \lceil n/\log(1/\beta) \rceil$,

$$\sum_{v \in \mathcal{C}} \hat{D}(v) \leq 1 + (2^n - 1) \cdot \beta^\ell < 1 + 2^n \cdot \beta^{2n/\log(1/\beta)} = 1 + 2^{-n}. \quad (8)$$

Suppose that (C, w) is a YES instance, which implies w is η -close to the code \mathcal{C} . Without loss of generality, assume $w = C \cdot x + e$ with $\text{wt}(e) \leq \eta m$,

$$\begin{aligned} \mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right] &= \frac{\sum_{v \in \mathcal{C}} \hat{D}(v + w)}{\sum_{v \in \mathcal{C}} \hat{D}(v)} > \frac{\hat{D}(e)}{1 + 2^{-n}} \\ &\geq \frac{(1 - 2\eta)^\ell}{1 + 2^{-n}} > \frac{1}{2} (1 - 2\eta)^{2(n/\log(1/\beta)+1)}, \end{aligned}$$

where the first inequality follows (8) and the positivity of the Fourier coefficients; the last inequality holds for all sufficiently large n . When $\beta, \eta < 1/6$, $(1 - 2\eta)^{2n/\log(1/\beta)} \gg 2^{-n}$.

For w being a NO instance, for all $v \in \mathcal{C}$, $\text{wt}(v + w) \in \frac{1}{2}(1 \pm \beta)m$, thus $\hat{D}(v + w) \leq \beta^\ell$,

$$\mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right] = \frac{\sum_{v \in \mathcal{C}} \hat{D}(v + w)}{\sum_{v \in \mathcal{C}} \hat{D}(v)} \leq \sum_{v \in \mathcal{C}} \hat{D}(v + w) \leq 2^n \cdot \beta^\ell < 2^{-n}.$$

The first inequality follows the lower bound given by (7). ◀

Quality of advice

In the following, we show that with high probability the preprocessing stage outputs *good* advice $(h_1, \dots, h_N) \in (\mathbb{F}_2^m)^N$: i.e., advice on which the algorithm `Decide` yields the correct output on all inputs w that satisfy the promise (for some settings of parameters (ℓ, N, t) depending on n, m, β, η). Informally, this will be advice for which inputs that are close to the code lead to an evaluation by `Decide` to a relatively high value, while inputs that are far from the code in turn lead to a significantly lower evaluation.

Recall that Lemma 9 shows that the expected value $\mathbb{E}_h \left[(-1)^{\langle h, w \rangle} \right]$, for h being sampled from D_C , serves as an ideal test in the sense that it is large if the vector is close to the code \mathcal{C} ; and is small when the vector is far. Using the advice, Algorithm 2 then essentially estimates

23:12 Decoding Balanced Linear Codes with Preprocessing

this expectation using the N samples in the advice, namely (h_1, \dots, h_N) . Since h_i s are independent samples, standard concentration arguments ensure that, with high probability, the advice works for all possible w (via an union bound) under the promise related to C . This yields the following lemma.

► **Lemma 10.** *Consider $m, n \in \mathbb{N}$ and $0 < \beta, \eta < 1/6$. Given any $\text{DBNCP}_{\beta, \eta}$ instance (C, w) , there is a setting of parameters (ℓ, N, t) , for which, with overwhelming probability, the preprocessing algorithm Pre outputs advice (h_1, \dots, h_N) such that for all w that are either η -close to or β -separated from the code C , the decision returned by Algorithm 2 is correct, where*

$$N = 72m \cdot (1 - 2\eta)^{-4(n/\log(1/\beta)+1)}.$$

Proof of Lemma 10. Let $\ell = 2 \cdot \lceil n/\log(1/\beta) \rceil$, by Lemma 9, for any $\text{DBNCP}_{m, \beta, \eta}$ instance (C, w) , we have either

$$\mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right] > \frac{1}{2} (1 - 2\eta)^{2(n/\log(1/\beta)+1)} \text{ or } \mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right] < 2^{-n},$$

for w being close to or separated from C , respectively. As $\eta, \beta < 1/6$, when n is sufficiently large, $(1 - 2\eta)^{2n/\log(1/\beta)} \gg 2^{-n}$. Set the threshold to be

$$t = \frac{1}{3} (1 - 2\eta)^{2(n/\log(1/\beta)+1)}$$

and let $\delta = \frac{1}{6} (1 - 2\eta)^{2(n/\log(1/\beta)+1)}$.

For a fixed w being η -close to the code C , the probability that Algorithm 2 decides w correctly with advice (h_1, \dots, h_N) on the randomness of preprocessing is equivalent to

$$\begin{aligned} \Pr_{h_i} \left[\sum_i (-1)^{\langle h_i, w \rangle} > t \cdot N \right] &> \Pr_{h_i} \left[\sum_i (-1)^{\langle h_i, w \rangle} > \left(\mathbb{E}_{h_i} (-1)^{\langle h_i, w \rangle} - \delta \right) \cdot N \right] \\ &\geq 1 - e^{-\delta^2 N/2} \end{aligned}$$

where the first inequality is obtained from $t < \mathbb{E}_{h_i} (-1)^{\langle h_i, w \rangle} - \delta$ and the second one follows from Hoeffding's inequality; the probability is at least $(1 - e^{-m})$ when $N = 2m \cdot \delta^{-2} = 72m \cdot (1 - 2\eta)^{-4(n/\log(1/\beta)+1)}$.

Similarly, for w being far from the code C , the success probability can be lower-bounded by

$$\begin{aligned} \Pr_{h_i} \left[\sum_i (-1)^{\langle h_i, w \rangle} \leq t \cdot N \right] &> \Pr_{h_i} \left[\sum_i (-1)^{\langle h_i, w \rangle} \leq \left(\mathbb{E}_{h_i} (-1)^{\langle h_i, w \rangle} + \delta \right) \cdot N \right] \\ &\geq 1 - e^{-\delta^2 N/2} \end{aligned}$$

where the first inequality follows from $t > \mathbb{E}_{h_i} (-1)^{\langle h_i, w \rangle} + \delta$ and the second inequality is ensured by Hoeffding's inequality.

