

Advanced cryptology (GBX9SY06)

*

Some coding theory aspects (useful) in cryptography

Pierre Karpman

2021-01

1 First definitions; examples

A *linear code* of length n and dimension k over a field \mathbb{K} is a k -dimensional subspace of \mathbb{K}^n . In these notes, we will focus on *binary* codes, for which $\mathbb{K} = \mathbb{F}_2$, or possibly an extension thereof.* An important characteristic of a code is the *minimum distance* (for the Hamming distance $hd(\cdot, \cdot)$) d between two (distinct) codewords. Define $wt(\mathbf{x})$, $\mathbf{x} \in \mathbb{F}_2^n$ as the number of non-zero coordinates of \mathbf{x} ; $hd(\mathbf{x}, \mathbf{y})$ as $wt(\mathbf{x} + \mathbf{y})$. Then the minimum distance of a code \mathcal{C} is $\min_{\mathbf{x} \in \mathcal{C}, \mathbf{y} \in \mathcal{C}, \mathbf{x} \neq \mathbf{y}} hd(\mathbf{x}, \mathbf{y})$, which by linearity of \mathcal{C} is equivalent to $\min_{\mathbf{x} \neq \mathbf{0} \in \mathcal{C}} wt(\mathbf{x})$. The main parameters of length, dimension, and minimum distance of a code \mathcal{C} are summarized by saying that \mathcal{C} is an $[n, k, d]_{\mathbb{F}_2}$ code. While determining n and k is usually straightforward, it is in general a hard problem to compute d .

Given a code \mathcal{C} , it will often be necessary to have an explicit *encoding* map $\mathbf{x} \in \mathbb{F}_2^k \mapsto \mathbf{x}' \in \mathcal{C}$ from *messages* to *codewords*. Such a map can be easily obtained by sampling k linearly-independent codewords $\mathbf{g}_1, \dots, \mathbf{g}_k$ and forming the *generator matrix* \mathbf{G} whose rows are the \mathbf{g}_i s.† The encoding map is then simply $\mathbf{x} \mapsto \mathbf{x} \times \mathbf{G}$. One should remark that in general the matrix \mathbf{G} (and thence the encoding map) will not be unique, as it depends on the selected codewords. A specific class of generator matrices are the ones in *systematic form*, corresponding to block matrices $(\mathbf{I}_k \ \mathbf{A})$, where \mathbf{I}_k is the k -dimensional identity matrix and \mathbf{A} is a *redundancy* block. A code always admits at least one systematic encoder, up to a permutation of the coordinates of its codewords. One may be obtained by selecting k linearly-independent columns of a generator matrix \mathbf{G} ; applying a column permutation on \mathbf{G} such that those columns are in the first k positions; and computing the reduced row-echelon form of \mathbf{G} . In other words, one obtains an encoder in systematic form by finding a permutation matrix \mathbf{P} such that \mathbf{GP} is of the form $\mathbf{G}' := (\mathbf{G}_1 \ \mathbf{G}_2)$ where \mathbf{G}_1 is invertible, and by computing $\mathbf{G}_1^{-1}\mathbf{G}'$.

From the existence of a systematic encoder, one deduces that the largest possible minimum distance (or weight) of an $[n, k]$ linear code is $d_{\text{MDS}} = n - k + 1$, which is a special case of the Singleton bound. Indeed, the maximum possible weight of any row of a systematic encoder is 1 on the left (identity) block, and $n - k$ on the right (redundancy) block. A code reaching this bound is called *maximum-distance separable*, or MDS.

Finally, note that for some codes, there may exist alternative encoders that do not explicitly use a generator matrix.

*Consequently, we may take the liberty of equating subtraction with addition in any formula or algorithm. In order to minimize the confusion, we will try to make this systematic. Nonetheless, most of the discussion seamlessly generalises to other (finite) fields.

†Note that we use the convention that vectors are *row* vectors, if not specified otherwise.

Example 1 (AES MixColumn). Let \mathbf{M} be the matrix used in the MixColumn operation of the AES block cipher. The code generated by $(\mathbf{I}_4 \ \mathbf{M})$ is an MDS code of parameters $[8, 4, 5]_{\mathbb{F}_{2^8}}$.

Example 2 (Binary Reed-Muller codes). The binary Reed-Muller code of order r and in m variables $\text{RM}(r, m)$ is the vector-space formed by the multi-point evaluations of m -variate Boolean functions of degree $\leq r$ over \mathbb{F}_2^m . In other words, a message is a Boolean function, and its associated (Reed-Muller) codeword is obtained by evaluating it over its entire domain.

The codewords of this code have length 2^m and form a space of dimension $k := \sum_{i=0}^r \binom{m}{i}$. It can be shown that the minimum weight of any codeword is 2^{m-r} [MS06, Ch. 13, Thm. 3]. The parameters of $\text{RM}(r, m)$ are thus $[2^m, k, 2^{m-r}]_{\mathbb{F}_2}$.

A (non-systematic) encoding of a message can be efficiently computed by using a fast Möbius transform. Due to its involutive nature, the same transform can also be used to decode a codeword back to a message. However, this does not correct any error and on its own it is thus of rather limited use.

Reed-Muller codes also follow a recursive “ $(u, u + v)$ ” decomposition. One has that $\text{RM}(r + 1, m + 1) = \{u \mid (u + v), u \in \text{RM}(r + 1, m), v \in \text{RM}(r, m)\}$. This follows from the fact that an $(m + 1)$ -variate Boolean function $F(X_1, \dots, X_{m+1})$ of degree at most $r + 1$ can be written as $F^0(X_1, \dots, X_m) + X_{m+1} F^1(X_1, \dots, X_m)$, with $\deg(F^0) \leq r + 1$ and $\deg(F^1) \leq r$. Furthermore, if we write F^{0+} the $m + 1$ -variate function whose monomials are identical to F^0, F^{1+} for $X^{m+1} F^1$, and $\vec{X}_{1,m}$ a given assignment for the indeterminates X_1, \dots, X_m , then we always have:

- $\text{eval}(F^0, (\vec{X}_{1,m})) = \text{eval}(F^{0+}, (\vec{X}_{1,m}, 0)) = \text{eval}(F^{0+}, (\vec{X}_{1,m}, 1));$
- $\text{eval}(F^1, (\vec{X}_{1,m})) = \text{eval}(F^{1+}, (\vec{X}_{1,m}, 1));$
- $\text{eval}(F^{1+}, (\vec{X}_{1,m}, 0)) = 0;$

and the decomposition follows. Finally, one may notice that this is essentially the same induction as the one used in the fast Möbius transform algorithm.

