

A short introduction to secret sharing from linear codes

Pierre Karpman

December 9, 2020

This short lecture note gives a simple introduction to secret sharing and shows how Shamir's scheme can easily be adapted to any MDS linear code. It then presents the code-based generalisation of Massey which introduces the important concept of *access structure* and *minimal codewords*. It concludes with a short presentation of a further generalisation, *viz.* perfectly-secure message transmission protocols.

1 Preliminaries: uniform masking in a finite group

We first recall a few essential results about the distribution of the sum of two random distributions.

Let $(\mathbb{G}, +)$ be a finite group of order N ; X, Y be independent random variables over \mathbb{G} . The distribution of $X + Y$ is given by:

$$\Pr[X + Y = \alpha] = \sum_{\beta \in \mathbb{G}} \Pr[X = \beta] \cdot \Pr[Y = \alpha - \beta], \quad (1)$$

for any $\alpha \in \mathbb{G}$. This is nothing but the *discrete convolution* of X and Y .

Now if Y is uniformly distributed, (1) simplifies to:

$$\Pr[X + Y = \alpha] = \sum_{\beta \in \mathbb{G}} \Pr[X = \beta] \cdot \frac{1}{N} = \frac{1}{N}, \quad (2)$$

i.e. the distribution of $X + Y$ is uniform regardless of the distribution of X , which can be arbitrary. Still assuming that Y is uniform, we also get the important property that the *a posteriori* knowledge on the distribution of X after learning a sample of $X + Y$ is equal to its *a priori* knowledge. This is easy to show using Bayes' formula since for every $\alpha \in \mathbb{G}$ one has that:

$$\Pr[X = \alpha | X + Y = \beta] = \frac{\Pr[X + Y = \beta | X = \alpha] \cdot \Pr[X = \alpha]}{\Pr[X + Y = \beta]} = \Pr[X = \alpha]. \quad (3)$$

It is also easy to show that the above is true for all (α, β) pair iff. X or Y is uniform.

2 Basic secret sharing

Definition 1. We define a secret-sharing scheme in the following, slightly informal way:

A (k, n) secret-sharing scheme over a general set \mathcal{S} is a randomised algorithm A that takes as input $s \in \mathcal{S}$ and returns $\mathbf{s} \in \mathcal{S}^n$ s.t. the two following properties hold:

1. (*Decoding $(n - k)$ erasures.*) There is a deterministic algorithm that returns s when given any k distinct entries of \mathbf{s} .

2. (*Uniformity of $(k - 1)$ tuples.*) The distribution of any $k - 1$ distinct entries of \mathbf{s} is uniform (over the randomness of A).

Example 2. A simple (n, n) scheme can be built over a finite group $(\mathbb{G}, +)$ as follows: for $i \in \llbracket 1, n - 1 \rrbracket$ take the \mathbf{s}_i s to be uniform i.i.d. variables, and then define \mathbf{s}_n as $s - \sum_{i=1}^{n-1} \mathbf{s}_i$. It is straightforward to see that s can be uniquely reconstructed from \mathbf{s} , and that the distribution of any $(n - 1)$ tuple is uniform follows from (2).

2.1 Shamir’s scheme from polynomial interpolation

We now present an elegant secret sharing scheme due to Shamir (1979) based on univariate polynomial interpolation. It allows to build (k, n) schemes for secrets in a finite field \mathbb{F}_q as long as $n \leq O(q)$. It works as follows: let $P \in \mathbb{F}_q[X]$ be a polynomial of degree $\leq k - 1$ whose $k - 1$ non-constant coefficients $\mathbf{c} \in \mathbb{F}_q^{k-1}$ are sampled uniformly at random and whose constant coefficient is the secret s to be shared, *i.e.* $P = s + \sum_{i=1}^{k-1} \mathbf{c}_i X^i$. Then let $\mathbf{a} \in \mathbb{F}_q^n$ be a vector whose entries are all pairwise distinct and not equal to zero, and define the sharing \mathbf{s} as $\text{eval}(P, \mathbf{a})$.^{*} We now prove that this defines a (k, n) secret sharing:

Proof. We show that the scheme satisfies the two properties of a (k, n) sharing.