Therefore, by the union bound, the probability that Algorithm 2 outputs correctly for all possible w is at least

$$\Pr_{h_i} [\forall w, \text{Decide}(w; h_1, \dots, h_N) \text{ is correct}] \geq 1 - 2^m \cdot e^{-m} > 1 - 2^{-0.4m},$$

when h_i is generated by the preprocessing algorithm. ◀

Proof of Proposition 6. From Lemma 10, we conclude that, with overwhelming probability, $\text{Pre}(C)$ returns good advice (h_1, \dots, h_N) , which works for all w that satisfies the promise, when $N = 72m \cdot (1 - 2\eta)^{-4(n/\log(1/\beta)+1)}$. This proves that the algorithm with preprocessing works correctly. The advice size and running time of the algorithm is: $O(N \cdot m) = m^2 \cdot \exp O(\eta m / \log(1/\beta))$. ◀

3.2 Search Algorithm

In this section, we describe the search algorithm and prove Theorem 7. With the same preprocessing procedure applied (as in Algorithm 1), the search algorithm proceeds as follows:

■ **Algorithm 3** Search algorithm with parameter N .

Input: Vector $w \in \mathbb{F}_2^m$ and advice (h_1, \dots, h_N)
Output: $\hat{x} \in \mathbb{F}_2^n$, such that $C \cdot \hat{x}$ is the nearest codeword to w
Algorithm Search $(w; h_1, \dots, h_N)$:

```

  Set  $S = \sum_k (-1)^{\langle h_k, w \rangle}$ ;
  for  $i \in \{1, \dots, m\}$  do
    Set  $w^{(i)} \leftarrow w$  with  $i$ -th bit flipped;
    if  $\sum_k (-1)^{\langle h_k, w^{(i)} \rangle} < S$  then
      Set  $\hat{e}_i \leftarrow 0$ ;
    end
  else
    Set  $\hat{e}_i \leftarrow 1$ ;
  end
end
Set  $\hat{e} = (\hat{e}_1, \dots, \hat{e}_m)$ ;
Solve the linear system  $(C \cdot \hat{x} + \hat{e} = w)$  for  $\hat{x}$ ;
return  $\hat{x}$ 

```

Suppose the input is $w = C \cdot x + e$ with advice (h_1, \dots, h_N) ; the algorithm proceeds to recover each bit of e by testing how the value $\sum_k (-1)^{\langle h_k, w \rangle}$ changes under corresponding bit flip. Intuitively, under the distribution D_C used to show Lemma 9, the value $\mathbb{E}_h [(-1)^{\langle h, w \rangle}]$ decreases monotonically when w is farther from \mathcal{C} . We can then toggle each coordinate of w to detect coordinates i for which $e_i = 1$. We show this formally below.

► **Lemma 11.** For $m, n \in \mathbb{N}$, $1/m \leq \beta, \eta < 1/8$, consider any $\text{BNCP}_{\beta, \eta}$ instance (C, w) with the promise that w is η -close to \mathcal{C} , and let $\ell = 2 \cdot \lceil n / \log(1/(\beta + 2\eta)) \rceil$ be an even integer. Assume that $w = C \cdot x + e$ with $\text{wt}(e) \leq \eta m$, denote $w^{(i)}, e^{(i)}$ as the vector w, e with their i -th bits flipped (respectively). Then, the sign of Δ_i indicates the value of the i -th bit e_i , where

$$\Delta_i = \mathbb{E}_{h \leftarrow D_C} [(-1)^{\langle h, w \rangle}] - \mathbb{E}_{h \leftarrow D_C} [(-1)^{\langle h, w^{(i)} \rangle}].$$

In particular, if $e_i = 0$, $\Delta_i > \delta$; if $e_i = 1$, $\Delta_i < -\delta$; with

$$\delta > \frac{\ell}{m} (1 - 4\eta)^\ell.$$

Proof of Lemma 11. Assume $\text{wt}(e) = \eta' m$ for $\eta' \leq \eta$. If the i -th bit of e is 0, the difference between $\hat{D}(e)$ and $\hat{D}(e^{(i)})$ is

$$\begin{aligned} \hat{D}(e) - \hat{D}(e^{(i)}) &= (1 - 2\eta')^\ell - \left(1 - 2\left(\eta' + \frac{1}{m}\right)\right)^\ell \\ &= \left(1 - 2\left(\eta' + \frac{1}{m}\right) + \frac{2}{m}\right)^\ell - \left(1 - 2\left(\eta' + \frac{1}{m}\right)\right)^\ell \\ &\geq \frac{2\ell}{m} \cdot \left(1 - 2\left(\eta' + \frac{1}{m}\right)\right)^{\ell-1} \\ &\geq \frac{2\ell}{m} (1 - 4\eta)^\ell. \end{aligned}$$

23:14 Decoding Balanced Linear Codes with Preprocessing

By (5), we have

$$\begin{aligned}\Delta_i &= \frac{\sum_{v \in \mathcal{C}} \hat{D}(v+w) - \sum_{v \in \mathcal{C}} \hat{D}(v+w^{(i)})}{\sum_{v \in \mathcal{C}} \hat{D}(v)} \\ &\geq \frac{\hat{D}(e) - \hat{D}(e^{(i)}) - (2^n - 1) \cdot (\beta + 2\eta)^\ell}{1 + (2^n - 1) \cdot \beta^\ell}.\end{aligned}$$

Let $\alpha = (\beta + 2\eta)$, $\ell = 2 \cdot \lceil n/\log(1/\alpha) \rceil$, then for sufficiently large n we have

$$2^n \cdot (\beta + 2\eta)^\ell < 2^{-n}.$$

When $\ell \approx 2n/\log(1/\alpha)$ and $\beta, \eta < 1/8$, $2^{-n} \ll (2\ell/m) \cdot (1 - 4\eta)^\ell$ for all sufficiently large n . Thus,

$$\Delta_i > \frac{1}{2} \left(\hat{D}(e) - \hat{D}(e^{(i)}) \right) \geq \frac{\ell}{m} (1 - 4\eta)^\ell,$$

which proves the claim for this case.