Given a code \mathcal{C} , it is often important to be able to determine if a vector of its ambient space is a codeword or not. This may be done using a map $\mathbf{x} \mapsto \mathbf{y}$ s.t. \mathbf{y} is “zero” iff. $\mathbf{x} \in \mathcal{C}$. One typically implements this with a *parity-check matrix* $\mathbf{H} \in \mathbb{F}_2^{n-k \times n}$ which is a basis of the (right) kernel of a generator matrix \mathbf{G} of \mathcal{C} ; the corresponding map, of codomain \mathbb{F}_2^{n-k} , is then $\mathbf{x} \mapsto \mathbf{H} \times \mathbf{x}$. Equivalently, \mathbf{H} is made of $(n - k)$ linearly-independent vectors of \mathbb{F}_2^n whose scalar product with any element of \mathcal{C} is zero, and $\mathbf{H}\mathbf{G}^t$ and $\mathbf{G}\mathbf{H}^t$ are both zero matrices. A parity-check matrix generates the *dual* of \mathcal{C} , written \mathcal{C}^\perp , which is thence an $[n, n - k]$ code. A code that is its own dual is called *self-dual*.

2 Information set decoding

In this section we focus on the problem of finding “low-weight” codewords of a code, which is also essentially equivalent to finding a close-by codeword to a given vector, i.e. to decode. This is generally a hard problem for codes that do not exhibit any particular structure (for instance if they are defined from a uniformly random generator matrix), as it is NP-hard [BMvT78], but efficient algorithms may exist for some specific codes. For now we will focus on “inefficient” generic algorithms that work for any code, but we will later present a good *list decoder* for (punctured and shortened) first-order Reed-Muller codes.

Let \mathcal{C} be an $[n, k, d]$ code for which \mathbf{G} is a generator matrix. Enumerating all the codewords of \mathcal{C} can trivially be done in time 2^k by multiplying \mathbf{G} by all the vectors of \mathbb{F}_2^k .

This immediately allows to find a weight- d codeword of \mathcal{C} or to decode to the (or one of the) closest codeword(s), but the cost is quickly prohibitive.

A first remark on the way to find better alternatives is that the problem that one needs to solve usually does not require to enumerate all the codewords of \mathcal{C} . For instance, one may not need to find a codeword of *minimum* weight, but finding one of weight less than a known bound may be enough. Similarly, in a decoding context, one may know an upper-bound on the error weight. It is then possible to use a probabilistic algorithm that stops when a “good-enough” solution has been found.

A second remark is that the decoding problem can be solved by finding low-weight codewords; this will be used to justify the fact that we solely focus on algorithms for the latter. This can be explained in the following way: let $\mathbf{c} \in \mathcal{C}$ be an initial codeword, and $\hat{\mathbf{c}} = \mathbf{c} + \mathbf{e}$ be a noisy codeword obtained by adding a noise \mathbf{e} of weight $w \leq \max \text{weight} < \lfloor (d-1)/2 \rfloor$. Then \mathbf{c} is the unique closest codeword to $\hat{\mathbf{c}}$, and \mathbf{e} is the unique vector of weight w in the affine subspace $\mathbf{e} + \mathcal{C}$. Furthermore, this latter vector can be found by searching for a “codeword” of weight w in the code generated by $\begin{pmatrix} \mathbf{G} \\ \hat{\mathbf{c}} \end{pmatrix}$. This codeword will even be unique, as for any $\mathbf{c}' \neq \mathbf{c} \in \mathcal{C}$, $wt(\mathbf{c}' + \hat{\mathbf{c}}) = wt(\mathbf{c}' + \mathbf{c} + \mathbf{e}) \geq wt(\mathbf{c}' + \mathbf{c}) - wt(\mathbf{e}) \geq d - \max \text{weight} > w$.

In practice decoding algorithms usually do not exploit this reduction; nonetheless the internal structures of generic algorithms used in *syndrome decoding* are essentially identical to the one of algorithms searching for low-weight codewords. Those solve the dual problem of finding a low-weight linear combination of vectors of a certain basis that sum to the null vector, where the weight of a combination is the number of vectors that appear in it with a non-zero coefficient.

The first probabilistic alternative to exhaustive search that we present is quite simple [Pra62, McE78]. Given \mathbf{G} , randomly select k linearly-independent columns; this is called an *information set*. Then permute these columns to the first k positions of \mathbf{G} , and compute the reduced row-echelon form (i.e. compute an alternative generator matrix \mathbf{G}' in systematic form, associated to the selected information set). Finally, check if any of the resulting k rows have a weight less than the input bound. The idea behind this algorithm is that any row of the obtained systematic encoder has by definition a very low weight of exactly one on its first k positions, and the weight on the remaining $n - k$ positions depends on a random codeword linear combination. One then hopes that for *some* information sets, the weight on these latter positions will also be small, resulting in an overall low-weight codeword. In other words, the algorithm will return a weight- w codeword after examining a given information set if it is s.t. there is a codeword of weight 1 over the information set and of weight $w - 1$ over its complement, the *redundancy set*.

There is also a simple interpretation of this algorithm if one directly thinks of it in a “decoding” sense. An information set is by definition a set of positions that carries enough information to fully determine the message corresponding to a codeword. Indeed, given the value of a codeword on an information set, one can reconstruct the entire (non-noisy) codeword by simply applying an encoder systematic w.r.t. this set; it is then easy to invert the encoding to go back to the original message. Thus, what the above does is (randomly) trying to find an information set over which the error vector is all-zero.