1. Let $\mathcal{I} \subseteq \llbracket 1, n \rrbracket$ be a subset of indices of size at least k . By Lagrange interpolation, $P' := \sum_{i \in \mathcal{I}} \mathbf{s}_i \prod_{j \in \mathcal{I} \setminus \{i\}} \frac{X - \mathbf{a}_j}{\mathbf{a}_i - \mathbf{a}_j}$ is the unique polynomial of degree $\leq k - 1$ that evaluates to \mathbf{s}_i on \mathbf{a}_i for $i \in \mathcal{I}$. It is thus equal to P , and one recovers s from $\text{eval}(P', 0)$.
2. Let $\mathcal{I} \subseteq \llbracket 1, n \rrbracket$ be a subset of indices of size $k - 1$ and denote by $\mathbf{s}_{\mathcal{I}}$ the vector of corresponding entries of \mathbf{s} . Then for fixed s and \mathbf{a} , the map $\mathbf{c} \mapsto \mathbf{s}_{\mathcal{I}}$ is bijective:
 - Given s and \mathbf{c} , $\mathbf{s}_{\mathcal{I}}$ is obtained from the evaluation of the fully-determined polynomial P on $\mathbf{a}_{\mathcal{I}}$.
 - ← Since $0 \notin \mathbf{a}$, one knows the evaluation of P on k distinct points $\{0\} \cup \mathbf{a}_{\mathcal{I}}$, which allows to uniquely interpolate P and then obtain \mathbf{c} .

Since this map is bijective and the distribution of \mathbf{c} is uniform and independent from s , then so is the one of $\mathbf{s}_{\mathcal{I}}$. □

2.2 Generalisation to arbitrary MDS linear codes

A (not so) close inspection of Shamir’s scheme reveals that it essentially relies on Reed-Solomon codes, which are specific instances of *maximum distance separable* (MDS) linear codes. We will now see how this scheme can be (somewhat) generalised and formulated entirely within the framework of such codes, using a purely linear-algebraic language. This is arguably the “correct” view of secret sharing, in particular since it later leads to natural and powerful generalisations such as sketched in the next section.

We first define linear and MDS codes and one of their characterisation that is essential to our purpose.

Definition 3 (Linear code). A *linear code* \mathcal{C} over \mathbb{F}_q of *length* n and *dimension* k (written $[n, k]_{\mathbb{F}_q}$) is a k -dimensional subspace of \mathbb{F}_q^n . Its *minimum distance* (or equivalently *minimum weight*) is defined as $\min_{\mathbf{x}, \mathbf{y} \neq \mathbf{x} \in \mathcal{C}} \text{wt}(\mathbf{x} - \mathbf{y})$, where $\text{wt}(\cdot)$ is the Hamming-weight function that counts how many entries of its input are non-zero.

^{*}We do not consider \mathbf{a} to be part of the sharing strictly speaking since it may be fixed and publicly known, yet its knowledge is essential to a successful recovery of the secret.

Every linear code can be *generated* by a matrix $\mathbf{G} \in \mathbb{F}_q^{k \times n}$ in the following sense: $\forall \mathbf{x} \in \mathcal{C}$, $\exists! \mathbf{m} \in \mathbb{F}_q^k$ s.t. $\mathbf{x} = \mathbf{m}\mathbf{G}$. Note that \mathbf{G} is in general not unique.

Definition 4 (Maximum distance separable code). An $[n, k]$ linear code is said to be *maximum distance separable* (MDS) iff. its minimum distance equals $n - k + 1$.

We now prove two fundamental results about MDS linear codes.

Theorem 5 (Existence). *For any $n, k \leq n \in \mathbb{N}$, $\exists q$ s.t. an $[n, k]_{\mathbb{F}_q}$ MDS linear code exists.*

Proof. We will prove this statement constructively by defining a family of MDS codes (*viz.* the Reed-Solomon codes).

