Electronic Colloquium on Computational Complexity, Report No. 146 (2012)

List decoding Reed-Solomon, Algebraic-Geometric, and Gabidulin subcodes up to the Singleton bound

Venkatesan Guruswami[*] Chaoping Xing[†]

October 2012

Abstract

We consider Reed-Solomon (RS) codes whose evaluation points belong to a subfield, and give a linear-algebraic list decoding algorithm that can correct a fraction of errors approaching the code distance, while pinning down the candidate messages to a well-structured affine space of dimension a constant factor smaller than the code dimension. By pre-coding the message polynomials into a subspace-evasive set, we get a Monte Carlo construction of a subcode of Reed-Solomon codes that can be list decoded from a fraction $(1 - R - \varepsilon)$ of errors in polynomial time (for any fixed $\varepsilon > 0$) with a list size of $O(1/\varepsilon)$. Our methods extend to algebraic-geometric (AG) codes, leading to a similar claim over constant-sized alphabets. This matches parameters of recent results based on folded variants of RS and AG codes, but our construction here gives subcodes of *Reed-Solomon and AG codes themselves* (albeit with restrictions on the evaluation points).

Further, the underlying algebraic idea also extends nicely to Gabidulin's construction of rank-metric codes based on linearized polynomials. This gives the ***first*** construction of positive rate rank-metric codes list decodable beyond half the distance, and in fact gives codes of rate R list decodable up to the optimal $(1 - R - \varepsilon)$ fraction of rank errors. A similar claim holds for the closely related subspace codes studied by Koetter and Kschischang.

We introduce a new notion called *subspace designs* as another way to pre-code messages and prune the subspace of candidate solutions. Using these, we also get a *deterministic* construction of a polynomial time list decodable subcode of RS codes. By using a cascade of several subspace designs, we extend our approach to AG codes, which gives the ***first deterministic*** construction of an algebraic code family of rate R with efficient list decoding from $1 - R - \varepsilon$ fraction of errors over an *alphabet of constant size* (that depends only on ε). The list size bound is almost a constant (governed by $\log^*$ (block length)), and the code can be constructed in quasi-polynomial time. Finding more efficient constructions of subspace designs is an interesting problem in pseudorandomness arising out of our work.

[*] Computer Science Department, Carnegie Mellon University, Pittsburgh, USA. Email: ⟨guruswami@cmu.edu⟩. Research supported in part by a Packard Fellowship and NSF CCF-0963975. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

[†] Division of Mathematical Sciences, School of Physical & Mathematical Sciences, Nanyang Technological University, Singapore. Email:⟨xingcp@ntu.edu.sg⟩

ISSN 1433-8092

Contents

1 Introduction

Reed-Solomon codes are a classical and widely used family of algebraic error-correcting codes. An $[n, k]$ Reed-Solomon (RS) code over a field $\mathbb{F}$ encodes polynomials $f \in \mathbb{F}[X]$ of degree $< k$ by their evaluations at a sequence $\alpha_1, \alpha_2, \ldots, \alpha_n$ of $n \leqslant |\mathbb{F}|$ distinct field elements. The rate R of this code equals k/n, and any two distinct codewords differ on more than $(1-R)n$ positions. Thus, a codeword can be unambiguously identified when up to a fraction $(1-R)/2$ of its symbols are corrupted. Polynomial time algorithms dating back to 1960 are known to correct a fraction $(1-R)/2$ of errors and recover the correct codeword [32]. We stress that in this work we focus on *worst-case symbol errors*, and the error fraction counts the proportion of symbols of the received word which differ from the corresponding codeword symbol.

When the error fraction ρ exceeds $(1-R)/2$, unique recovery of the correct codeword may not be possible, but one can hope to *list decode* a small set of codewords that includes all codewords within distance ρn from the noisy input string. In fact, such a list decoding task can be accomplished in polynomial time for Reed-Solomon codes for ρ as large as $1-\sqrt{R}$ [39, 18]. This remains the best known bound on list decodable error-fraction for RS codes. The $1-\sqrt{R}$ bound is best possible (in the sense that a larger noise level might necessitate super-polynomial list size) in some more general settings like list recovery [16], but for list decoding itself the only known limit is the trivial bound of $1-R$.

Recently, variants of RS codes, such as folded Reed-Solomon codes and derivative codes [17, 19, 24], have been used to decode up to a fraction of errors approaching $1-R$. The $1-R$ bound is the Singleton bound on relative distance of codes, and information-theoretically optimal for error-correction as one cannot hope to correct an error fraction larger than the proportion of redundant symbols in the codeword.

In this work, we show that certain *subcodes* of RS codes where the evaluation points $\alpha_1, \ldots, \alpha_n$ belong to a *subfield* of $\mathbb{F}$ can also be decoded up to the $1-R$ radius. In fact, we show that these RS codes *themselves* can be list decoded up to radius $(1-R-\varepsilon)$, pinning down the candidate messages to a subspace of dimension εk.[1] The *list size*, i.e., bound on number of codewords that might be output, is $\approx |\mathbb{F}|^{\varepsilon k}$; though exponential, note that it is non-trivially smaller than the total number $|\mathbb{F}|^k$ of possible messages.

To bring down the list size and decoding complexity, we use a subcode of the RS code that only encodes polynomials whose coefficients are restricted to belong to a carefully chosen subset that is (a) large (so we don't lose much in rate) and (b) *subspace-evasive.* Specifically, we ensure that this subset has a small intersection with every subspace of the sort output by the list decoder, and further allows for polynomial time computation of this intersection. (Note that testing all $|\mathbb{F}|^{\varepsilon k}$ candidates in the subspace for membership in the subspace-evasive subset would take too long.)

An explicit construction of a large subspace-evasive set in $\mathbb{F}^k$ that intersects every d-dimensional subspace in $d^{O(d)}$ points was given by Dvir and Lovett [6]. The intersection size was improved to $2^{O(d)}$ at the expense of worse construction complexity [1]. In our applications, the subspaces we need to evade have $\Omega(k)$ dimension, so we cannot afford an intersection size that is exponential in the dimension. We instead exploit the structural properties of the subspaces we encounter in list decoding to construct subsets with much smaller intersection. We do this in two ways: (i) using hierarchically subspace-evasive (h.s.e) as in our previous work [20][2]; this in fact achieves $O_\varepsilon(1)$ intersection size, but we only know a randomized construction, and (ii) using *subspace designs*, a notion apparently new to this work, which we can construct deterministically, and which ensures that the intersection is itself a subspace of (nearly) $O_\varepsilon(1)$ dimension. Further details on our techniques, both algebraic and pseudorandomness related, are discussed in Section 2.