A variant of the above first algorithm due to Lee and Brickell [LB88] consists, for each information set, in checking the weight of all linear combinations of rows of \mathbf{G}' of weight less than a small value p (typically 2 or 3). This somehow amortizes the cost of computing \mathbf{G}' by considering more codewords for each matrix, as now the algorithm returns on a given information set if a codeword’s weight splits as $(i, w - i)$, $1 \leq i \leq p$ over itself and its complement. Also, note that for binary codes, computing all of these can be done particularly efficiently by using Gray codes.

Another variant due to Leon aims to reduce the practical cost of checking the weight of a codeword, and is thus mostly useful for long codes. The idea is simply to first check if a codeword generated from the above procedure has a small weight on a few positions (i.e. to first consider a short *punctured* code), and only to look at the entire codeword in that case. For instance, if one requires the punctured codeword to have weight zero on its redundancy set of size l , one is in effect searching for codewords whose weight splits as $(i, 0, w - i), 1 \leq i \leq p$ over the information set, and the non-punctured (resp. punctured) redundancy positions.

While this approach looks at fewer candidates per information set as the Lee-Brickell algorithm, this is hoped to be counter-balanced by more efficient implementations.

Another algorithm due to Stern still follows the overall same approach, but improves the time complexity at the cost of some memory [Ste88]. The key idea is to split the search space into two lists and to exploit collisions to obtain a quadratic speed-up at some stage of the search. Starting from the initial algorithm, one splits the information set into two subsets I_1 and I_2 , and forms the lists Λ_1 and Λ_2 of codewords of weight less than p on each subset respectively. Then one only fully checks the weight of codewords formed by the sum of elements of Λ_1 and Λ_2 that are identical over l prescribed positions Z of the redundancy set. One is then searching for codewords whose weight splits as $(i, j, 0, w - i - j), 1 \leq i, j \leq p$ over I_1, I_2, Z and the remainder.

It is essential to notice that for a given information set, checking for each of the $\#\Lambda_1\#\Lambda_2$ candidate codewords if it is of the above form indeed takes a cost linear in $\#\Lambda_{1,2}$ (by using an appropriate data structure).

Finally, one may remark that this algorithm takes more input parameters than the previous ones. This, together with the fact that it is not memory-less may make it harder to determine what parameter choice is best suited to a given code.

An important observation made by Canteaut and Chabaud [CC98] is that the most expensive step in the above algorithms is the computation of the systematic encoder for a given information set. They then suggest that instead of selecting a new independent information set at every iteration, one may “update” the current set by randomly replacing one of its columns by one column of the redundancy set, which is much more efficient (with the obvious drawback that not all the information set is replaced anymore). Furthermore, one can easily be convinced that after a few iterations, the obtained information set will be essentially independent from the starting one, hence there is no risk that one gets stuck in a small subset of the search space considered by the other algorithms. To carry out a rigorous analysis of the algorithm, it is appropriate to model this random walk as a Markov process.

We will conclude by describing how to efficiently update an information set. Let $\mathbf{G} = (\mathbf{I} \ \mathbf{A})$ be a systematic generator matrix; our objective is to compute $\mathbf{G}' = (\mathbf{I} \ \mathbf{A}')$ which is a generator matrix for the same code and equal to the reduced row-echelon form of a matrix obtained from \mathbf{G} by swapping one column $\mathbf{I}_{\cdot,i}$ of the identity with one column $\mathbf{A}_{\cdot,j}$ of the redundancy matrix. First notice that this latter process only results in a systematic matrix if $\mathbf{A}_{\cdot,j}$ is linearly independent from $\mathbf{I} \setminus \mathbf{I}_{\cdot,i}$, which is equivalent to requiring that $\mathbf{A}_{i,j} = 1 \neq 0$. Second, the matrix \mathbf{A}' is simply obtained from \mathbf{A} by adding the row \mathbf{A}_i to every row $\mathbf{A}_{i'}$ where $\mathbf{A}_{i',j} = 1$. Indeed, this corresponds to the “reduction” step one needs to perform after swapping the above two columns.

3 Learning Parity with Noise cryptosystems

In this section, we introduce the *Learning Parity with Noise* (LPN) problem, and its application to the design of cryptosystems. The LPN problem is rather attractive because of its very concise description. Let $\mathbf{s} \xleftarrow{\$} \mathbb{F}_2^k, \mathbf{a} \xleftarrow{\$} \mathbb{F}_2^k, e \leftarrow \text{Ber}_\eta$, where Ber_η is the Bernoulli

distribution of parameter η ; that is, $\Pr[e = 1] = \eta$. Then the LPN problem is to guess the value of the scalar product $\mathbf{s} \cdot \mathbf{a}$ when given $(\mathbf{a}, \mathbf{s} \cdot \mathbf{a} + e)$. An algorithm is then said to solve this problem with advantage ε if it answers correctly with probability p and $2|p - 1/2| = \varepsilon$. Another, equivalent formulation of the problem is to ask one to distinguish $(\mathbf{a}, \mathbf{s} \cdot \mathbf{a} + e)$ from $(\mathbf{a}, r \stackrel{\$}{\leftarrow} \mathbb{F}_2)$.

It is clear that if $\eta = 1/2$, the distribution of $\mathbf{s} \cdot \mathbf{a} + e$ is independent of \mathbf{a} and \mathbf{s} , and no algorithm can succeed with a non-zero advantage. Similarly, without prior knowledge about \mathbf{s} and except when $\mathbf{a} = \mathbf{0}$, one cannot hope to solve the problem given a *single* query of the above form, even in the absence of noise (i.e. even when $\eta = 0$) since in that case $\mathbf{s} \cdot \mathbf{a}$ is uniformly distributed. It is then natural to extend the problem to arbitrarily-many queries q , where the unknown (secret) vector \mathbf{s} is kept constant. One may then reformulate the problem as letting $\mathbf{A} \stackrel{\$}{\leftarrow} \mathbb{F}_2^{k \times q}$, $\mathbf{e} \leftarrow \text{Ber}_{\eta, q}$ (that is, each bit of $\mathbf{e} \in \mathbb{F}_2^q$ has independent probability η to be equal to one), and asking to distinguish $(\mathbf{A}, \mathbf{sA} + \mathbf{e})$ from (\mathbf{A}, \mathbf{u}) where $\mathbf{u} \stackrel{\$}{\leftarrow} \mathbb{F}_2^q$. An algorithm is then said to (t, q) solve $\text{LPN}_{k, \eta}$ with advantage ε if it makes q queries and has running time t (in some metric).