If the i -th entry of e is 1,

$$\begin{aligned}\hat{D}(e^{(i)}) - \hat{D}(e) &= \left(1 - 2 \left(\eta' - \frac{1}{m} \right) \right)^\ell - (1 - 2\eta')^\ell \\ &= \left(1 - 2\eta' + \frac{2}{m} \right)^\ell - (1 - 2\eta')^\ell \\ &> \frac{2\ell}{m} (1 - 2\eta')^\ell.\end{aligned}$$

Using similar arguments as above, we can obtain

$$\begin{aligned}-\Delta_i &= \frac{\sum_{v \in \mathcal{C}} \hat{D}(v+w^{(i)}) - \sum_{v \in \mathcal{C}} \hat{D}(v+w)}{\sum_{v \in \mathcal{C}} \hat{D}(v)} \\ &\geq \frac{\hat{D}(e^{(i)}) - \hat{D}(e) - (2^n - 1) \cdot (\beta + 2\eta)^\ell}{1 + (2^n - 1) \cdot \beta^\ell} \\ &> \frac{\ell}{m} (1 - 2\eta)^\ell,\end{aligned}$$

which proves the claim for this case as well. \blacktriangleleft

We will also require the following concentration claim in our main proof. This follows from standard inequalities.

► **Lemma 12.** *For any $w_1, w_2 \in \mathbb{F}_2^m$, if there exists $\delta > 0$, such that*

$$\mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w_1 \rangle} \right] - \mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w_2 \rangle} \right] > \delta,$$

then for $N \geq 8m \cdot \delta^{-2}$, we have that

$$\Pr_{h_1, \dots, h_N \leftarrow D_C} \left[\sum_k (-1)^{\langle h_k, w_1 \rangle} > \sum_k (-1)^{\langle h_k, w_2 \rangle} \right] > 1 - 2 \exp(-m).$$

Proof of Lemma 12. By Hoeffding's inequality,

$$\Pr_{h_1, \dots, h_N} \left[\sum_k (-1)^{\langle h_k, w_1 \rangle} > \left(\mathbb{E}_h (-1)^{\langle h, w_1 \rangle} - \frac{\delta}{2} \right) \cdot N \right] > 1 - \exp(-\delta^2 N/8) = 1 - \exp(-m),$$

where $N = 8m \cdot \delta^{-2}$. Similarly,

$$\Pr_{h_1, \dots, h_N} \left[\sum_k (-1)^{\langle h_k, w_2 \rangle} < \left(\mathbb{E}_h (-1)^{\langle h, w_2 \rangle} + \frac{\delta}{2} \right) \cdot N \right] > 1 - \exp(-\delta^2 N/8) = 1 - \exp(-m).$$

Since $\mathbb{E}_h (-1)^{\langle h, w_1 \rangle} - \frac{\delta}{2} > \mathbb{E}_h (-1)^{\langle h, w_2 \rangle} + \frac{\delta}{2}$, we have

$$\Pr_{h_1, \dots, h_N} \left[\sum_k (-1)^{\langle h_k, w_1 \rangle} > \sum_k (-1)^{\langle h_k, w_2 \rangle} \right] > 1 - 2 \cdot \exp(-m). \quad \blacktriangleleft$$

Finally we can show that the preprocessing step again generates good advice such that the search algorithm produces the correct output. This lemma plays a similar role to Lemma 10 for the decisional setting.

► **Lemma 13.** *For $m, n \in \mathbb{N}$, $0 < \beta, \eta < 1/8$. Given any $\text{BNCP}_{\beta, \eta}$ instance (C, w) , there is a setting of parameters (ℓ, N) , for which, with high probability, the preprocessing algorithm outputs advice (h_1, \dots, h_N) , such that, for all w that is η -close to \mathcal{C} , Algorithm 3 outputs x such that $C \cdot x$ is the closest codeword to w , where*

$$N = (2m^3 \log^2(1/\alpha)/n^2) \cdot (1 - 4\eta)^{-4(n/\log(1/\alpha)+1)}, \text{ with } \alpha = \beta + 2\eta.$$

Proof of Lemma 13. Consider any fixed $w = C \cdot x + e$ with $\text{wt}(e) \leq \eta m$, denote $w^{(i)}, e^{(i)}$ as the vector w, e with their i -th bit flipped. Let $\ell = 2 \cdot \lceil n/\log(1/\alpha) \rceil$ and $\alpha = \beta + 2\eta$. By Lemma 11, when $e_i = 0$

$$\mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right] - \mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w^{(i)} \rangle} \right] > \frac{\ell}{m} (1 - 4\eta)^\ell.$$

Let $\delta = \frac{\ell}{m} (1 - 4\eta)^\ell$ and $N = 8m \cdot \delta^{-2} = (2m^3 \log^2(1/\alpha)/n^2) \cdot (1 - 4\eta)^{-4(n/\log(1/\alpha)+1)}$, by Lemma 12, the probability that Search in Algorithm 3 sets \hat{e}_i to be 0 for h_k 's being sampled from D_C can be lower-bounded by

$$\Pr_{h_1, \dots, h_N} [\hat{e}_i = 0 | e_i = 0] = \Pr_{h_1, \dots, h_N} \left[\sum_k (-1)^{\langle h_k, w \rangle} > \sum_k (-1)^{\langle h_k, w^{(i)} \rangle} \right] > 1 - 2 \cdot \exp(-m).$$

When $e_i = 1$, according to Lemma 11,

$$\mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w^{(i)} \rangle} \right] - \mathbb{E}_{h \leftarrow D_C} \left[(-1)^{\langle h, w \rangle} \right] > \frac{\ell}{m} (1 - 4\eta)^\ell,$$

then with the same parameters, Lemma 12 implies that

$$\Pr_{h_1, \dots, h_N} [\hat{e}_i = 1 | e_i = 1] = \Pr_{h_1, \dots, h_N} \left[\sum_k (-1)^{\langle h_k, w \rangle} \leq \sum_k (-1)^{\langle h_k, w^{(i)} \rangle} \right] > 1 - 2 \cdot \exp(-m).$$