[1] To be accurate, it will be a subspace over the subfield $\mathbb{K}$ of $\mathbb{F}$ of dimension $\varepsilon k[\mathbb{F} : \mathbb{K}]$.

[2] Actually, we use a combination of h.s.e sets with the Dvir-Lovett construction; see Section 2 for details.

Code	Construction	Alphabe size	List size	Decoding time	Reference
Folded RS/derivative	Explicit	$n^{O(1/\varepsilon^2)}$	$n^{O(1/\varepsilon)}$	$n^{O(1/\varepsilon)}$	[17, 19], [24]
Folded RS subcode	Monte Carlo	$n^{O(1/\varepsilon^2)}$	$O(1/\varepsilon)$	$n^{O(1/\varepsilon)}$	[19]
Folded RS subcode	Explicit	$n^{O(1/\varepsilon^2)}$	$(1/\varepsilon)^{O(1/\varepsilon)}$	$n^{O(1)}2^{1/\varepsilon^{O(1)}}$	[6]
Folded cyclotomic	Las Vegas	$(\log n)^{O(1/\varepsilon^2)}$	$n^{O(1/\varepsilon^2)}$	$n^{O(1/\varepsilon^2)}$	[13]
Folded AG subcode	Monte Carlo	$\exp(\tilde{O}(1/\varepsilon^2))$	$O(1/\varepsilon)$	$n^{O(1)}2^{1/\varepsilon^{O(1)}}$	[20]
RS subcode	Monte Carlo	$n^{O(1/\varepsilon^2)}$	$O(1/\varepsilon)$	$n^{O(1)}2^{1/\varepsilon^{O(1)}}$	Thm. 1.1
RS subcode	Explicit	$n^{O(1/\varepsilon^2)}$	$n^{O(1/\varepsilon^2)}$	$n^{O(1/\varepsilon^2)}$	Thm 1.1
AG subcode	Monte Carlo	$\exp(\tilde{O}(1/\varepsilon^2))$	$O(1/\varepsilon)$	$n^{O(1)}2^{1/\varepsilon^{O(1)}}$	Thm. 1.2
AG subcode	Explicit†	$\exp(\tilde{O}(1/\varepsilon^2))$	$\ell = 2^{2^{(\log^* n)^2}}$	$n^{O(1)}(1/\varepsilon)^{O(\ell)}$	Thm. 1.2

Figure 1: Parameters of various constructions of codes that enable list decoding $(1 - R - \varepsilon)$ fraction of errors, with rate R. The last four lines are from this work. “Explicit” means the code can be constructed in deterministic polynomial time. The † refers to quasi-polynomial construction time. The rows with first column in boldface are not dominated by other constructions. The last row gives the first deterministic construction of algebraic codes for efficient optimal rate list decoding over constant-sized alphabets.

1.1 Our results for Reed-Solomon and Algebraic-geometric codes

Below is a statement of our result on list decoding RS codes (the details can be found in Section 7.1).

Theorem 1.1 (List decoding Reed-Solomon (sub)-codes)**.** *Let $\varepsilon > 0$, and k, n be integers with $1 \leqslant k < n$. Let $\mathbb{F}_q$ be a field of characteristic 2, with $n \leqslant q \leqslant \mathrm{poly}(n)$. Let $m = \Theta(1/\varepsilon^2)$. Consider the $[n, k]$ Reed-Solomon code over alphabet $\mathbb{F}_{q^m}$ of rate $R = k/n$ whose evaluation points lie in the subfield $\mathbb{F}_q$. This code can be list decoded in polynomial time up to $(1 - \varepsilon)(n - k)$ errors pinning down the candidate messages to an subspace over $\mathbb{F}_q$ of dimension at most εmk.*

Further, there are subcodes of this RS code, of rate $(1 - \varepsilon)R$, list-decodable in polynomial time (for fixed $\varepsilon > 0$) from fraction $(1 - R - \varepsilon)$ of errors using list size ℓ, with

(i) $\ell = O(1/(R\varepsilon))$. (This subcode is non-linear and admits a Monte Carlo construction in $(n/\varepsilon)^{O(1)}$ time.)

(ii) $\ell = n^{O(1/\varepsilon^2)}$, with the list contained in a subspace of dimension $O(1/\varepsilon^2)$ over $\mathbb{F}_q$. (This subcode is $\mathbb{F}_q$-linear and can be constructed in $n^{\mathrm{poly}(1/\varepsilon)}$ time.)

As a comparison, folded RS and derivative codes offer a list size guarantee similar to the deterministic construction (ii) above (in fact, the bound on dimension is better and equals $O(1/\varepsilon)$) [19]. Those codes also admit a randomized subcode construction (using appropriate subspace-evasive sets) that brings down the list size to $O_\varepsilon(1)$, similar to (i) above [19], and an explicit construction to reduce the list size to $\exp(\tilde{O}(1/\varepsilon))$ [6]. These and other previous results for list decoding from $1 - R - \varepsilon$ error fraction are listed in Figure 1. The main point of Theorem 1.1 above is not the parameters, but that we can construct subcodes of Reed-Solomon codes themselves that can be list decoded up to the optimal error fraction with polynomial complexity. Perhaps more importantly, the methods extend to (i) Algebraic-geometric codes, leading to explicit codes offering new trade-offs (the last row of Figure 1), and (ii) Gabildulin codes for the rank metric, giving the first algorithm to list decode beyond half the distance with positive rate, as discussed in Section 1.2.