It is now again clear that if $\eta = 0$, the problem is trivially solvable as long as $\text{rank}(\mathbf{A}) = k$, as it is then enough to identify k linearly-independent columns of \mathbf{A} to recover \mathbf{s} (which then allows to predict all the other queries with advantage 1). In the more meaningful case where $\eta > 0$, recovering \mathbf{s} becomes equivalent to decoding a noisy codeword for some random code of length q (which is a parameter that may be chosen by the solving algorithm). We will discuss this matter in more details in the next section, and focus for now on some LPN-based cryptosystems whose security depends on the computational hardness of this problem.

We first describe LPN-C, which is a family of symmetric encryption schemes defined by Gilbert et al. [GRS08]. While the confidentiality of an LPN-C instance reduces to the hardness of a corresponding LPN problem, the scheme is inherently malleable and must thus be used in conjunction with a MAC.

An LPN-C instance is parameterized by a random code length k , a noise level η , a message length r , and the parameters of an $[m, r, d]$ binary code \mathcal{C} , assumed to be efficiently decodable up to w errors and s.t. $\Pr[wt(\mathbf{e}) > w : \mathbf{e} \leftarrow \text{Ber}_{\eta, m}]$ is small. The scheme works as follows. The sender and the receiver first share a secret random matrix $\mathbf{M} \in \mathbb{F}_2^{k \times m}$. Then, to encrypt an r -bit message \mathbf{x} , the sender draws a vector $\mathbf{a} \stackrel{\$}{\leftarrow} \mathbb{F}_2^k$ and $\mathbf{e} \leftarrow \text{Ber}_{\eta, m}$, computes $\mathbf{y} = \mathcal{C}(\mathbf{x}) + \mathbf{aM} + \mathbf{e}$, and sends (\mathbf{a}, \mathbf{y}) to the receiver. To decrypt, the receiver computes $\hat{\mathbf{y}} = \mathbf{y} + \mathbf{aM} = \mathcal{C}(\mathbf{x}) + \mathbf{e}$, and uses the decoder of \mathcal{C} to recover \mathbf{x} . The designers of LPN-C proposed some parameters for secure instantiations of LPN-C, but without specifying which code \mathcal{C} to choose. An example is to take $k = 768$, $m = 160$, $\eta = 1/20$, $r = 75$, $d = 25$.

An informal way to argue about the security of this scheme is that if one uses parameters for which the LPN problem is hard to solve, then the (encoded) message $\mathcal{C}(\mathbf{x})$ is whitened by a pseudo-random mask $\mathbf{aM} + \mathbf{e}$ which is hard to distinguish from random, and is thus encrypted by a secure “stream cipher”.[‡]

We now turn to an LPN-based public-key cryptosystem due to Alekhovich [Ale03]. This is a highly impractical design, as it only encrypts a single bit and the decryption of “1” fails with probability $1/2$. It is however of theoretical interest, and is rather simple to describe.

This scheme is parameterized by an integer n , from which one derives $m = 2n$, $k =$

[‡]There is a slight difference between this setting and the actual definition of LPN that we have used: here the *matrix* is secret and the vector \mathbf{a} is public. Yet, this simply corresponds to m “single” queries for m independent secrets batched together into \mathbf{M} .

$n^{1/2-\varepsilon}$, $\eta = k/n$. The public key is a matrix $\mathbf{A}' \in \mathbb{F}_2^{n+1 \times m}$, generated as $\begin{pmatrix} \mathbf{A} \\ \hat{\mathbf{c}} \end{pmatrix}$ where $\mathbf{A} \stackrel{\$}{\leftarrow} \mathbb{F}_2^{m \times m}$, $\hat{\mathbf{c}} = \mathbf{x}\mathbf{A} + \mathbf{e}$, $\mathbf{x} \stackrel{\$}{\leftarrow} \mathbb{F}_2^n$, $\mathbf{e} \leftarrow \text{Ber}_{\eta,m}$. The private key is the vector \mathbf{e} . In other words, one defines a random $[2n, n]$ code with generating matrix \mathbf{A} and augments it with a low-weight codeword \mathbf{e} to the code \mathcal{C} generated by \mathbf{A}' . One can indeed check that \mathbf{e} is in the span of \mathbf{A}' , as it is equal to $(\mathbf{x} \ 1) \mathbf{A}'$.[§] This augmentation is however done in a “hidden” way, as recovering \mathbf{e} from \mathbf{A}' is a (hard) decoding problem for the code defined by \mathbf{A} .

To encrypt one bit for the public key \mathbf{A}' , the sender proceeds as follows. To encrypt the bit 1, it sends a vector $\alpha_1 \stackrel{\$}{\leftarrow} \mathbb{F}_2^m$. To encrypt the bit 0, it computes and sends $\alpha_0 = \mathbf{d} + \mathbf{e}'$, where $\mathbf{e}' \leftarrow \text{Ber}_{\eta,m}$ and \mathbf{d} is a uniformly random element of \mathcal{C}^\perp . The receiver computes the decryption of α as $\alpha \cdot \mathbf{e}$.