Let $\mathbb{F}_q[X]_{<k}$ denote the k -dimensional vector space of univariate polynomials over \mathbb{F}_q of degree $< k$, and $\mathbf{a} \in \mathbb{F}_q^n$ be a vector of $k \leq n \leq q$ pairwise-distinct elements of \mathbb{F}_q , then the map $\mathbb{F}_q^k \rightarrow \mathbb{F}_q^n$, $P \in \mathbb{F}_q[X]_{<k} \mapsto \text{eval}(P, \mathbf{a})$ is linear. Since P has at most $k - 1$ distinct roots, any input that is not the zero polynomial has an image of weight at least $n - k + 1$, and this map thus defines an $[n, k, n - k + 1]$ MDS linear code. \square

Theorem 6 (Optimality for erasure decoding). *An $[n, k]$ linear code, is MDS iff. all its generator matrices are s.t. every k of their columns are linearly independent.*

Proof. \Rightarrow Let $\mathbf{G} \in \mathbb{F}_q^{k \times n}$ be a generator matrix of an $[n, k]$ MDS code, and suppose that there is a subset $\mathcal{I} \subseteq \llbracket 1, n \rrbracket$ of k indices such that the columns indexed by \mathcal{I} are not linearly independent. Equivalently this means that the submatrix $\mathbf{G}_{\mathcal{I}}$ has rank $< k$; thus its left kernel has dimension ≥ 1 and $\exists \mathbf{m} \neq \mathbf{0} \in \mathbb{F}_q^k$ s.t. $\mathbf{m}\mathbf{G}_{\mathcal{I}} = \mathbf{0}$, and thusly a non-zero codeword $\mathbf{m}\mathbf{G}$ of weight at most $n - k < n - k + 1$, which is a contradiction with the fact that the code generated by \mathbf{G} is MDS.

\Leftarrow Let $\mathbf{G} \in \mathbb{F}_q^{k \times n}$ be as required. Then given $\mathbf{x} := \mathbf{m}\mathbf{G}$, one can recover \mathbf{m} from the restriction $\mathbf{x}_{\mathcal{I}}$ of \mathbf{x} to any $\mathcal{I} \subseteq \llbracket 1, n \rrbracket$ of size k by inverting the full-rank submatrix $\mathbf{G}_{\mathcal{I}}$. Consequently, the encodings of two distinct messages cannot be equal in more than $k - 1$ positions and the minimum distance of the code generated by \mathbf{G} is thus at least (in fact exactly) $n - k + 1$. \square

We now show how to define a (k, n) secret sharing scheme in \mathbb{F}_q from an $[n + 1, k]$ MDS linear code over the same field. Let \mathbf{G} be a generator matrix for such a code, and reduce it in its first row to obtain the block matrix $(\mathbf{e}_1 \ \mathbf{G}')$, where $\mathbf{e}_1 = (1 \ 0 \ \dots \ 0)^t$ is the first vector of the canonical basis and (by Thm. (6)) $\mathbf{G}' \in \mathbb{F}_q^{k \times n}$ generates an $[n, k]$ MDS code. Then a sharing of s is obtained by sampling \mathbf{c} uniformly at random in \mathbb{F}_q^{k-1} , setting $\mathbf{m} := (s \ \mathbf{c})$, and computing \mathbf{s} as $\mathbf{m}\mathbf{G}'$ (or equivalently \mathbf{s} is formed from the last n coordinates of $\mathbf{m}\mathbf{G}$ *i.e.* it is a codeword obtained after *puncturing* the first coordinate). We now prove that this defines a (k, n) secret sharing:

Proof. We proceed as previously.

1. Immediate from Thm. (6) and the fact that \mathbf{G}' is MDS.
2. Let $\mathcal{I} \subseteq \llbracket 1, n \rrbracket$ be a subset of indices of size $k - 1$, then for a fixed s and \mathbf{G}' , the map $\mathbf{c} \mapsto \mathbf{s}_{\mathcal{I}}$ is bijective:
 - \rightarrow Given s and \mathbf{c} , one can easily compute $\mathbf{m}\mathbf{G}'_{\mathcal{I}} = \mathbf{s}_{\mathcal{I}}$.
 - \leftarrow By construction and the MDS property of the codes generated by \mathbf{G} and \mathbf{G}' , $\mathbf{G}'_{\mathcal{I}}$ is of rank $k - 1$ and \mathbf{e}_1 is not in its span, hence $\mathbf{R} := (\mathbf{e}_1 \ \mathbf{G}'_{\mathcal{I}})$ is invertible, and \mathbf{c} can be computed from $(s \ \mathbf{s}_{\mathcal{I}}) \cdot \mathbf{R}^{-1} = \mathbf{m}$.