For a fixed w , the probability that Algorithm 3 finds the correct x is at least

$$\Pr_{h_1, \dots, h_N} \left[\hat{x} = \arg \min_x \text{wt}(C \cdot x + w) \right] = \Pr_{h_1, \dots, h_N} [\forall i, \hat{e}_i = e_i] > 1 - 2m \cdot \exp(-m).$$

Consider all possible w that satisfies the promise, by taking the union bound,

$$\Pr_{h_1, \dots, h_N} \left[\forall w, \hat{x} = \arg \min_x \text{wt}(C \cdot x + w) \right] > 1 - 2m \cdot 2^m \cdot \exp(-m) > 1 - m \cdot 2^{-0.4m+1}. \quad \blacktriangleleft$$

23:16 Decoding Balanced Linear Codes with Preprocessing

Proof of Theorem 7. Lemma 13 indicates that $\text{Pre}(C)$ outputs a good advice matrix H with overwhelming probability, given a balanced code \mathcal{C} , such that Algorithm 3 works correctly for all possible input w , when $N = (m^3 \log^2(1/\alpha)/n^2) \cdot \exp O(\eta n / \log(1/\alpha))$ with $\alpha = \beta + 2\eta$, then the advice size and running time are $O(m \cdot N) = (m^4 \log^2(1/\alpha)/n^2) \cdot \exp O(\eta n / \log(1/\alpha))$. ◀

3.3 Search to Decision Reduction

► **Proposition 14.** *If there exists an algorithm for $\text{DBNCP}_{\beta+2\eta,\eta}$ with message length $n - 1$ and block length m with advice size a and time complexity t , then there exists an algorithm for $\text{BNCP}_{\beta,\eta}$ with advice size a and time complexity tn .*

As a consequence of Proposition 14, the search problem $\text{BNCP}_{\beta,\eta}$ can be solved using the decision algorithm in time $m^2 n \exp O(\eta n / \log(1/\alpha))$, where $\alpha = \beta + 2\eta$.

Proof of Proposition 14. Given a $\text{BNCP}_{\beta,\eta}$ instance (C, w) , where $C \in \mathbb{F}_2^{m \times n}$ and $w \in \mathbb{F}_2^m$ the algorithm: (1) constructs n codes $C^{(i)} \in \mathbb{F}_2^{m \times (n-1)}$ by removing the i -th column of C , (2) queries the $\text{DBNCP}_{\beta+2\eta,\eta}$ oracle on $(C^{(i)}, w)$ for all i , and (3) outputs \hat{x} where

$$\hat{x}_i = \begin{cases} 0, & \text{if } \text{DBNCP}_{\beta+2\eta,\eta}(C^{(i)}, w) \text{ answers YES,} \\ 1, & \text{if not.} \end{cases}$$

All codes $C^{(i)}$ are subcodes of C and therefore β -balanced.

Assuming w is η -close to C , it equals $C \cdot x + e$ for some e of weight less than ηm . We show that \hat{x} must equal x . If $x_i = 0$ then Cx equals $C^{(i)}x^{(i)}$, where $x^{(i)}$ is x with its i -th entry removed. Therefore $w = C^{(i)}x^{(i)} + e$ is also η -close to $C^{(i)}$ and $\text{DBNCP}_{\beta+2\eta,\eta}$ answers YES.

If $x_i = 1$, then w equals $C^{(i)}x^{(i)} + c^{(i)} + e$, where $c^{(i)}$ is the i -th column of C . It remains to argue that $c^{(i)} + e$ is $(\beta + 2\eta)$ -separated from $C^{(i)}$. This guarantees a NO answer from the oracle. If, for contradiction, $c^{(i)} + e$ was $\frac{1}{2}(1 - \beta - 2\eta)$ -close or $\frac{1}{2}(1 + \beta + 2\eta)$ -far from some codeword c in $C^{(i)}$, then $c^{(i)}$ would be $\frac{1}{2}(1 - \beta)$ -close to or $\frac{1}{2}(1 + \beta)$ -far from c . $c + c^{(i)}$ would then be a β -unbalanced codeword of C , violating the promise. ◀

4 On the Optimality of Our Algorithms

In this section, we show the optimality of the algorithm presented in the previous section. We emphasize that all our impossibility results in this section are stated for the *decision problem* $\text{DBNCP}_{\beta,\eta}$ with parameters set to $m = \text{poly}(n)$ and $\beta = 3\sqrt{n/m}$. This value of β allows an overwhelming fraction of linear codes to meet the balance requirement (see (1)).

4.1 Limits of Threshold Distinguishers

Recall that Algorithm 2 proceeds as follows: on input $w \in \mathbb{F}_2^m$, it takes as advice a matrix $H = (h_1, \dots, h_N) \in \mathbb{F}_2^{N \times m}$ associated with the code \mathcal{C} , it calculates $\sum_i (-1)^{\langle h_i, w \rangle}$ and compares it with a threshold T . We call this type of algorithm a *threshold distinguisher* for code \mathcal{C} and refer to N as its size.

More generally, we allow a threshold distinguisher to apply some affine shift $b \in \mathbb{F}_2^N$. That is, it computes $\sum_i (-1)^{\langle h_i, w \rangle + b_i}$ where b_i denotes the i -th coordinate of b . In Section 4.2 we show that threshold distinguishers are somewhat more powerful than they seem.

Our main lower bound shows that threshold distinguishers have limited power.

► **Theorem 15.** For $n, m, d \in \mathbb{N}$, $\beta = 3\sqrt{n/m}$, for any code \mathcal{C} of message length n , blocklength m , and dual distance d , no threshold distinguisher of size $\frac{1}{7} \exp[\min\{\eta d/6, 3n\}]$ correctly decides whether w is η -close to or β -separated from \mathcal{C} for all w .

All but $2^{-n/2}$ linear codes have dual distance at most $d = n/(2 \log(m))$. The reason is that there are at most m^d non-zero strings of weight at most d , and each of them is a dual codeword with probability 2^{-n} . By a union bound all of them fail to be in the dual except with probability $2^{-n/2}$. This gives the following corollary:

► **Corollary 16.** For all but a $2^{-n/2}$ fraction of linear codes \mathcal{C} , no threshold distinguisher of size $\frac{1}{7} \exp \min\{(\eta n/12 \log m), 3n\}$ decides whether w is η -close or $3\sqrt{n/m}$ -separated from \mathcal{C} for all w .