It is quite immediate to see that upon receiving α_1 , decryption will fail with probability $1/2$. It is a bit less obvious that the decryption of α_0 is more successful. In that case, the receiver computes $\alpha_0 \cdot \mathbf{e} = \mathbf{d} \cdot \mathbf{e} + \mathbf{e}' \cdot \mathbf{e}$. Because $\mathbf{d} \in \mathcal{C}^\perp$, $\mathbf{e} \in \mathcal{C}$, this simplifies to $\mathbf{e} \cdot \mathbf{e}'$. Finally, the probability that \mathbf{e} and \mathbf{e}' have a non-disjoint support is $\approx (1 - \eta)^k$ (as it is enough for \mathbf{e} and \mathbf{e}' to have disjoint supports, and that their expected weights are k), which is asymptotically less than 0.5.

We already have informally argued that computing the secret key \mathbf{e} from the public key \mathbf{A}' reduces to a hard decoding problem. We may now also remark that distinguishing α_0 from α_1 reduces to an LPN-like problem for a generating matrix of \mathcal{C}^\perp . Alekhnovich then showed that the latter is computationally indistinguishable from a random matrix (informally this also follows from the hardness of LPN, since in that case $\hat{\mathbf{c}}$ “looks random”, and then so does \mathbf{A}' and a basis of its kernel), which allows to conclude the reduction to LPN.

We conclude this overview by describing an LPN-based symmetric authentication protocol named Lapin [HKL⁺12]. Strictly speaking, Lapin is based on the Ring-LPN variant of the problem, whose aim is to decrease the communication complexity. We will first describe it in the LPN framework, and will address this difference next.

A challenger and a verifier share two secret vectors $\mathbf{s}, \mathbf{s}' \in \mathbb{F}_2^n$. To authenticate the challenger, the verifier draws $\mathbf{C} \stackrel{\$}{\leftarrow} \mathbb{F}_2^{n \times n}$ and send it to the former. The challenger then draws $\mathbf{R} \stackrel{\$}{\leftarrow} \text{GL}(n, \mathbb{F}_2)$, $\mathbf{e} \leftarrow \text{Ber}_{\eta,n}$, and sends $(\mathbf{R}, (\mathbf{s}\mathbf{C} + \mathbf{s}')\mathbf{R} + \mathbf{e})$. The verifier then recovers \mathbf{e} and validates the challenge if it is of weight less than $\mu\eta n$ for some small acceptance threshold μ .

The Ring-LPN variant of the protocol works similarly, but works over modular rings of the form $\mathbb{F}_2[X]/\langle f \rangle$, for some polynomial f . The challenge matrix \mathbf{C} is replaced by a ring element $\pi(c)$, $c \in \mathbb{F}_2^\lambda$, where π is a mapping verifying some conditions; the matrix \mathbf{R} is replaced by an invertible element of the ring r ; the secrets are now also ring elements s and s' , and so is the error e (still drawn from a Bernoulli distribution, when seen as a vector). The message sent by the challenger is then simply $(r, (s\pi(c) + s')r + e)$. The advantage of this variant over the matrix version of the protocol is that the ring elements (esp. r) have a much more compact representation of size $\approx n$ than $n \times n$ random matrices, which essentially decreases the communication cost by a factor n . However, the security now depends on the hardness of decoding codes possessing some structure, which may allow for more efficient algorithms.

Finally, a (still informal) way to argue about the security of this protocol is to notice that the challenger’s answer can be written (in the original LPN case) as $\mathbf{s}\mathbf{C}\mathbf{R} + \mathbf{s}'\mathbf{R} + \mathbf{e}$. In other words, one is masking the challenge-and-secret-dependent string $\mathbf{s}\mathbf{C}\mathbf{R}$ by an LPN query $\mathbf{s}'\mathbf{R} + \mathbf{e}$, which is by assumption computationally indistinguishable from random.

[§]One may also remark that for the chosen parameters, $\mathbf{e} \notin \text{span}(\mathbf{A})$ with high probability. In the unlikely event where this would be the case, one can simply choose another vector and try again.

4 LPN-solving algorithms

It is quite clear that solving an LPN instance can be done by using generic decoding algorithms. Indeed, one may simply try to solve the *search* variant of the LPN problem, which on $\mathbf{sA} + \mathbf{e}$ tries to recover \mathbf{s} , i.e. tries to decode a noisy codeword of the code generated by \mathbf{A} . As solving the search problem allows to solve the distinguishing problem as well, the hardness of LPN is not more than the one of decoding.

One subtlety that we already mentioned is that an LPN-solving algorithm is free to choose the number of oracle queries it wants to use, i.e. the *length* of the code it wishes to decode. Surely this has to be sufficiently large to even make a successful decoding possible (for instance one may try to ensure that there is at least one error-free information set w.h.p.), but there is no similar constraint on the maximum number of queries.

We will now describe an algorithm that specifically exploits the ability to make many queries in an LPN problem to decrease the time complexity for solving the (search) problem (at the cost of a potentially huge increase in memory and query complexity). This algorithm was first described in this context by Blum, Kalai and Wasserman [BKW03], and is traditionally called *BKW*; we heavily base our presentation on the variant described by Bernstein and Lange [BL12].

We first describe a simple variant of BKW. An intuition behind the algorithm is that given sufficiently many LPN queries $\mathbf{s} \cdot \mathbf{a} + e$ with a *fixed*, vector \mathbf{a} , one can efficiently recover one bit of \mathbf{s} (viz. $\mathbf{s} \cdot \mathbf{a}$) using a majority vote among the samples. This requires $\approx c^{-2}$ samples, where $c = 1 - 2\eta$ ($0 < \eta < 1/2$) is the correlation of the noise. Repeating this process for k linearly independent masks \mathbf{a} leads to a full recovery of \mathbf{s} .

Of course, as the masks \mathbf{a} of LPN queries are uniformly random, collecting enough samples for a single one is equivalent to observing a c^{-2} -multi-collision, which requires $2^{k \cdot (c^{-2}-1)/c^{-2}} \approx 2^k$ samples, so this does not really improve on the exhaustive search of the secret. The idea behind BKW is then to artificially create samples of the above form by combining many random ones; this may provide enough samples to apply majority decoding, but each sample is now “noisier”, and one must then find a proper tradeoff.