AG subcodes. In our work [20], we extended the linear-algebraic list decoding algorithm to folded algebraic-geometric (AG) codes, and showed that (pseudorandom) subcodes of certain folded AG codes achieve similar parameters to part (i) of Theorem 1.1 and in addition have an alphabet size $\exp(\tilde{O}(1/\varepsilon^2))$ (the alphabet size using RS codes in Theorem 1.1 is $n^{O(1/\varepsilon^2)}$). Here, we extend our approach for RS codes with evaluation points in a subfield to algebraic-geometric codes based on constant extensions. Using pseudorandomly constructed subcodes of such AG codes, in this work we match these parameters obtained in [20]. Perhaps more significantly, we also give a deterministic subcode construction (in quasi-polynomial time) with near-constant list size. This gives the ***first deterministic construction*** of an *algebraic* code family with optimal rate list decoding (i.e., list decoding $1 - R - \varepsilon$ fraction of errors with rate R) over an *alphabet of constant size* (that depends only on ε). Previously, such codes were known only via code concatenation with inner codes found by brute-force combined with expander-based symbol redistribution [17]. Additionally, the list size in our construction is bounded by a very slowly growing function of the block length. The one minus point is that the construction time is quasi-polynomial in the block length. Below is a more formal statement that we can prove for AG list decoding (the details appear in Section 7.2).

Theorem 1.2 (List decodable algebraic-geometric (sub)-codes)**.** *For arbitrary $R, \varepsilon \in (0,1)$, pick a prime power $q = \Theta(1/\varepsilon^2)$ and integer $m = \Theta(1/\varepsilon^2)$. Then, we can construct a family of algebraic-geometric codes over $\mathbb{F}_{q^m}$ of rate R that can be list decoded from a fraction $(1-R)(1-\varepsilon)$ of errors, pinning down the candidate messages to an subspace over $\mathbb{F}_q$ of dimension at most εmk (where k is the dimension of the code).*

Further, there are subcodes of this AG code, of rate at least $(1-\varepsilon)R$ with the following guarantees for list decoding from fraction $(1 - R - \varepsilon)$ of errors:

(i) List size $O(1/(R\varepsilon))$, decoding complexity $\text{poly}(N, \exp(1/\varepsilon^2))$*. This subcode is non-linear and admits a Monte Carlo construction in $(N/\varepsilon)^{O(1)}$ time where N is the block length.*

(ii) An $(N/\varepsilon)^{O(1)}$ time decoder that finds a subspace of dimension $\exp(O(\log^* N)^2)$ *over $\mathbb{F}_q$ containing the list.*[3] *This subcode is* $\mathbb{F}_q$-linear *and can be constructed deterministically in $N^{O(\log_q^3 N)}$ time.*

A point worth noting about our deterministic constructions (part (ii) of Theorems 1.1 and 1.2) is that the subcode is *linear* over the subfield $\mathbb{F}_q$ of the alphabet $\mathbb{F}_{q^m}$. Further, the list decoder will prune the $\approx \varepsilon mk$-dimensional to a near-constant dimensional subspace by imposing additional $\mathbb{F}_q$*-linear* constraints on the message.

To summarize, our basic construction is based on Reed-Solomon codes themselves, and not any folded version or other variant.[4] The underlying approach can be extended to AG codes, and yields codes matching previous parameters for randomized constructions, and a deterministic construction with improved parameters. Figure 1 compares parameters of different constructions for optimal rate list decoding.

As mentioned earlier, a further advantage of our approach is that it extends to give similar guarantees for *Gabidulin codes* [8], which are the rank-metric analog of Reed-Solomon codes, based on linearized polynomials. In fact, we originally discovered the new algorithm in the context of rank-metric codes, and later realized it also applied for RS and AG codes. We describe rank-metric codes and the prior and our results for list decoding them next.

[3] $\log^* n$ denotes the number of iterated logarithms to the base 2 needed to reach a number below 1.

[4] We note though that restricting the evaluation points to a subfield makes our construction an "interleaved" RS code construction over the subfield $\mathbb{F}_q$. Decoding algorithms for interleaved RS codes for close to a fraction $(1-R)$ of *random* errors were given in [4, 3].

1.2 Rank-metric codes

A rank-metric code $\mathcal{C}$ is a collection of matrices of certain dimension over a finite field (say, $\mathcal{C} \subseteq \mathbb{M}_{n\times t}(\mathbb{F}_h)$, with $n \leqslant t$, where $\mathbb{M}_{n\times t}(\mathbb{F}_h)$ denotes the set of $n \times t$ matrices with entries in $\mathbb{F}_h$). The rate of $\mathcal{C}$ is defined to be $\log_h |\mathcal{C}|/(nt)$. The notion of distance $d(A, B)$ between matrices A, B is the rank distance $\operatorname{rank}(A - B)$, and the (rank) distance d of $\mathcal{C}$ equals $\min_{A\neq B\in\mathcal{C}} \operatorname{rank}(A - B)$. Gabidulin gave a construction of rank-metric codes which are the analog of RS codes in the world of linearized polynomials [8]. The messages of Gabidulin codes are h-linearized polynomials over $\mathbb{F}_{h^t}$ of h-degree less than k, and such a polynomial f is encoded into $(f(\alpha_1), f(\alpha_2), \ldots, f(\alpha_n))^T \in \mathbb{M}_{n\times t}(\mathbb{F}_h)$, where $\alpha_1, \ldots, \alpha_n \in \mathbb{F}_{h^t}$ are linearly independent over $\mathbb{F}_h$, and we think of $f(\alpha_j)$ as a column vector in $\mathbb{F}_h^t$ under a fixed $\mathbb{F}_h$-basis of $\mathbb{F}_{h^t}$. The rate of this code is k/n, and its rank distance is $n - k + 1$, which is optimal and meets the Singleton bound for rank-metric codes.

Prior and recent work. The rank metric was first considered in the context of coding theory by Delsarte (who used the terminology of bilinear forms) [5]. In addition to being a natural concept of inherent interest, rank-metric codes are motivated by applications such as reliable communication of messages in linear network coding [37], crisscross error-correction in magnetic tape recording or memory chip arrays [34], space-time coding in wireless communications [28], public key cryptosystems [10, 27], etc. The Gabidulin codes play a preeminent role in the subject, and the problem of unique decoding them up to $(n - k)/2$ rank errors (this means recovering a codeword matrix M given $M + E$ where the error matrix E has rank at most $(n - k)/2$) has received a lot of attention. In fact, it has been solved several times, by adapting the different approaches for unique decoding Reed-Solomon codes to the linearized setting, starting with Gabidulin's original paper [8], and later in [9, 34, 33, 26, 36, 23].