In particular, if m is polynomial in n and $\eta \gg (\log n)^2/n$, polynomial-size threshold distinguishers fail to solve $\text{DBNCP}_{\eta, \beta}$ with $\beta = 3\sqrt{n/m}$.

Proof of Theorem 15. By flipping all b_i if necessary, we may assume that the distinguisher accepts when $A_{H,b}(w) \geq T$, and rejects otherwise, where $A_{H,b}(w) = \sum_i (-1)^{\langle h_i, w \rangle + b_i}$.

We argue that $A_{H,b}$ cannot “tell apart” a randomly corrupted random codeword from a truly random string. Let $w = C \cdot x + e$ with random x and the bits of e independent $\text{Ber}(\eta/2)$. We first show that:

$$\mathbb{E}A_{H,b}(w) \leq \mathbb{E}A_{H,b}(r) + N \exp(-\eta d),$$

where r is a truly random string. By linearity of expectation, it is sufficient to argue that $\mathbb{E}(-1)^{\langle h, w \rangle + b}$ is $(1 - \eta)^d$ -close to $\mathbb{E}(-1)^{\langle h, r \rangle + b}$ for every row (h, b) . If h is the all zero string both expectations are one. If h not a dual codeword both are zero. The only difference comes from non-zero dual codewords. The difference is then $(1 - \eta)^{\text{wt}(h)} \leq (1 - \eta)^d \leq \exp(-\eta d)$.

Now assume $A_{H,b}$ outputs at least T for all strings η -close to the code and at most $T - 1$ for all β -separated strings. By the law of total probability,

$$\mathbb{E}A_{H,b}(w) \geq T(1 - \Pr(\text{FAR})) + \mathbb{E}[A_{H,b}(w) \mid \text{FAR}] \Pr(\text{FAR}) \geq T - 2N \Pr(\text{FAR}),$$

where FAR is the event that w is η -far from \mathcal{C} . By a Chernoff bound, this has probability at most $\exp(-\eta m/6)$. On the other hand,

$$\mathbb{E}A_{H,b}(r) \leq (T - 1)(1 - \Pr(\overline{\text{SEP}})) + \mathbb{E}[A_{H,b}(w) \mid \overline{\text{SEP}}] \Pr(\overline{\text{SEP}}) \leq (T - 1) + 2N \Pr(\overline{\text{SEP}}),$$

where $\overline{\text{SEP}}$ is the event that r is not β -separated from \mathcal{C} . By a union bound and a Chernoff bound, $\overline{\text{SEP}}$ had probability at most $2^n \cdot 2 \exp(-\beta^2 m/2) \leq 2 \exp(-3n)$. In summary,

$$\begin{aligned} T - 2N \exp(-\eta m/6) &\leq \mathbb{E}A_{H,b}(w) \\ &\leq \mathbb{E}A_{H,b}(r) + N \exp(-\eta d) \\ &\leq (T - 1) + N \exp(-\eta d) + 2N \cdot 2 \exp(-3n). \end{aligned}$$

It follows that

$$\frac{1}{N} \leq \exp(-\eta d) + 2 \exp(-\eta m/6) + 4 \exp(-3n) \leq 7 \max\{\exp(-\eta d), \exp(-\eta m/6), \exp(-3n)\}.$$

As $d \leq m$, $\exp(-\eta d)$ and $\exp(-\eta m/6)$ are both bounded by $\exp(-\eta d/6)$. ◀

4.2 Limits of Interval Distinguishers

Instead of a threshold, a distinguisher could potentially base its decision on some other function of the measure $A_{H,b}(w)$. For example, it may be sensible to accept inputs whose value is close to zero and reject those whose value is far from zero, be positive or negative. This is a special case of an interval distinguisher. In general, a k -interval distinguisher partitions the range of $A_{H,b}$ into k intervals and decides based on the identity of the interval that $A_{H,b}(w)$ belongs to.

► **Corollary 17.** For $n, m \in \mathbb{N}$, $\beta = 3\sqrt{n/m}$, for any code \mathcal{C} of message length n , blocklength m , and dual distance d , no k -interval distinguisher of size $\frac{1}{14} \exp[\min\{\eta d/6, 3n\}/(k-1)]$ correctly decides whether w is η -close to or β -separated from \mathcal{C} for all w .

We prove it by reducing an interval distinguisher to a threshold distinguisher.

► **Lemma 18.** Every k -interval distinguisher of size N can be simulated by a threshold distinguisher of size at most $(2N)^{k-1}$.

Proof. For every partition of the line into k intervals there is a polynomial p of degree $k-1$ that alternates sign among the intervals. Given a k -interval distinguisher $A_{H,b}$, we construct a threshold distinguisher $A_{\hat{H},\hat{b}}$ so that

$$A_{\hat{H},\hat{b}}(w) = p(A_{H,b}(w)). \quad (9)$$

The value $A_{\hat{H},\hat{b}}(w)$ is positive precisely when $A_{H,b}(w)$ falls into a positive interval. The polynomial p factorizes as

$$p(z) = \prod_{i=1}^{k-1} (z - a_i),$$

where the breakpoints a_i are integers between $-N$ and N .

The advice \hat{H}, \hat{b} will be constructed by “applying” p to H, b . To do so we describe how to subtract a constant from a threshold distinguisher, and how to multiply two threshold distinguishers.