Since this map is bijective and the distribution of \mathbf{c} is uniform and independent from s , then so is the one of $\mathbf{s}_{\mathcal{I}}$.

□

Finally, let us remark that there is nothing special about the choice of \mathbf{e}_1 in the above construction of \mathbf{G}' from \mathbf{G} , and that any non-null vector would be similarly suitable to embed the secret as the corresponding linear combination of elements of \mathbf{m} . For instance one can check that using instead the all-1 vector and setting \mathbf{m} to the $(k-1, k)$ additive sharing from Ex. (2) also results in a (k, n) sharing.

3 Access structures and minimal codewords

A natural generalisation of secret sharing as described above given by Massey (1993) is to make it possible to recover the secret from subsets of shares of different sizes. More generally, we would ideally be able to have a fine-grained control over which specific subsets allow to do so. For instance we may want a scheme s.t. $\mathbf{s}_{\{1,2\}}$ and $\mathbf{s}_{\{2,3,4\}}$ allow to recover the secret \mathbf{s}_1 , with the constraint that only sets with one of $\{1, 2\}$ or $\{2, 3, 4\}$ as a subset allow to do so. We will see how to relate such a requirement to the structure of a linear code used to perform the sharing; however to achieve such an imbalance between the different shares' roles the codes will not be MDS any more.

In all of the following we will drop the k parameter used previously in the definition of (k, n) secret sharing schemes, since the knowledge of which subsets of shares determine the secret will be *a priori* independent from the size of the subset. We replace this with the notion of *access structure* defined below.

3.1 Access structures

We start with the following:

Definition 7 (Access structure). The *access structure* for a secret-sharing scheme with n shares over a general set \mathcal{S} is a subset \mathcal{A} of $\wp(\mathcal{S})$ s.t. for all sharings $\mathbf{s} \in \mathcal{S}^n$ of an arbitrary secret s one has that:

1. There is a deterministic algorithm that returns s on input of any $\mathcal{X} \in \mathcal{A}$.
2. The distribution of the elements of any $\mathcal{X} \in \wp(\mathcal{S})$ for which $\nexists \mathcal{X}' \in \mathcal{A}$ s.t. $\mathcal{X}' \subseteq \mathcal{X}$ is uniform over the randomness used to compute the sharing.

In words, this means that the access structure contains all the “minimal” subsets of shares that allow to recover the secret: any of these is enough to do so, and any subset that does not include at least one of them is statistically independent from the secret (and thence does not allow to learn anything about it). It is however not obvious that this definition is not contradictory, *i.e.* that a secret sharing scheme does possess an access structure. In the following we will show that this is always the case for schemes built over linear codes.

3.2 Dual characterisation

We first show how to use the *dual* of the code used to build the secret sharing to determine if a set of shares allows to recover the secret, and address the minimality condition next. We consider as in the previous section (and w.l.o.g.) a sharing obtained as the n last coordinates of $(s \ \mathbf{c}) \mathbf{G}$ for a uniform \mathbf{c} and \mathbf{G} a generator matrix of an $[n+1, k]$ linear code \mathcal{C} whose first column is equal to \mathbf{e}_1 .

Let us give the following:

Definition 8 (Dual code). The *dual* C^\perp of an $[n, k]$ linear code C is the $[n, n - k]$ code s.t. $\forall \mathbf{x} \in C, \forall \mathbf{y} \in C^\perp, \mathbf{x} \cdot \mathbf{y} = 0$. Equivalently, it is the code generated by a parity-check matrix for C .

And then:

Lemma 9. Consider an n -share secret sharing scheme defined from a linear code C as above; a subset of shares of indices in \mathcal{X} allows to recover the secret iff. $\exists \mathbf{x} \in C^\perp$ of support $\{1\} \cup \mathcal{X}'$, $\mathcal{X}' \subseteq \mathcal{X}$. Furthermore, in the negative case the subset of shares is statistically independent from the secret.

Proof. \Rightarrow Let $\mathbf{x} \in C^\perp$ be as required. By definition of the dual code this means that for every sharing \mathbf{s} of s , since $(s \ \mathbf{s}) \in C$ one has $(s \ \mathbf{s}) \cdot \mathbf{x} = 0$. Let us write \mathbf{x} as $(\xi \ \mathbf{x}')$ where $\xi \neq 0$, then we have $s\xi = -\mathbf{s} \cdot \mathbf{x}'$ which means that s can be uniquely recovered as an appropriate linear combination of the known shares \mathcal{X} .