Despite this interest and many results paralleling RS codes, an algorithm for *list* decoding Gabidulin codes beyond half the distance has remained elusive. Recently, by adapting a construction from [2], Wachter-Zeh showed that for Gabidulin codes, correcting more than a fraction $1 - \sqrt{R}$ of rank errors is not possible with a polynomial sized list [40]. This is in contrast with the situation for Reed-Solomon codes where we still do not know if the list size can become super-polynomial beyond the $1 - \sqrt{R}$ bound. Perhaps this indicates the difference in behavior of list decoding Gabidulin and RS codes. In this context, it is worth mentioning that no analog of the Johnson bound (which implies a small list size up to radius $1 - \sqrt{R}$ for RS codes) is known for rank-metric codes. Therefore, we currently do not even know if list decoding Gabidulin codes up to radius $1 - \sqrt{R}$, or for that matter any error fraction exceeding $(1 - R)/2$, is *combinatorially* feasible (in that the number of close-by codewords is guaranteed to be small).

Recently, a folded variant of Gabidulin codes (paralleling the folded RS codes of [17]) was considered in [14] (and independently in [30]), and a linear-algebraic list decoding algorithm along the lines of [19] was given for these codes. The results of [14] are stated for the model of "subspace codes" and deal with the analog of Gabidulin codes in this setting defined by [41, 23]. But subspace codes are closely related to rank-metric codes (see [37]), and the results of [14] can be readily translated to rank-metric codes. In particular, the authors of [14] give a code construction that can correct a fraction $(1 - \varepsilon)$ of rank errors for any $\varepsilon > 0$, but the rate is polynomially small. The loss in rate occurs because in order to keep the list size small, the h-linearized message polynomials must be restricted to have coefficients in the base field $\mathbb{F}_h$ instead of $\mathbb{F}_{h^t}$, and this makes the rate a factor t smaller. The same drawback applies to [30] (though the title claims decodability up to the Singleton bound, in this case the output list size would be exponentially large in the dimensions of the matrix).

List decoding of subspace codes that are in some sense the linearized analog of Parvaresh-Vardy codes [31] was studied in [29], but their algorithm could only correct "insertions" (and not removal of basis elements from the subspace), and so does not immediately apply to the rank-metric setting.

Our results. In this work, we consider Gabidulin codes where the n evaluation points $\alpha_1, \alpha_2, \ldots, \alpha_n$ belong to a subfield $\mathbb{F}_{h^n}$ of the field $\mathbb{F}_{h^t}$ over which the message polynomials are defined (so we require $n|t$). The evaluation points thus form an $\mathbb{F}_h$-basis of $\mathbb{F}_{h^n}$. We give a list decoding algorithm for these Gabidulin codes, and combine them with suitable hierarchical subspace-evasive sets, to prove the following statement paralleling Theorem 1.1.

Theorem 1.3. *Let $\mathbb{F}_h$ be a finite field of characteristic 2, $\varepsilon > 0$, and k, n, t be integers with $1 \leqslant k < n < O(\varepsilon^2 t)$ and $n|t$. Consider the Gabidulin code $\mathcal{G} \subseteq \mathbb{M}_{n \times t}(\mathbb{F}_h)$ consisting of evaluations of h-linearized polynomials in $\mathbb{F}_{h^t}[X]$ of h-degree at most $k-1$ at an $\mathbb{F}_h$-basis of $\mathbb{F}_{h^n}$. The code $\mathcal{G}$ can be list decoded in polynomial time up to $(1-\varepsilon)(n-k)$ rank errors pinning down the candidate messages to an $\mathbb{F}_h$-subspace of dimension at most εtk.*

Further, there is a Monte Carlo construction of a subcode of this Gabidulin code, of rate $R = (1-\varepsilon)k/n$, which can be list decoded from $(1-\varepsilon)(n-k)$ errors in $\operatorname{poly}(n, \log h, \exp(1/\varepsilon^2))$ time, outputting a list of size at most $O(1/(R\varepsilon))$.

Thus we are able to give a Monte Carlo construction of a rank-metric code of rate R that is efficiently list decodable up to a fraction $(1 - R - \varepsilon)$ of rank errors. Note that we list decode up to the best possible radius, approaching the Singleton bound. Further, to our knowledge, this is the first construction of rank-metric codes with rate bounded away from zero that can be list decoded beyond the half-the-distance bound.

Subspace codes. We can also obtain a result for subspace codes using similar methods, yielding list decodable codes with trade-offs almost matching existential bounds. For simplicity we focus on rank-metric codes in the body of the paper, and discuss subspace codes only in Appendix A.

1.3 Organization

We give an overview of some of the main ideas in this work in Section 2. We abstract the key "periodicity" property of the subspaces arising in our linear-algebraic list decoding in Section 3. Section 4 discusses the list decoding algorithms for Reed-Solomon and algebraic-geometric codes, including a separate treatment of Garcia-Stichtenoth codes in Section 4.3. We put forth the notion of subspace designs and study their constructions as well as cascaded subspace designs in Section 5. Hierarchical subspace evasive (h.s.e) sets are defined and constructed in Section 6. With the necessary pseudoranom ingredients in place, in Section 7, we construct appropriate subcodes of RS and AG codes with improved list size. In Section **??**, we develop a linear-algebraic list decoder for Gabidulin codes, and construct a good subcode of that code by pre-coding using h.s.e sets.

2 Our techniques

We now discuss at a high level some of the new ingredients in this work.

Algebraic ideas. We begin by describing how restricting evaluation points to a subfield enables correcting more errors, which is the algebraic starting point of our work. The idea behind list decoding results for folded RS (or derivative) codes in [17, 19] is that the encoding of a message polynomial $f \in \mathbb{F}_Q[X]$ includes the values of f and closely related polynomials at the evaluation points. Given a string not too far from the encoding of f, one can use this property together with the "interpolation method" to find an algebraic condition that f (and its closely related polynomials) must satisfy, eg. $A_0(X) + A_1(X)f(X) + A_2(X)f'(X) + \cdots + A_s(X)f^{(s-1)}(X) = 0$ in the case of derivative codes [19] (here $f^{(i)}$ denotes the i'th formal derivative of f, and the $A_0, A_1, \ldots, A_s$ are low-degree polynomials found by the decoder). The solutions $f(X)$ to this equation form an affine space, which can be efficiently found (and later pruned for list size reduction when we pre-code messages into a subspace-evasive set).