To subtract a constant a from H, b we pad H with $|a|$ zero rows and pad b with $|a|$ ones if $a > 0$ and $|a|$ zeros if $a < 0$. The resulting distinguisher \hat{H}, \hat{b} satisfies

$$A_{\hat{H},\hat{b}}(w) = A_{H,b}(w) - a.$$

To multiply two distinguishers H, b and H', b' , we create a new distinguisher \hat{H}, \hat{b} whose rows are the XORS of all pairs of rows of (H, b) and (H', b') . Then

$$\begin{aligned} A_{\hat{H},\hat{b}}(w) &= \sum_{h,h'} (-1)^{\langle h+h', w \rangle + b+b'} \\ &= \left(\sum_h (-1)^{\langle h, w \rangle + b} \right) \left(\sum_{h'} (-1)^{\langle h', w \rangle + b'} \right) \\ &= A_{H,b}(w) \cdot A_{H',b'}(w). \end{aligned}$$

Applying N subtractions and $k-2$ multiplications we obtain a distinguisher satisfying (9). Subtraction affects size by at most an additive N . As the original distinguisher has size N it at most doubles it. Multiplication results in a distinguisher whose size is the product of its parts. The size of the final distinguisher is therefore at most $(2N)^{k-1}$. ◀

4.3 Optimality of Our Distribution

The specific measure D over m -bit strings defined in Section 3.1 is key for our analysis. By (6), w has nonnegligible bias against tests of relative weight $O((\log n)^2/n)$ and tiny bias against all $n^{-\Omega(1)}$ balanced tests (assuming C is $n^{-\Omega(1)}$ -balanced). (In contrast, a symmetric product measure with the same expected Hamming weight has negligible bias beyond $w = O(\log n/n)$.)

Is there a better choice of D that remains biased against tests of weight $\omega((\log n)^2/n)$ yet remains pseudorandom against all balanced tests? Such a D would have improved the decoding radius η in Proposition 6. The improved algorithm would have then stood in contradiction to our lower bound Theorem 15. We summarize the conclusion:

► **Lemma 19.** *Assume $m > c(\ln 2/\beta)(\ln 1/\gamma)/\beta^2\eta$ for some absolute constant c . For any distribution D over $\{-1, 1\}^m$, if $\hat{D}(w) \geq \gamma$ for all w of weight at most ηm (for even ηm and $\eta \leq 1/4$), there must exist a β -balanced w for which*

$$\hat{D}(w) \geq \exp \left[-\frac{\ln 2/\beta \cdot \ln 1/\gamma}{\eta} \right].$$

In our application to decoding, when $\beta = n^{-\Theta(1)}$ and $\gamma = n^{-\Theta(1)}$, the lower bound reads $\hat{D}(w) \geq \exp -O((\log n)^2/\eta)$. When η is $\Theta((\log n)^2/n)$ some β -balanced test has bias at least $\exp -O(n)$, matching (6).

Lemma 19 is a statement about distributions over m -bit strings that makes no reference to codes and algorithms. Qualitatively, it says that if a distribution has noticeable bias against all light tests, then it must have some bias against some balanced test. We give a direct analytical proof of this.

Proof of Lemma 19. Let X be a sample of D and $t = \eta m$. Let I be a uniformly random index in $\{1, \dots, m\}$. By monotonicity of moments, for every $K \geq t$,

$$\mathbb{E}(\mathbb{E}X_I|X)^K \geq (\mathbb{E}(\mathbb{E}X_I|X)^t)^{K/t}.$$

We can expand $\mathbb{E}(\mathbb{E}X_I|X)^t$ as $\mathbb{E}X_{I_1} \cdots X_{I_t}$, where I_1, \dots, I_t are independent replicas of I . By our assumption on the light tests, for every fixing of I_1, \dots, I_t this expression is at least γ . Therefore

$$(\mathbb{E}(\mathbb{E}X_I|X)^t)^{K/t} \geq \gamma^{K/t}.$$

Let K be a Poisson random variable of rate λm (to be determined). Averaging over K we obtain

$$\mathbb{E}(\mathbb{E}X_I|X)^K \geq \mathbb{E}\gamma^{K/t} \cdot \mathbb{1}(K \geq t) \geq \mathbb{E}\gamma^{K/t} - \Pr(K < t).$$

as $\gamma \leq 1$. Using the formula for the Poisson moment generating function,

$$\mathbb{E}\gamma^{K/t} = \exp \lambda m (\gamma^{1/t} - 1) \geq \gamma^{\lambda m/t} = \gamma^{\lambda/\eta}.$$

and so

$$\mathbb{E}(\mathbb{E}X_I|X)^K \geq \gamma^{\lambda/\eta} - \Pr(\text{Poisson}(\lambda m) < \eta m). \quad (10)$$

The left-hand side also expands into a product of a K independent replicas $\mathbb{E}X_{I_1} \cdots X_{I_K}$. Let N_i be the number of times X_i occurs in this expression, i.e. N_i is the number of indices

j for which $I_j = i$. By the additivity of Poisson random variables, the N_i are independent Poissons of rate λ . Therefore

$$\mathbb{E}(\mathbb{E}X_I|X)^K = \mathbb{E}X_1^{N_1} \dots X_m^{N_m} = \mathbb{E}X_1^{\oplus N_1} \dots X_m^{\oplus N_m} = \mathbb{E}\hat{D}(\oplus N), \quad (11)$$

where $\oplus N_i = N_i \bmod 2$ and $\oplus N = (\oplus N_1, \dots, \oplus N_m)$, because $X_i^2 = 1$. The random variables $\oplus N_i$ inherit their independence from N_i . Marginally they are Bernoulli of bias

$$\mathbb{E}(-1)^{\oplus N_i} = \mathbb{E}(-1)^{N_i} = \exp -2\lambda.$$

again using the formula for the Poisson moment generating function.