\Leftarrow By the assumption that there is no \mathbf{x} satisfying the requirement we have that the first column \mathbf{G}_1 of \mathbf{G} is not in the column space of the submatrix $\mathbf{G}_{\mathcal{X}}$ with columns in \mathcal{X} . Thus the shares $\mathbf{s}_{\mathcal{X}}$ in which s appears write as $s + \boldsymbol{\lambda} \cdot \mathbf{c}$ for some $\boldsymbol{\lambda} \neq \mathbf{0}$. Let $t \leq \#\mathcal{X}$ be the number of such linearly-independent shares and $\mathbf{\Lambda} := (\boldsymbol{\lambda}_{(1)} \ \cdots \ \boldsymbol{\lambda}_{(t)})$ be the corresponding full-rank matrix, then from the change of variable $\mathbf{c}'_i := \boldsymbol{\lambda}_{(i)} \cdot \mathbf{c}$ the same (linearly-independent) shares can be rewritten as $s + \mathbf{c}'_i$ and the uniformity of \mathbf{c} (which carries over to \mathbf{c}' from the bijectivity of the change of variable) allows to conclude that they all are uniformly-distributed random variables independent from s . □

3.3 Minimal codewords

The previous lemma can be used to determine if a subset of shares allows to recover the secret or if it is statistically independent from it; to further determine if it is an access structure (and to also justify the well-foundedness of this definition) we will need the additional concept of *minimal* codewords. In all of the following we use the notation $\mathbf{x} \preceq \mathbf{y}$ (resp. $\mathbf{x} \prec \mathbf{y}$) to mean that $\text{supp}(\mathbf{x}) \subseteq \text{supp}(\mathbf{y})$ (resp. $\text{supp}(\mathbf{x}) \subset \text{supp}(\mathbf{y})$).

Definition 10 (Minimal codeword). A *minimal codeword* of a linear code C is a codeword \mathbf{x} with its first non-zero coordinate equal to 1 s.t.: 1) there is no distinct $\mathbf{x}' \in C$ with its first non-zero coordinate equal to 1 and $\mathbf{x}' \preceq \mathbf{x}$; 2) or equivalently there is no non-zero $\mathbf{x}' \in C$ s.t. $\mathbf{x}' \prec \mathbf{x}$.

Proof. (1) \Rightarrow (2) $\exists \lambda$ s.t. $(\mathbf{x} - \lambda \mathbf{x}') \prec \mathbf{x}$.

(2) \Rightarrow (1) $\exists \lambda$ s.t. the first coordinate of $\lambda \mathbf{x}'$ equals 1. □

Using a similar argument one shows that no two distinct minimal codewords may have the same support.

Massey then proves the following:

Lemma 11. Every codeword \mathbf{x} of C is a linear combination of minimal codewords $\{\mathbf{x}_i\}$ s.t. for all i one has $\mathbf{x}_i \preceq \mathbf{x}$.

Proof. Let $\mathbf{x} \in C$ be a non-zero codeword. If it is minimal (up to scaling by a scalar) then we are done; otherwise there must exist a non-zero codeword $\mathbf{x}_1 \prec \mathbf{x}$ which by induction may be assumed to be minimal. Thus $\exists \lambda$ s.t. $(\mathbf{x} - \lambda \mathbf{x}_1) \prec \mathbf{x}$ and the first part of the result follows by induction; the second part immediately follows from the base case and the two inclusions $\mathbf{x}_1 \prec \mathbf{x}$ and $(\mathbf{x} - \lambda \mathbf{x}_1) \prec \mathbf{x}$. □

This leads to the following:

Corollary 12. *For every codeword $\mathbf{x} \in \mathcal{C}$ that is not minimal (up to rescaling), there is a minimal codeword $\mathbf{x}' \prec \mathbf{x}$ whose first non-zero coordinate is at the same index as the first non-zero coordinate of \mathbf{x} .*

This finally leads to the main result we want to show:

Theorem 13. *The access structure of a secret sharing scheme built from a code \mathcal{C} is the set of the supports of the minimal codewords of \mathcal{C}^\perp whose first coordinate is non-zero (with 1 then excluded from these sets).*

Proof. The fact that subsets of shares of this form allow to reconstruct the secret is a consequence of Lem. (9); it remains to show that no subset that does not include one of these allows to do so.