For Reed-Solomon codes as in Theorem 1.1, the encoding only includes the values of f at $\alpha_1, \alpha_2, \ldots, \alpha_n$. But since $\alpha_i \in \mathbb{F}_q$, we have $f(\alpha_i)^q = f^\sigma(\alpha_i)$ where f^σ is the polynomial obtained by the action of the Frobenius automorphism that maps $y \mapsto y^q$ on f (formally, $f^\sigma(X) = \sum_{j=0}^{k-1} f_j^q X^j$ if $f(X) = \sum_{j=1}^{k-1} f_j X^j$). Thus the decoder can "manufacture" the values of f^σ (and similarly $f^{\sigma^2}, f^{\sigma^3}$, etc.) at the α_i. Applying the above approach then enables finding a relation $A_0(X) + A_1(X)f(X) + A_2(X)f^\sigma(X) + \cdots + A_s(X)f^{\sigma^{s-1}}(X) = 0$, which is again an $\mathbb{F}_q$-linear condition on f that can be used to solve for f.

To extend this idea to algebraic-geometric codes, we work with constant extensions $\mathbb{F}_{q^m} \cdot F$ of algebraic function fields $F/\mathbb{F}_q$. The messages belong to a Riemann-Roch space over $\mathbb{F}_{q^m}$, but they are encoded via their evaluations at $\mathbb{F}_q$-rational points. Similarly to [20], for decoding we recover the message function f in terms of the coefficients of its local expansion at some rational point P. (The Reed-Solomon setting is a special case when $F = \mathbb{F}_q(X)$, and P is 0, i.e., the zero of X.) To get the best trade-offs, we use AG codes based on a tower of function fields due to Garcia and Stichtenoth [11, 12] which achieve the optimal trade-off between the number of $\mathbb{F}_q$-rational points and the genus. For this case, we recover messages in terms of their local expansion around the point at infinity P_∞ which is also used to define the Riemann-Roch space of messages. So we treat this setting separately (Section 4.3), after describing the framework for general AG codes first (Section 4.2).

Subspace designs and subspace-evasive sets. In the case of folded RS codes, the solutions to the equation $A_0(X) + A_1(X)f(X) + A_2(X)f(\gamma X) + \cdots + A_s(X)f(\gamma^{s-1}X) = 0$ are restricted to an s-dimensional space over $\mathbb{F}_Q$ with $Q = q^m$, if $f \in \mathbb{F}_Q[X]$ [19], and a similar statement holds for derivative codes. For the Reed-Solomon codes with evaluation points in $\mathbb{F}_q$ considered in this work, *each* coefficient f_j, $j = 0, 1, \ldots, k-1$, of $f = f_0 + f_1 X + \cdots + f_{k-1}X^{k-1} \in \mathbb{F}_{q^m}[X]$ will be restricted to an s-dimensional $\mathbb{F}_q$-subspace of $\mathbb{F}_{q^m}$. This would lead to a list size bound of $(q^s)^k$, which is exponentially large. Thus, to get polynomial complexity, we need to prune this space by intersecting it with a large subspace-evasive subset of $\mathbb{F}_q^{mk}$ (where we treat the coefficient vector $(f_0, f_1, \ldots, f_{k-1})$ as an element of $\mathbb{F}_q^{mk}$ by fixing some $\mathbb{F}_q$-basis of $\mathbb{F}_{q^m}$). Despite the large dimension, the solution subspace has a nice "periodic" structure; namely, once $f_0, f_1, \ldots, f_{i-1}$ are fixed, the i'th "block" f_i belongs to an s-dimensional subspace of $\mathbb{F}_q^m$. Exploiting this, we can use hierarchically subspace-evasive (h.s.e) sets of the kind constructed in [20] to randomly construct a subcode achieving list size as small as $O(1/\varepsilon)$.[5]

Naively computing the intersection of the solution space with the h.s.e set will involve trying all possibilities in $\mathbb{F}_q^m$ to compute the allowed extensions f_i to each partial solution $f_0, \ldots, f_{i-1}$. The resulting q^m time complexity will be a large polynomial (like n^{1/ε^2}) in the Reed-Solomon case, and worse still, super-polynomial in the case of Gabidulin codes for rank-metric. To circumvent this problem, we compose the h.s.e set with an "inner" subspace-evasive subset $\mathcal{I} \subset \mathbb{F}_q^m$ in each of the k blocks. (That is, we insist f_i belongs to $\mathcal{I}$, in addition to $(f_0, \ldots, f_{k-1})$ belonging to the h.s.e set.)

For the subset $\mathcal{I}$, we use the subspace-evasive variety constructed by Dvir and Lovett [6] (with a different choice of degree parameters to accommodate any field size). The intersection of an s-dimensional subspace with this variety can be found in time polynomial in the intersection size. This allows us to find the allowed

[5] We actually observe and use a simplification of this construction, by defining the set based on values of random polynomials instead of their zero sets, following [19, Sec. 4.1].

extensions f_i to $f_0, \ldots, f_{i-1}$ efficiently without searching over all q^m possibilities, and leads to the claimed runtime bounds for decoding the (randomized) subcodes of RS and Gabidulin codes in Theorems 1.1 and 1.3.

The h.s.e sets are constructed randomly and lead to Monte Carlo constructions of the associated subcodes. We next turn to our *deterministic* subcode constructions (parts (ii) of Theorems 1.1 and 1.2). The starting point for this is an observation we make that the periodic property of the subspace of candidate solutions is even nicer than what was used in [20]. Specifically, there is a subspace $W \subset \mathbb{F}_q^m$ such that once $f_0, f_1, \ldots, f_{i-1}$ are fixed, f_i belongs to a coset of W (the point is that this W is the *same* for every block i). Our idea then is to restrict f_i to belong to a subspace H_i where $H_1, H_2, \ldots, H_k$ are a collection of subspaces in $\mathbb{F}_q^m$ such that for any s-dimensional subspace $W \subset \mathbb{F}_q^m$, only a small number of them have non-trivial intersection with W. More precisely, we require that $\sum_{i=1}^k \dim(W \cap H_i)$ is small. We call such a collection as a *subspace design* in $\mathbb{F}_q^m$. We feel that the concept of subspace designs is interesting in its own right, and view the introduction of this notion in Section 5 as a key contribution in this work.