More precisely, BKW proceeds as follows. To solve an instance of LPN with noise level η (i.e. noise correlation $c = |1 - 2\eta|$) and dimension k with q queries making a pool P_0 , we start by fixing a block-size parameter b . Then one creates a table T of size 2^b and an updated pool P_1 , both initially empty. Next, for each sample $x = (\mathbf{a}, v = \mathbf{a} \cdot \mathbf{s} + e) \in P$, do the following:

1. Call i the integer value corresponding to the last b bits of \mathbf{a} . If $T[i] = \emptyset$, update it as $T[i] \leftarrow (\mathbf{a}, v)$.
2. Else, retrieve (\mathbf{a}', v') from $T[i]$ and update x as $(\mathbf{a}, v) \leftarrow (\mathbf{a} + \mathbf{a}', v + v')$ and store it in P_1 .

At the end of this process, and provided that $q \gg 2^b$, there are $q - 2^b$ samples in P_1 which all have their masks \mathbf{a} equal to zero on their last b bits. However, the corresponding dot products indeed became noisier, the correlation c having been squared.[¶] One may now clear the table T of all its entries and start this process again, obtaining a pool P_2 of $q - 2^{b+1}$ samples with masks whose $2b$ last bits are zero and noise correlation c^4 , etc., up to a pool P_t of $q - t2^b$ samples with masks with tb zeroes and noise correlation c^{2^t} .

The original algorithm chooses t and b s.t. $tb = k - 1$, that is the last pool is made of samples with masks \mathbf{a} of the form $(0, \dots, 0)$, which are useless, and $(1, 0, \dots, 0)$, whose corresponding v values can be used to find the first bit of \mathbf{s} by a majority vote. This latter

[¶]One can check that the probability of error is $2\eta(1 - \eta) = 2(\eta - \eta^2)$, giving a correlation $1 - 4(\eta - \eta^2) = (1 - 2\eta)^2$.

step will succeed w.h.p. if $\#P_t \approx c^{-2^{t+1}}$. Finally, the remaining bits of \mathbf{s} are iteratively retrieved using the same process.

We now introduce a first optimization to the above algorithm, due to Levieil and Fouque [LF06]. Let l be a second “block size” parameter, corresponding to an exhaustive search step. The goal is to guess \mathbf{s} by blocks of l bits instead of blocks of 1; the first iteration of BKW is then stopped at $tb = k - l$. At that point, one has $N := q - t2^b$ noisy dot products with masks \mathbf{a}_i whose l first bits only may be non-zero.^{||} Let $\mathbf{y} \in \mathbb{F}_2^k$ be a vector whose $k - l$ last bits are zero; the idea is to notice that if \mathbf{y} and \mathbf{s} are equal on their first l bits, then $\mathbf{a}_i \cdot \mathbf{y} = \mathbf{a}_i \cdot \mathbf{s}$. Consequently, for every sample (\mathbf{a}_i, v_i) , one has $\mathbf{a}_i \cdot \mathbf{y} + v_i = e_i$ (where e_i is an aggregated error of some correlation c' coming from the pool creation) which is then one with non-uniform probability $(1 - c')/2$. On the other hand, if \mathbf{y} and \mathbf{s} disagree on their last bits, $\mathbf{a}_i \cdot \mathbf{y} + v_i = \mathbf{a}_i \cdot (\mathbf{y} + \mathbf{s}) + e_i$ for a non-zero term $\mathbf{y} + \mathbf{s}$, and from the uniformity of \mathbf{a}_i on its first l bits, this expression is one with probability $1/2$ exactly. If $N \approx c'^{-2}$, it is thus possible to distinguish the value of \mathbf{y} that matches \mathbf{s} to recover its last l bits.

The full process for one block of size l is as follows:

1. For all 2^l candidates \mathbf{y}_i for the first l bits of \mathbf{s} , compute $\hat{c}_{\mathbf{y}_i} = \sum_{j=0}^N (-1)^{\mathbf{a}_j \cdot \mathbf{y}_i + v_j}$ (where the sum is over \mathbb{Z}).
2. Return the \mathbf{y}_i for which $|\hat{c}_{\mathbf{y}_i}|$ is largest.

A key algorithmic observation is that, the above is similar to the computation of the Hadamard-Walsh spectrum of the Boolean function $i \mapsto v_i$, for which fast algorithms exist. One can indeed remark that for \mathbf{y} and \mathbf{y}' that differ only on a single bit, on position l (w.l.o.g.), if we write \mathbf{y}'' the common part of the two vectors of length $l - 1$ and \mathbf{a}_j^1 (resp. \mathbf{a}_j^0) the masks whose l^{th} bit is 1 (resp. 0); then we have the following; let

$$A := \sum_{\mathbf{a}_j^1} (-1)^{\mathbf{a}_j \cdot \mathbf{y} + v_j}, \quad B := \sum_{\mathbf{a}_j^0} (-1)^{\mathbf{a}_j \cdot \mathbf{y}' + v_j}, \quad -A = \sum_{\mathbf{a}_j^1} (-1)^{\mathbf{a}_j \cdot \mathbf{y}' + v_j},$$

and $\hat{c}_{\mathbf{y}} = A + B$, $\hat{c}_{\mathbf{y}'} = -A + B$, and this equality can be applied l times recursively.