Again from Lem. (9) we have that if \mathcal{X} allows to reconstruct the secret then there must exist $\mathbf{x} \in \mathcal{C}^\perp$ with support included in $\{1\} \cup \mathcal{X}$ and whose first coordinate is non-zero. Now either \mathbf{x} is minimal (up to rescaling) and the previous inclusion is an equality and we are done, or by Cor. (12) there is a minimal codeword $\mathbf{x}' \prec \mathbf{x}$ whose first coordinate is a 1; from Lem. (9) $\text{supp}(\mathbf{x}') \subseteq \mathcal{X}$ still allows to reconstruct the secret and we can conclude. \square

This result is useful in that it justifies the fine-grained notion of access structure for linear secret sharing schemes and gives a conceptually simple characterisation thereof. In turn one may hope to use this characterisation to build secret sharing schemes with specific access structures as soon as this one can be mirrored by the minimal codewords of some linear code.

4 Perfectly-secure message transmission

We now briefly introduce *perfectly-secure message transmission* (or *PSMT*) as a generalisation of secret sharing as presented above. The generalisation is in two ways: 1) we may now consider interactive protocols where several participants may exchange information in several “rounds” of communication; 2) the adversary (somewhat implicit in the above secret-sharing scenario) is now *active*, so that it may tamper with exchanged messages.

In short the goal of PSMT is for two (or more) participants to exchange a secret in a way such that the adversary does not learn anything about it. We wish to do this w.r.t. a very strong security objective, *viz.* that whatever the adversary may learn is statistically independent from the secret. To achieve this goal, the participants have access to a certain number of n channels with the guarantee (or in fact, the assumption) that the adversary controls at most t of them.

Remark 14. If in the above scenario the adversary is further assumed to be passive, the problem is immediately solved in all the relevant cases $t < n$ in only one round of communication from the existence of (n, n) secret sharing.

This remark justifies the fact that we only need to further develop techniques against active adversaries.

The goal of this section is to present an elegant scheme from Spini and Zémor (2016) for two participants \mathfrak{A} and \mathfrak{B} , in its simple and unoptimised form; it is a two-round protocol in the setting $n = 2t + 1$, which for that number of rounds is minimal (Dolev *et al.*, 1993).

A first simple observation is that since in the chosen setting $n > 2t$, reliable communication between \mathfrak{A} and \mathfrak{B} is trivial to achieve as it is enough to broadcast any message over the n channels and decode with a majority vote. It is however equally obvious that this method does not provide any confidentiality since the adversary may read any message from one of the channels it controls.

The idea is then for \mathfrak{B} to transmit a number of independent codewords from an $[n, t + 1, t + 1]_{\mathbb{F}_q}$ MDS code, where each coordinate $i \in \llbracket 1, n \rrbracket$ is sent over the corresponding channel. Note that from the previous sections these codewords all naturally induce $(t+1, n)$ secret-sharing, and \mathfrak{A} may pick one codeword to recover the secret and use the latter as a one-time-pad to reliably *and* confidentially broadcast a message. However because of adversarial tampering and since $t > (d-1)/2 = t/2$, \mathfrak{A} will in general not be able to recover the original codeword and the derived secret it uses may thus be “noisy”. The crux of the method is then to provide \mathfrak{B} with sufficient information so that it may itself remove the noise from the secret and finally recover the message, without revealing anything to the adversary.

We now describe the techniques used by Spini and Zémor to resolve this last point.

4.1 Decoding from a syndrome-spanning subset

We first recall the following:

Definition 15 (Parity-check matrix). Let \mathcal{C} be an $[n, k]_{\mathbb{F}_q}$. A *parity-check matrix* for \mathcal{C} is any full-rank matrix $\mathbf{H} \in \mathbb{F}_q^{(n-k) \times n}$ s.t. for any generator matrix $\mathbf{G} \in \mathbb{F}_q^{k \times n}$ of \mathcal{C} one has $\mathbf{GH}^t = \mathbf{0}^{k \times (n-k)}$.