There are well known explicit constructions of "spreads" which are a collection of $\approx q^{m/2}$ subspaces of $\mathbb{F}_q^m$ which pairwise intersect only at 0 [21]. These would ensure that $\sum_{i=1}^k \dim(W \cap H_i) \leqslant \dim(W) \leqslant s$. But the subspaces in such spreads necessarily have dimension at most $m/2$, so restricting $f_i \in H_i$ for such subspaces would incur a factor two loss in rate. We instead resort to random choices of the subspaces. We prove that $q^{\Omega(\varepsilon m)}$ random subspaces of dimension $(1-\varepsilon)m$ have small total intersection with every s-dimensional W. Fortunately, we are able to derandomize this construction using conditional expectations to also get a deterministic construction. This leads to our explicit subcode of Reed-Solomon codes promised in Theorem 1.1, part (ii).

For explicit subcodes of algebraic-geometric codes (Section 7.2), we need additional ideas. The dimension k in the case of AG codes is much larger than the alphabet size q^m (that's the whole point of generalizing to AG codes). So we cannot have a subspace design in $\mathbb{F}_q^m$ with k subspaces. We therefore use several "layers" of subspace designs in a cascaded fashion – the first one in $\mathbb{F}_q^m$, the next one in $\mathbb{F}_q^{m_1}$ for $m_1 \gg q^{\sqrt{m}}$, the third one in $\mathbb{F}_q^{m_2}$ for $m_2 \gg q^{\sqrt{m_1}}$ and so on. Since the m_i's increase exponentially, we only need about $\log^* k$ levels of subspace designs. Each level incurs about a factor $1/\varepsilon$ increase in the dimension of the "period subspace" (which is W when we begin). With a careful technical argument and choice of parameters, we are able to obtain the bounds of Theorem 1.2, part (ii).

The field size in Gabidulin codes is too large to accommodate a deterministic construction of subspace designs using our methods. So for the Gabidulin case, we only construct subcodes using h.s.e sets.

3 Periodic subspaces

In this section we formalize a certain "periodic" property of affine subspaces that will arise in our list decoding application. A property of similar nature was formulated in our earlier work [20]; here we give a more restrictive definition which turns out to more accurately capture the kind of subspaces we encounter. This in turn facilitates pruning the list of candidate solutions in the subspace via appropriate pre-coding of the messages.

We begin with some notation. For a vector $\mathbf{y} = (y_1, y_2, \ldots, y_m) \in \mathbb{F}_q^m$ and positive integers $t_1 \leqslant t_2 \leqslant m$, we denote by $\mathrm{proj}_{[t_1,t_2]}(\mathbf{y}) \in \mathbb{F}_q^{t_2-t_1+1}$ its projection onto coordinates t_1 through t_2, i.e., $\mathrm{proj}_{[t_1,t_2]}(\mathbf{y}) = (y_{t_1}, y_{t_1+1}, \ldots, y_{t_2})$. When $t_1 = 1$, we use $\mathrm{proj}_t(\mathbf{y})$ to denote $\mathrm{proj}_{[1,t]}(\mathbf{y})$. These notions are extended to subsets of strings in the obvious way: $\mathrm{proj}_{[t_1,t_2]}(S) = \{\mathrm{proj}_{[t_1,t_2]}(\mathbf{x}) \mid \mathbf{x} \in S\}$.

Definition 1 (Periodic subspaces)**.** *For positive integers r, b, Λ and $\kappa := b\Lambda$, an affine subspace $H \subset \mathbb{F}_q^{\kappa}$ is said to be (r, Λ, b)-periodic if there exists a subspace $W \subseteq \mathbb{F}_q^{\Lambda}$ of dimension at most r such that for every $j = 1, 2, \ldots, b$, and every "prefix" $\mathbf{a} \in \mathbb{F}_q^{(j-1)\Lambda}$, the projected affine subspace of $\mathbb{F}_q^{\Lambda}$ defined as*

$$\{\mathrm{proj}_{[(j-1)\Lambda+1, j\Lambda]}(\mathbf{x}) \mid \mathbf{x} \in H \textit{ and } \mathrm{proj}_{(j-1)\Lambda}(\mathbf{x}) = \mathbf{a}\}$$

is contained in an affine subspace of $\mathbb{F}_q^{\Lambda}$ given by $W + \mathbf{v_a}$ for some vector $\mathbf{v_a} \in \mathbb{F}^{\Lambda}$ dependent on $\mathbf{a}$.[6]

The motivation of the above definition will be clear when we present our linear-algebraic list decoders, which will pin down the messages that must be output within an $(s-1, m, k)$-periodic (affine) subspace of $\mathbb{F}_q^{mk}$ (where q^m will be the alphabet size of the code, k its dimension, and s a parameter of the algorithm that governs how close the decoding performance approaches the Singleton bound).

The following properties of periodic affine spaces follow from the definition.

Claim 3.1. *Let H be an (r, Λ, b)-periodic affine subspace. Then for each $j = 1, 2, \ldots, b$,*

1. *the projection of H to the first j blocks of Λ coordinates, $\mathrm{proj}_{j\Lambda}(H) = \{\mathrm{proj}_{j\Lambda}(\mathbf{x}) \mid \mathbf{x} \in H\}$, has dimension at most jr. (In particular H has dimension at most br.)*
2. *for each $\mathbf{a} \in \mathbb{F}_q^{(j-1)\Lambda}$, there are at most q^r extensions $\mathbf{y} \in \mathrm{proj}_{j\Lambda}(H)$ such that $\mathrm{proj}_{(j-1)\Lambda}(\mathbf{y}) = \mathbf{a}$.*

For an affine space H, its *underlying subspace* is the subspace S such that H is a coset of S.

Definition 2 (Representing periodic affine subspaces)**.** *The canonical representation of an (r, Λ, b)-periodic subspace H consists of a matrix $B \in \mathbb{F}_q^{\Lambda \times \Lambda}$ such that $\ker(B)$ has dimension at most r, and vectors $a_i \in \mathbb{F}_q^{\Lambda}$ and matrices $A_{i,j} \in \mathbb{F}_q^{\Lambda \times \Lambda}$ for $1 \leqslant i \leqslant b$ and $1 \leqslant j < i$, such that $\mathbf{x} \in H$ if and only if for every $i = 1, 2, \ldots, b$ the following holds:*

$$a_i + \Big(\sum_{j=1}^{i-1} A_{i,j} \cdot \mathrm{proj}_{[(j-1)\Lambda+1, j\Lambda]}(\mathbf{x})\Big) + B \cdot \mathrm{proj}_{[(i-1)\Lambda+1, i\Lambda]}(\mathbf{x}) = 0 \ .$$

Ultra-periodic subspaces. For our result on pre-coding algebraic-geometric codes with subspace designs, we will exploit an even stronger property that holds for the subspaces output by the linear-algebraic list decoder. We formalize this notion below.