The above idea can be further optimized by first applying a transformation on the LPN samples in order to reduce the problem to finding an equivalent *low weight* secret. This idea is due to Kirchner [Kir11] and works as follows. Denote \mathbf{A} the full matrix of the masks of q queries to an LPN oracle of dimension k . Let \mathbf{A}_1 be an information set of k linearly-independent columns of \mathbf{A} , and let \mathbf{A}_2 denote k columns of \mathbf{A} not in \mathbf{A}_1 . The key observation is that the sum $(\mathbf{s}\mathbf{A}_1 + \mathbf{e}_1)\mathbf{A}_1^{-1}\mathbf{A}_2 + \mathbf{s}\mathbf{A}_2 + \mathbf{e}_2$ simplifies to $\mathbf{e}_1\mathbf{A}_1^{-1}\mathbf{A}_2 + \mathbf{e}_2$. In other words, one can transform the $2k$ queries w.r.t. masks \mathbf{A}_1 and \mathbf{A}_2 for the secret \mathbf{s} into k queries for the “secret” \mathbf{e}_1 , whose expected Hamming weight is given by the noise level η , and is then strictly less than $k/2$. The same transformation can be applied many times by changing \mathbf{A}_2 , and one can then run a BKW algorithm to retrieve \mathbf{e}_1 (which obviously immediately leads to \mathbf{s}). There is no particular advantage in using this transformation when applying the original BKW algorithm, but when one guesses the secret by block, it becomes enough to guess secrets of low weight, which reduces the search space. Finally, note that the joint computation of the many correlations $\hat{c}_{\mathbf{y}}$ can still be done efficiently for those sparse secrets.

We conclude this section by mentioning that even prior the publication of the BKW algorithm, Bleichenbacher described a similar algorithm to exploit biased DSA signatures [Ble00, AFG⁺14]. Recall that in a DSA (or Schnorr) signature, the signer provides

^{||}Note that as $N \gg l$, *non-noisy* approximations would indeed allow to uniquely recover the first l bits of \mathbf{s} .

a pair $(c, r + cx)$, where x is a secret exponent, c is random and depends on the message to be signed, and r is a random mask. If r is uniform, the secret cx is blinded by a one-time-pad and nothing can be learned. But if r is biased, then one obtains a problem of noisy decoding similar to LPN. Bleichenbacher’s algorithm to retrieve x from many signatures consists in combining samples to zero some of their bits, and to apply a fast Fourier transform to recover, say, 40 bits of the secret. The original bias exploited by Bleichenbacher was a small modulo bias: instead of taking r uniformly in \mathcal{S} , $\#\mathcal{S} \approx 2^{160}$, it was uniform over $[0, \dots, 2^{160} - 1]$ and then reduced modulo $\#\mathcal{S}$. Finally, we remark that similarly to information-set decoding style algorithms, one can also recover a DSA secret using fewer sample signatures with biased masks by finding short vectors in a Euclidean lattice [HS01].

5 The Goldreich-Levin theorem

In this section, we present the application of list decoding to a proof of existence of *hard-core* predicates for one-way functions, due to Goldreich and Levin [GL89]. Informally, the objective is to show that if $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a one-way function, in the sense that it is hard to find a preimage x given $F(x)$, then it is also hard to predict (with probability significantly away from $1/2$) the value $\mathbf{a} \cdot \mathbf{x}$, for an arbitrary $\mathbf{a} \in \mathbb{F}_2^n$ (where $\mathbf{x} \in \mathbb{F}_2^n$ is the canonical embedding of $x \in \{0, 1\}^n$). A possible proof is to show that if one is given a prediction oracle for (several) $\mathbf{a} \cdot \mathbf{x}$ with correlation c , then one can reconstruct x w.h.p. with time and memory complexity c^{-2} . Thus, if “it costs at least T to invert F ”, one has that it is impossible to predict the value of an arbitrary predicate of the above form with correlation (\approx advantage) better than $1/\sqrt{T}$.**

One first subtlety that deserves to be mentioned is that because of the nature of the result we want to prove, the predicate oracle that we will use to invert F can only be called *once* for a given mask \mathbf{a} : indeed, it makes no sense to define different predictions $\mathbf{a} \cdot \mathbf{x}$ several times as \mathbf{x} itself is fixed. On the other hand, the mask \mathbf{a} on which we call the oracle can be chosen freely. This is to be contrasted with, say, an LPN setting, where one is given $\mathbf{a} \cdot \mathbf{s}$ for random (and uncontrolled) masks \mathbf{a} , that may still be potentially equal, in case of unlikely collisions. Yet in either case the effect is the same: it is (by definition or computationally) feasible to recover \mathbf{x} by recovering enough biased predictions or samples with the same n linearly-independent masks. Finally, an accurate modelisation of the oracle’s power in terms of codes is to say that if $\mathbf{A} \in \mathbb{F}_2^{n \times q}$ has q pairwise distinct columns, querying a prediction oracle with q masks \mathbf{A} and advantage ε (meaning that a prediction is correct with probability $p > 1/2$; $2p - 1 = \varepsilon$) is equivalent to obtaining a noisy codeword $\mathbf{x}\mathbf{A} + \mathbf{e}$ with $\mathbf{e} \leftarrow \text{Ber}_{\eta, q}$, $\eta = 1/2 - \varepsilon/2$.

A second essential remark is that F and $F(x)$ are both known to the adversary. Thus, as soon as one knows a “small” list L that contains x w.h.p., one can recover x uniquely by mapping F to L and comparing the result to $F(x)$ (here we assume that L does not contain collisions for F , which is true w.h.p. if it is small).

Putting the two remarks together, what we need is an efficient *list-decoding* algorithm for a code generated by a matrix \mathbf{A} with pairwise-distinct columns. In the remainder, we will give exactly such an algorithm for a punctured first-order Reed-Muller code.