We choose λ so that $\beta = 2 \exp -2\lambda$. By Chernoff bounds, $\oplus N$ is β -balanced except with probability $\exp -\Omega(\beta^2 m)$. By the total probability theorem,

$$\mathbb{E}\hat{D}(\oplus N) \leq \max_{w \text{ is } \beta\text{-balanced}} \hat{D}(w) + \exp -\Omega(\beta^2 m).$$

Plugging into (10) and (11),

$$\max_{w \text{ is } \beta\text{-balanced}} \hat{D}(w) \geq \gamma^{(\ln 2/\beta)/2\eta} - \exp -\Omega(\beta^2 m) - \Pr(\text{Poisson}(\lambda m) < \eta m).$$

Under our assumptions on m and η and the Poisson large deviation inequality [24, Theorem A.8] the leading term dominates and gives the desired bound. ◀

References

- 1 Damiano Abram, Giulio Malavolta, and Lawrence Roy. Trapdoor hash functions and PIR from low-noise LPN. Cryptology ePrint Archive, Paper 2025/416, 2025. URL: <https://eprint.iacr.org/2025/416>.
- 2 Dorit Aharonov and Oded Regev. Lattice problems in NP cap coNP. *J. ACM*, 52(5):749–765, 2005. doi:10.1145/1089023.1089025.
- 3 Michael Alekhnovich. More on average case vs approximation complexity. In *44th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2003*, pages 298–307. IEEE Computer Society, 2003. doi:10.1109/SFCS.2003.1238204.
- 4 Noga Alon. Problems and results in extremal combinatorics—i. *Discrete Mathematics*, 273(1):31–53, 2003. doi:10.1016/S0012-365X(03)00227-9.
- 5 Noga Alon, Jehoshua Bruck, Joseph Naor, Moni Naor, and Ron M. Roth. Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. *IEEE Transactions on Information Theory*, 38(2):509–516, 1992. Earlier version in ISIT '91. doi:10.1109/18.119713.
- 6 Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k -wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992. Earlier version in FOCS '90. doi:10.1002/rsa.3240030308.
- 7 Noga Alon, Rina Panigrahy, and Sergey Yekhanin. Deterministic approximation algorithms for the nearest codeword problem. In *12th International Workshop on Randomization and Computation, APPROX-RANDOM 2009*, volume 5687 of *Lecture Notes in Computer Science*, pages 339–351. Springer, 2009. doi:10.1007/978-3-642-03685-9_26.
- 8 Benny Applebaum, Naama Haramaty, Yuval Ishai, Eyal Kushilevitz, and Vinod Vaikuntanathan. Low-Complexity Cryptographic Hash Functions. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, volume 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:31, Dagstuhl, Germany, 2017. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2017.7.

- 9 Sanjeev Arora, László Babai, Jacques Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. In *34th Annual Symposium on Foundations of Computer Science, FOCS 1993*, pages 724–733. IEEE Computer Society, 1993. doi:10.1109/SFCS.1993.366815.
- 10 Anja Becker, Antoine Joux, Alexander May, and Alexander Meurer. Decoding random binary linear codes in $2^{n/20}$: How $1 + 1 = 0$ improves information set decoding. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 520–536. Springer, 2012. doi:10.1007/978-3-642-29011-4_31.
- 11 Avraham Ben-Aroya and Amnon Ta-Shma. Constructing small-bias sets from algebraic-geometric codes. *Theory Comput.*, 9:253–272, 2013. doi:10.4086/TQC.2013.V009A005.
- 12 Piotr Berman and Marek Karpinski. Approximating minimum unsatisfiability of linear equations. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '02*, pages 514–516. Society for Industrial and Applied Mathematics, 2002. URL: <http://dl.acm.org/citation.cfm?id=545381.545450>.
- 13 Daniel J. Bernstein. Grover vs. McEliece. In Nicolas Sendrier, editor, *Post-Quantum Cryptography, Third International Workshop, PQCrypto 2010*, volume 6061 of *Lecture Notes in Computer Science*, pages 73–80. Springer, 2010. doi:10.1007/978-3-642-12929-2_6.
- 14 Vijay Bhattiprolu, Venkatesan Guruswami, and Xuandi Ren. PCP-free APX-hardness of nearest codeword and minimum distance. *CoRR*, abs/2503.11131, 2025. doi:10.48550/arXiv.2503.11131.
- 15 Nir Bitansky and Sapir Freizeit. Statistically sender-private OT from LPN and derandomization. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part III*, volume 13509 of *Lecture Notes in Computer Science*, pages 625–653. Springer, 2022. doi:10.1007/978-3-031-15982-4_21.
- 16 Nir Bitansky, Prahladh Harsha, Yuval Ishai, Ron D. Rothblum, and David J. Wu. Dot-product proofs and their applications. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024*, pages 806–825. IEEE Computer Society, 2024. doi:10.1109/FOCS61266.2024.00057.
- 17 Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, volume 773 of *Lecture Notes in Computer Science*, pages 278–291. Springer, 1993. doi:10.1007/3-540-48329-2_24.
- 18 Leif Both and Alexander May. Optimizing bjmm with nearest neighbors: full decoding in $22/21n$ and mceliece security. In *WCC workshop on coding and cryptography*, volume 214, 2017.
- 19 Leif Both and Alexander May. Decoding linear codes with high error rate and its impact for LPN security. In Tanja Lange and Rainer Steinwandt, editors, *Post-Quantum Cryptography - PQCrypto 2018*, volume 10786 of *Lecture Notes in Computer Science*, pages 25–46. Springer, 2018. doi:10.1007/978-3-319-79063-3_2.
- 20 Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Nicolas Resch, and Peter Scholl. Correlated pseudorandomness from expand-accumulate codes. In *Advances in Cryptology - CRYPTO 2022: 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part II*, pages 603–633. Berlin, Heidelberg, 2022. Springer-Verlag. doi:10.1007/978-3-031-15979-4_21.
- 21 Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Correlated pseudorandom functions from variable-density LPN. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 1069–1080. IEEE, 2020. doi:10.1109/FOCS46700.2020.00103.