Definition 3. *An affine subspace H of $\mathbb{F}_q^{\kappa}$ is said to be (r, Λ)-*ultra periodic *if for every integer ℓ, $1 \leqslant \ell \leqslant \frac{\kappa}{\Lambda}$, setting $b_\ell = \lfloor \frac{\kappa}{\ell\Lambda} \rfloor$, we have $\mathrm{proj}_{b_\ell \cdot \ell\Lambda}(H)$ is $(\ell r, \ell\Lambda, b_\ell)$-periodic.*

The definition captures the fact that the subspace is periodic not only for blocks of size Λ, but also for block sizes that are multiples of Λ. Thus the subspace looks periodic in all "scales" simultaneously.

4 List decoding of Reed-Solomon and algebraic-geometric codes

In this section, we will present a linear-algebraic list decoding algorithm for algebraic-geometric (AG) codes based on evaluations of functions at rational points over a subfield. The algorithm will manage to

[6] In fact, in our applications this affine space will either be empty or *equal* to a coset of W, but for a simpler definition we just require that it is always contained in a coset of W.

correct a large fraction of errors, and pin down the possible messages to a well-structured affine subspace of dimension much smaller than that of the code. For simplicity, we begin with the case of Reed-Solomon codes in Section 4.1. We then extend it to a general framework for decoding AG codes based on constant field extensions in Section 4.2. Finally, in Section 4.1, we instantiate the general framework (with a slight twist) to codes based on the Garcia-Stichtenoth tower.

4.1 Decoding Reed-Solomon codes

Our list decoding algorithm will apply to Reed-Solomon codes with evaluation points in a subfield, defined below.

Definition 4 (Reed-Solomon code with evaluations in a subfield)**.** *Let $\mathbb{F}_q$ be a finite field with q elements, and m a positive integer. Let n, k be positive integers satisfying $1 \leqslant k < n \leqslant q$. The Reed-Solomon code $\mathsf{RS}^{(q,m)}[n,k]$ is a code over alphabet $\mathbb{F}_{q^m}$ that encodes a polynomial $f \in \mathbb{F}_{q^m}[X]$ of degree at most $k-1$ as*

$$f(X) \mapsto (f(\alpha_1), f(\alpha_2), \cdots, f(\alpha_n))$$

where $\alpha_1, \alpha_2, \ldots, \alpha_n$ are an arbitrary sequence of n distinct elements of $\mathbb{F}_q$.

Note that while the message polynomial has coefficients from $\mathbb{F}_{q^m}$, the encoding only contains its evaluations at points in the subfield $\mathbb{F}_q$. The above code has rate k/n, and minimum distance $(n-k+1)$.

We now present a list decoding algorithm for the above Reed-Solomon codes. Suppose the codeword $(f(\alpha_1), f(\alpha_2), \cdots, f(\alpha_n))$ is received as $(y_1, y_2, \ldots, y_n) \in \mathbb{F}_{q^m}^n$ with at most $e = \tau n$ errors (i.e., $y_i \neq f(\alpha_i)$ for at most e values of $i \in \{1, 2, \ldots, n\}$). The goal is to recover the list of all polynomials of degree less than k whose encoding is within Hamming distance e from y. As is common in algebraic list decoders, the algorithm will have two steps: (i) interpolation to find an algebraic equation the message polynomials must satisfy, and (ii) solving the equation for the candidate message polynomials.

Interpolation step. Let $1 \leqslant s \leqslant m$ be an integer parameter of the algorithm. Choose the "degree parameter" D to be

$$D = \left\lfloor \frac{n-k+1}{s+1} \right\rfloor . \tag{1}$$

Definition 5 (Space of interpolation polynomials)**.** *Let $\mathcal{P}$ be the space of polynomials $Q \in \mathbb{F}_{q^m}[X, Y_1, Y_2, \ldots, Y_s]$ of the form*

$$Q(X, Y_1, Y_2, \ldots, Y_s) = A_0(X) + A_1(X)Y_1 + A_2(X)Y_2 + \cdots + A_s(X)Y_s\ , \tag{2}$$

with each $A_i \in \mathbb{F}_{q^m}[X]$ and $\deg(A_0) \leqslant D + k - 1$ and $\deg(A_i) \leqslant D$ for $i = 1, 2, \ldots, s$.

The lemma below follows because for our choice of D, the number of degrees of freedom for polynomials in $\mathcal{P}$ exceeds the number n of interpolation conditions (3). We include the easy proof for completeness.

Lemma 4.1. *There exists a nonzero polynomial $Q \in \mathcal{P}$ such that*

$$Q(\alpha_i, y_i, y_i^q, y_i^{q^2}, \cdots, y_i^{q^{s-1}}) = 0 \quad \textit{for} \quad i = 1, 2, \ldots, n\ . \tag{3}$$

Further such a Q can be found using $O(n^3)$ operations over $\mathbb{F}_{q^m}$.

Proof. Note that $\mathcal{P}$ is an $\mathbb{F}_{q^m}$-vector space of dimension

$$(D+k)+s(D+1)=(D+1)(s+1)+k-1>n,$$

where the last inequality follows from our choice (1). The interpolation conditions required in the lemma impose n homogeneous linear conditions on Q. Since this is smaller than the dimension of $\mathcal{P}$, there must exist a nonzero $Q \in \mathcal{P}$ that meets the interpolation conditions

$$Q(\alpha_i, y_i, y_i^q, y_i^{q^2}, \cdots, y_i^{q^{s-1}})=0 \quad \text{for} \quad i=1,2,\ldots,n\ .$$

Finding such a Q amounts to solving a homogeneous linear system over $\mathbb{F}_{q^m}$ with n constraints and at most $\mathsf{dim}(\mathcal{P}) \leqslant n+s+2$ unknowns, which can be done in $O(n^3)$ time. □

Lemma 4.3 below shows that any polynomial Q given by Lemma 4.1 yields an algebraic condition that the message functions f we are interested in list decoding must satisfy.