Let $\mathbf{V} \in \mathbb{F}_2^{m \times (r+1)}$ be a projection matrix of rank $r + 1$, where $2^r \approx (1 - 2\eta)^{-2} = \varepsilon^{-2}$ is the number of samples required to distinguish the uniform distribution from Ber_{η} w.h.p., and whose first column is a unit vector \mathbf{b}_i (corresponding to the bit x_i that one wishes to recover); let $\mathbf{y} = \mathbf{x}\mathbf{V} \in \mathbb{F}_2^{(r+1)}$ be the projected message; let $\mathbf{W} \in \mathbb{F}_2^{(r+1) \times 2^r}$ be the matrix

**Note that this does not mean that *no* predicate could be predicted; for instance F could be one-way and still fully reveal the first bit of its input. However one is expected to be unsuccessful on arbitrary inputs, so also for instance on random ones.

whose first row is all one and whose $r \times 2^r$ lower block is made of all the vectors of \mathbb{F}_2^r ; let $\mathbf{z} = \mathbf{y}\mathbf{W}$, and remark that if \mathbf{y} is seen as a degree-1 Boolean function in r variables, then \mathbf{z} corresponds to the first-order Reed-Muller code $\text{RM}(1, r)$ encoding of \mathbf{y} . Now the problem is: given a noisy codeword $\hat{\mathbf{z}} = \mathbf{z} + \mathbf{e}$, recover x_i . If the noise level is low, one can simply use an efficient decoding algorithm for a Reed-Muller code. However, if η is close to $1/2$, the expected weight of \mathbf{e} is $2^r \eta \approx 2^{r-1}$ which corresponds to the minimum distance of $\text{RM}(1, r)$, and unique decoding is not possible anymore. An observation that may not seem extremely useful at first is that despite a high-level of noise, one can obtain a list of 2^r possible values for x_i in the following way: 1) guess the value of $\mathbf{y}[1, \dots, r]$ and build a matching vector \mathbf{y}' , with $\mathbf{y}'[0] = 0$; 2) compute $\mathbf{z}' = \hat{\mathbf{z}} + \mathbf{y}'\mathbf{W}$; 3) if $wt(\mathbf{z}') > 2^{r-1}$, guess $x_i = 1$, else guess $x_i = 0$.^{††} Now the crucial point is that this process can be in fact jointly applied to all the bits of x by using the *same* projection matrix \mathbf{V} , up to its first column.^{‡‡} By doing so, the guess for the value of the last r bits of \mathbf{y} can be reused for every bit of x to obtain a list of 2^r consistent predictions for x in its entirety. An example high-level implementation of this list decoder is given at https://github.com/P1K/LectureNotes/blob/master/CRY-M2-ADV/sw/list_decoding/simple_encoder_decoder.sage.

References

- [AFG⁺14] Diego F. Aranha, Pierre-Alain Fouque, Benoît Gérard, Jean-Gabriel Kammerer, Mehdi Tibouchi, and Jean-Christophe Zavalowicz. GLV/GLS Decomposition, Power Analysis, and Attacks on ECDSA Signatures with Single-Bit Nonce Bias. In *ASIACRYPT (1)*, volume 8873 of *Lecture Notes in Computer Science*, pages 262–281. Springer, 2014.
- [Ale03] Michael Alekhnovich. More on Average Case vs Approximation Complexity. In *FOCS*, pages 298–307. IEEE Computer Society, 2003.
- [BKW03] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, 2003.
- [BL12] Daniel J. Bernstein and Tanja Lange. Never Trust a Bunny. In *RFIDSec*, volume 7739 of *Lecture Notes in Computer Science*, pages 137–148. Springer, 2012.
- [Ble00] Daniel Bleichenbacher. On the generation of one-time keys in DL signature schemes. In *Presentation at IEEE P1363 working group meeting*, page 81, 2000.
- [BMvT78] Elwyn R. Berlekamp, Robert J. McEliece, and Henk C. A. van Tilborg. On the inherent intractability of certain coding problems (Corresp.). *IEEE Trans. Inf. Theory*, 24(3):384–386, 1978.
- [CC98] Anne Canteaut and Florent Chabaud. A New Algorithm for Finding Minimum-Weight Words in a Linear Code: Application to McEliece’s Cryptosystem and to Narrow-Sense BCH Codes of Length 511. *IEEE Trans. Inf. Theory*, 44(1):367–378, 1998.
- [GL89] Oded Goldreich and Leonid A. Levin. A Hard-Core Predicate for all One-Way Functions. In *STOC*, pages 25–32. ACM, 1989.

^{††}This entire process can be sped-up using a fast Walsh-Hadamard transform.

^{‡‡}We leave as an exercise the proof that the corresponding queries are pairwise-distinct.

- [GRS08] Henri Gilbert, Matthew J. B. Robshaw, and Yannick Seurin. How to Encrypt with the LPN Problem. In *ICALP (2)*, volume 5126 of *Lecture Notes in Computer Science*, pages 679–690. Springer, 2008.
- [HKL⁺12] Stefan Heyse, Eike Kiltz, Vadim Lyubashevsky, Christof Paar, and Krzysztof Pietrzak. Lapin: An Efficient Authentication Protocol Based on Ring-LPN. In *FSE*, volume 7549 of *Lecture Notes in Computer Science*, pages 346–365. Springer, 2012.
- [HS01] Nick Howgrave-Graham and Nigel P. Smart. Lattice Attacks on Digital Signature Schemes. *Des. Codes Cryptogr.*, 23(3):283–290, 2001.
- [Kir11] Paul Kirchner. Improved Generalized Birthday Attack. *IACR Cryptol. ePrint Arch.*, 2011:377, 2011.
- [LB88] Pil Joong Lee and Ernest F. Brickell. An Observation on the Security of McEliece’s Public-Key Cryptosystem. In *EUROCRYPT*, volume 330 of *Lecture Notes in Computer Science*, pages 275–280. Springer, 1988.
- [LF06] Éric Leveil and Pierre-Alain Fouque. An Improved LPN Algorithm. In *SCN*, volume 4116 of *Lecture Notes in Computer Science*, pages 348–359. Springer, 2006.
- [McE78] Robert J McEliece. A public-key cryptosystem based on algebraic. *DSN Progress Report*, 1978.
- [MS06] Florence Jessie MacWilliams and Neil James Alexander Sloane. *The Theory of Error-Correcting Codes*. North-Holland Mathematical Library. North-Holland, 12 edition, 2006.
- [Pra62] Eugene Prange. The use of information sets in decoding cyclic codes. *IRE Trans. Inf. Theory*, 8(5):5–9, 1962.
- [Ste88] Jacques Stern. A method for finding codewords of small weight. In *Coding Theory and Applications*, volume 388 of *Lecture Notes in Computer Science*, pages 106–113. Springer, 1988.