- 22 Zvika Brakerski, Alex Lombardi, Gil Segev, and Vinod Vaikuntanathan. Anonymous ibe, leakage resilience and circular security from new assumptions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 535–564. Springer, 2018. doi:10.1007/978-3-319-78381-9_20.
- 23 Zvika Brakerski, Vadim Lyubashevsky, Vinod Vaikuntanathan, and Daniel Wichs. Worst-case hardness for LPN and cryptographic hashing via code smoothing. In *Advances in Cryptology - EUROCRYPT 2019*, volume 11478 of *Lecture Notes in Computer Science*, pages 619–635. Springer, 2019. doi:10.1007/978-3-030-17659-4_21.
- 24 Clément L. Canonne. Topics and techniques in distribution testing: A biased but representative sample. *Foundations and Trends in Communications and Information Theory*, 19(6):1032–1198, 2022. doi:10.1561/0100000114.
- 25 Kevin Carrier, Thomas Debris-Alazard, Charles Meyer-Hilfiger, and Jean-Pierre Tillich. Statistical decoding 2.0: Reducing decoding to lpn. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 477–507. Springer, 2022.
- 26 Kévin Carrier, Thomas Debris-Alazard, Charles Meyer-Hilfiger, and Jean-Pierre Tillich. Reduction from sparse lpn to lpn, dual attack 3.0. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 286–315. Springer, 2024. doi:10.1007/978-3-031-58754-2_11.
- 27 Geoffroy Couteau, Peter Rindal, and Srinivasan Raghuraman. Silver: Silent vole and oblivious transfer from hardness of decoding structured ldpc codes. In *Advances in Cryptology - CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part III*, pages 502–534, Berlin, Heidelberg, 2021. Springer-Verlag. doi:10.1007/978-3-030-84252-9_17.
- 28 Thomas Debris-Alazard and Nicolas Resch. Worst and average case hardness of decoding via smoothing bounds. In *Public-Key Cryptography - PKC 2025 - 28th IACR International Conference on Practice and Theory of Public-Key Cryptography, Røros, Norway, May 12–15, 2025, Proceedings, Part II*, volume 15675 of *Lecture Notes in Computer Science*, pages 363–392. Springer, 2025. doi:10.1007/978-3-031-91823-0_12.
- 29 Léo Ducas, Andre Esser, Simona Etinski, and Elena Kirshanova. Asymptotics and improvements of sieving for codes. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024*, volume 14656 of *Lecture Notes in Computer Science*, pages 151–180. Springer, 2024. doi:10.1007/978-3-031-58754-2_6.
- 30 Léo Ducas and Ludo N Pulles. Accurate score prediction for dual-sieve attacks. *Cryptology ePrint Archive*, 2023.
- 31 Uriel Feige and Daniele Micciancio. The inapproximability of lattice and coding problems with preprocessing. *Journal of Computer and System Sciences*, 69(1):45–67, 2004. doi:10.1016/j.jcss.2004.01.002.
- 32 Joel Friedman and Jean-Pierre Tillich. Generalized Alon–Boppana theorems and error-correcting codes. *SIAM Journal on Discrete Mathematics*, 19(3):700–718, 2005. doi:10.1137/S0895480102408353.
- 33 Oded Goldreich and Shafi Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.*, 60(3):540–563, 2000. doi:10.1006/JCSS.1999.1686.
- 34 A Al Jabri. A statistical decoding algorithm for general linear block codes. In *IMA International Conference on Cryptography and Coding*, pages 1–8. Springer, 2001.
- 35 Ghazal Kachigar and Jean-Pierre Tillich. Quantum information set decoding algorithms. In Tanja Lange and Tsuyoshi Takagi, editors, *Post-Quantum Cryptography - PQCrypto 2017*, volume 10346 of *Lecture Notes in Computer Science*, pages 69–89. Springer, 2017. doi:10.1007/978-3-319-59879-6_5.
- 36 Vladimir I Levenshtein. Bounds for packings of metric spaces and some of their applications. *Problemy Kibernet*, 40:43–110, 1983.

- 37 Yi-Kai Liu, Vadim Lyubashevsky, and Daniele Micciancio. On bounded distance decoding for general lattices. In *9th International Workshop on Randomization and Computation, APPROX-RANDOM 2006*, volume 4110 of *Lecture Notes in Computer Science*, pages 450–461. Springer, 2006. doi:10.1007/11830924_41.
- 38 Alexander May, Alexander Meurer, and Enrico Thomae. Decoding random linear codes in $O(2^{0.054n})$. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology - ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 107–124. Springer, 2011. doi:10.1007/978-3-642-25385-0_6.
- 39 Alexander May and Ilya Ozerov. On computing nearest neighbors with applications to decoding of binary linear codes. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 203–228. Springer, 2015. doi:10.1007/978-3-662-46800-5_9.
- 40 R. McEliece, E. Rodemich, H. Rumsey, and L. Welch. New upper bounds on the rate of a code via the Delsarte-MacWilliams inequalities. *IEEE Trans. Inf. Theor.*, 23(2):157–166, September 2006. doi:10.1109/TIT.1977.1055688.
- 41 Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993. Earlier version in STOC '90. doi:10.1137/0222053.
- 42 Raphael Overbeck. Statistical decoding revisited. In *Australasian Conference on Information Security and Privacy*, pages 283–294. Springer, 2006. doi:10.1007/11780656_24.
- 43 Eugene Prange. The use of information sets in decoding cyclic codes. *IRE Trans. Inf. Theory*, 8(5):5–9, 1962. doi:10.1109/TIT.1962.1057777.
- 44 Oded Regev. Improved inapproximability of lattice and coding problems with preprocessing. *IEEE Trans. Inf. Theory*, 50(9):2031–2037, 2004. Preliminary version in CCC 2003. doi:10.1109/TIT.2004.833350.
- 45 Noah Stephens-Davidowitz and Vinod Vaikuntanathan. SETH-hardness of coding problems. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019*, pages 287–301. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00027.
- 46 Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC '17, pages 238–251, New York, NY, USA, 2017. ACM. Also available as ECCC Technical Report TR17-041. doi:10.1145/3055399.3055408.
- 47 Lloyd R. Welch. Lower bounds on the maximum cross correlation of signals. *IEEE Trans. Inf. Theory*, 20:397–399, 1974.
- 48 Yu Yu and Jiang Zhang. Smoothing out binary linear codes and worst-case sub-exponential hardness for LPN. In *Advances in Cryptology - CRYPTO 2021*, volume 12827 of *Lecture Notes in Computer Science*, pages 473–501. Springer, 2021. doi:10.1007/978-3-030-84252-9_16.
- 49 Yu Yu, Jiang Zhang, Jian Weng, Chun Guo, and Xiangxue Li. Collision resistant hashing from sub-exponential learning parity with noise. In *Advances in Cryptology - ASIACRYPT 2019: 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8–12, 2019, Proceedings, Part II*, pages 3–24, Berlin, Heidelberg, 2019. Springer-Verlag. doi:10.1007/978-3-030-34621-8_1.