Definition 16 (Syndrome of a vector). Let \mathcal{C} be an $[n, k]_{\mathbb{F}_q}$ and \mathbf{H} one of its parity-check matrix. The *syndrome* (usually written \mathbf{s}) of a vector $\mathbf{x} \in \mathbb{F}_q^n$ w.r.t. \mathbf{H} is defined as $\mathbf{H}\mathbf{x}^t$.

Note that from the definition of a parity-check matrix, $\mathbf{H}\mathbf{x}^t = \mathbf{0}^{n-k}$ iff. $\mathbf{x} \in \mathcal{C}$. In particular if one has $\mathbf{y} := \mathbf{x} + \mathbf{e}$ for some $\mathbf{e} \in \mathbb{F}_q^n$ and some $\mathbf{x} \in \mathcal{C}$, then $\mathbf{H}\mathbf{y}^t = \mathbf{H}\mathbf{e}^t$.

We now wish to show that \mathfrak{B} is able to recover an error \mathbf{e}_1 from its syndrome $\mathbf{H}\mathbf{e}_1^t$ as long as the former belongs to a suitable subspace for which \mathfrak{B} possesses a basis (or “spanning subset”) and its image through the syndrome map. We will then show that this can be ensured in the above protocol, which will allow us to conclude.

We start with:

Lemma 17. *Let \mathcal{C} be an $[n, k, d]_{\mathbb{F}_q}$ linear code for which \mathbf{H} is a parity-check matrix. Let \mathcal{E} be a linear subspace of \mathbb{F}_q^n whose elements all have weight $d - 1$ or less. Then the map $\sigma_{\mathcal{E}} : \mathcal{E} \rightarrow \mathbb{F}_q^{n-k}$, $\mathbf{e} \mapsto \mathbf{H}\mathbf{e}^t$ is injective.*

Proof. By definition the (right) kernel of \mathbf{H} is \mathcal{C} . From $\mathcal{C} \cap \mathcal{E} = \{\mathbf{0}^n\}$ the kernel of $\sigma_{\mathcal{E}}$ is trivial and the map is thence injective. \square

Now since $\sigma_{\mathcal{E}}$ is injective and linear it must admit an efficiently-computable inverse $\sigma_{\mathcal{E}}^{-1} : \text{Im}(\sigma_{\mathcal{E}}) \subset \mathbb{F}_q^{n-k} \rightarrow \mathcal{E}$ which is also linear. We formalise this as:

Proposition 18. *Let \mathcal{C} , \mathbf{H} , \mathcal{E} , $\sigma_{\mathcal{E}}$ be as above and $t := \dim(\mathcal{E})$. Let $\mathbf{B}_{\mathcal{E}} := (\mathbf{e}_1 \ \cdots \ \mathbf{e}_t)$ be a basis for \mathcal{E} (where the blocks are here row-wise; note also that in this case \mathbf{e}_i may not necessarily be equal to the i^{th} vector of the canonical basis of the ambient space) and $\mathbf{B}_{\mathcal{S}} := (\mathbf{s}_1 \ \cdots \ \mathbf{s}_t) = \mathbf{H}\mathbf{B}_{\mathcal{E}}^t$. Then one may efficiently compute $\sigma_{\mathcal{E}}$ and $\sigma_{\mathcal{E}}^{-1}$ from the knowledge of $\mathbf{B}_{\mathcal{E}}$ and $\mathbf{B}_{\mathcal{S}}$.*

Proof. This follows immediately from the definitions and by linearity:

$$\begin{aligned} - \sigma_{\mathcal{E}}(\mathbf{e} := \sum_{i=1}^t \lambda_i \mathbf{e}_i) &= \lambda_i \sum_{i=1}^t \sigma_{\mathcal{E}}(\mathbf{e}_i) = \lambda_i \sum_{i=1}^t \mathbf{s}_i. \\ - \sigma_{\mathcal{E}}^{-1}(\mathbf{s} := \sum_{i=1}^t \lambda_i \mathbf{s}_i) &= \lambda_i \sum_{i=1}^t \sigma_{\mathcal{E}}^{-1}(\mathbf{s}_i) = \lambda_i \sum_{i=1}^t \mathbf{e}_i. \end{aligned}$$

\square

4.2 A perfectly-secure message transmission protocol

We are now ready to fully describe and analyse the two-round PSMT protocol from Spini and Zémor, in its simple form.

Using all notations as above, one does the following:

0. In order to transmit one message defined over \mathbb{F}_q , \mathfrak{A} , \mathfrak{B} and the adversary agree on an $[n, t + 1, t + 1]_{\mathbb{F}_q}$ code \mathcal{C} and on a parity-check matrix \mathbf{H} . They also equip \mathcal{C} with a $(t + 1, n)$ secret-sharing scheme; from the previous sections and w.l.o.g. we assume that the secret is recovered from a codeword \mathbf{x} as $\mathbf{x} \cdot \mathbf{h}$ where $\mathbf{h} \in \mathbb{F}_q^n$ is a suitable fixed, publicly-known vector.
1. \mathfrak{B} selects $t + 1$ codewords $\{\mathbf{x}_i\}_{1 \leq i \leq t+1}$ independently and uniformly at random and sends them to \mathfrak{A} using the n channels that they share together. Specifically, the i^{th} channel is used to transmit the i^{th} coordinate of all the $t + 1$ codewords.
2. These codewords are received by \mathfrak{A} as $\{\mathbf{y}_i := \mathbf{x}_i + \mathbf{e}_i\}_{1 \leq i \leq t+1}$, where the \mathbf{e}_i 's are error vectors of \mathbb{F}_q^n of weight less than t introduced by the adversary. \mathfrak{A} computes $\{\mathbf{s}_i := \mathbf{H}\mathbf{y}_i^t\}_{1 \leq i \leq t+1}$ and picks i s.t. $\mathbf{s}_i \in \langle \mathbf{s}_j \rangle_{1 \leq j \neq i \leq t+1}$;[†] in the following we assume w.l.o.g. that $i = 1$. It then computes $s := \mathbf{y}_1 \cdot \mathbf{h}$, encrypts the message $m \in \mathbb{F}_q$ that it wants to send as $c := m + s$ and broadcasts it over the n channels. Finally it further broadcasts $\{\mathbf{s}_i\}_{1 \leq i \leq t+1}$ and $\{\mathbf{y}_i\}_{2 \leq i \leq t+1}$.
3. \mathfrak{B} computes $\{\mathbf{e}_i\}_{2 \leq i \leq t+1}$ from $\{\mathbf{y}_i\}_{2 \leq i \leq t+1}$ and its knowledge of $\{\mathbf{x}_i\}_{2 \leq i \leq t+1}$. Since the errors all have weight $t < t + 1$, knowing $\{\mathbf{e}_i\}_{2 \leq i \leq t+1}$ and $\{\mathbf{s}_i\}_{2 \leq i \leq t+1}$ is then enough to recover \mathbf{e}_1 from $\mathbf{s}_1 \in \langle \mathbf{s}_i \rangle_{2 \leq i \leq t+1}$ by Prop. 18. From then it can easily recompute s as $(\mathbf{x}_1 + \mathbf{e}_1) \cdot \mathbf{h}$ and decrypt c .

The correctness of the protocol follows from its description and the previous results, and we conclude with a short, somewhat informal security analysis. We only need to show that the distribution of s is uniform and independent from the data known to the adversary at the end of a run. This follows immediately from the remarks that: 1) the broadcast values $\{\mathbf{s}_i\}_{1 \leq i \leq t+1}$ are already known by the adversary since for all i one has $\mathbf{s}_i = \mathbf{H}\mathbf{y}_i^t = \mathbf{H}\mathbf{e}_i^t$ and the \mathbf{e}_i 's are error terms that it itself introduces; 2) the non-revealed noisy codeword \mathbf{y}_1 is independent from the revealed ones $\{\mathbf{y}_i\}_{2 \leq i \leq t+1}$; 3) $s = \mathbf{y}_1 \cdot \mathbf{h} = (\mathbf{x}_1 + \mathbf{e}_1) \cdot \mathbf{h} = (\mathbf{x}_1 \cdot \mathbf{h}) + (\mathbf{e}_1 \cdot \mathbf{h})$, where $(\mathbf{e}_1 \cdot \mathbf{h})$ is known by the adversary, and $\mathbf{x}_1 \cdot \mathbf{h}$ uniform and statistically independent from the (at most) t shares of the $(t + 1, n)$ sharing that it controls. \square

[†]This is always possible since the errors span a subspace of dimension at most t .