Definition 6 (Frobenius action on polynomials)**.** *For a polynomial $f \in \mathbb{F}_{q^m}[X]$ with $f(X)=f_0+f_1X+\cdots+f_{k-1}X^{k-1}$, define the polynomial $f^\sigma \in \mathbb{F}_{q^m}[X]$ as $f^\sigma(X)=f_0^q+f_1^qX+\cdots+f_{k-1}^qX^{k-1}$.*

For $i \geqslant 2$, we define f^{σ^i} recursively as $(f^{\sigma^{i-1}})^\sigma$.

The following simple fact is key to our analysis.

Fact 4.2. *If $\alpha \in \mathbb{F}_q$, then $f(\alpha)^{q^j}=(f^{\sigma^j})(\alpha)$ for all $j=1,2,\ldots$.*

Lemma 4.3. *Suppose $Q \in \mathcal{P}$ satisfies the interpolation conditions* (3)*. Suppose $f \in \mathbb{F}_{q^m}[X]$ of degree less than k satisfies $f(\alpha_i) \neq y_i$ for at most e values of $i \in \{1,2,\ldots,n\}$ with $e \leqslant \frac{s}{s+1}(n-k)$. Then $Q(X, f(X), f^\sigma(X), f^{\sigma^2}(X), \cdots, f^{\sigma^{s-1}}(X))=0$.*

Proof. Define the polynomial $\Phi \in \mathbb{F}_{q^m}[X]$ by $\Phi(X):=Q(X,f(X),f^\sigma(X),f^{\sigma^2}(X),\cdots,f^{\sigma^{s-1}}(X))$. By the construction of Q and the fact that $\deg(f) \leqslant k-1$, we have $\deg(\Phi) \leqslant D+k-1 \leqslant \frac{n-k+1}{s+1}+k-1=\frac{n}{s+1}+\frac{s}{s+1}(k-1)$.

Suppose $y_i=f(\alpha_i)$. By Fact 4.2, we have $y_i^q=f(\alpha_i)^q=(f^\sigma)(\alpha_i)$, and similarly $y_i^{q^j}=(f^{\sigma^j})(\alpha_i)$ for $j=2,3,\ldots$. Thus for each i such that $f(\alpha_i)=y_i$, we have $\Phi(\alpha_i)=Q(\alpha_i,f(\alpha_i),f^\sigma(\alpha_i),\cdots,f^{\sigma^{s-1}}(\alpha_i))=Q(\alpha_i,y_i,y_i^q,\cdots,y_i^{q^{s-1}})=0$. Thus Φ has at least $n-e \geqslant \frac{n}{s+1}+\frac{s}{s+1}k$ zeroes. Since this exceeds the upper bound on the degree of Φ, Φ must be the zero polynomial. □

Finding candidate solutions. The previous two lemmas imply that the polynomials f whose encodings differ from $(y_1,\cdots,y_n)$ in at most $\frac{s}{s+1}(n-k)$ positions can be found amongst the solutions of the functional equation $A_0+A_1f+A_2f^\sigma+\cdots+A_sf^{\sigma^{s-1}}=0$. We now prove that these solutions form a well-structured affine space over $\mathbb{F}_q$.

Lemma 4.4. *For integers $1 \leqslant s \leqslant m$, the set of solutions $f=\sum_{i=0}^{k-1} f_iX^i \in \mathbb{F}_{q^m}[X]$ to the equation*

$$A_0(X)+A_1(X)f(X)+A_2(X)f^\sigma(X)+\cdots+A_s(X)f^{\sigma^{s-1}}(X)=0 \tag{4}$$

when at least one of $\{A_0,A_1,\ldots,A_s\}$ is nonzero is an affine subspace over $\mathbb{F}_q$ of dimension at most $(s-1)k$. Further, fixing an $\mathbb{F}_q$-basis of $\mathbb{F}_{q^m}$ and viewing each f_i as an element of $\mathbb{F}_q^m$, the solutions are an $(s-1,m,k)$-periodic subspace of $\mathbb{F}_q^{mk}$. A canonical representation of this periodic subspace (in the sense of Definition 2) can be computed in $\mathrm{poly}(k,m,\log q)$ *time.*

Proof. If f, g are two solutions to (4), then so is $\alpha f + \beta g$ for any $\alpha, \beta \in \mathbb{F}_q$ with $\alpha + \beta = 1$. So the solutions to (4) form an affine $\mathbb{F}_q$-subspace. We now proceed to analyze the structure of the subspace.

First, by factoring out a common powers of X that divide all of $A_0(X), A_1(X), \ldots, A_s(X)$, we can assume that at least one $A_{i^*}(X)$ for some $i^* \in \{0, 1, \ldots, s\}$ is not divisible by X, and has nonzero constant term. Further, if $A_1(X), \ldots, A_s(X)$ are all divisible by X, then so is $A_0(X)$, so we can take $i^* > 0$.

Let us denote $A_\iota(X) = a_{\iota,0} + a_{\iota,1}X + a_{\iota,2}X^2 + \cdots$ for $\iota = 0, 1, 2, \ldots, s$. For $l = 0, 1, 2, \ldots, k-1$, define the linearized polynomial

$$B_l(X) = a_{1,l}X + a_{2,l}X^q + a_{3,l}X^{q^2} + \cdots + a_{s,l}X^{q^{s-1}} \ . \tag{5}$$

We know that $a_{i^*,0} \neq 0$, and therefore $B_0 \neq 0$. This implies that the solutions $\beta \in \mathbb{F}_{q^m}$ to $B_0(\beta) = 0$ is a subspace, say W, of $\mathbb{F}_{q^m}$ of dimension at most $s-1$.

Fix an $i \in \{0, 1, \ldots, k-1\}$. Expanding the equation (4) and equating the coefficient of X^i to be 0, we get

$$a_{0,i} + B_i(f_0) + B_{i-1}(f_1) + \cdots + B_1(f_{i-1}) + B_0(f_i) = 0 \ . \tag{6}$$

This implies $f_i \in W + \theta_i$ for some $\theta_i \in \mathbb{F}_{q^m}$ that is determined by $f_0, f_1, \ldots, f_{i-1}$. Therefore, for each choice of $f_0, f_1, \ldots, f_{i-1}$, f_i must belong to a fixed coset of the subspace W of dimension at most $s-1$. Thus, the solutions belong to an $(s-1, m, k)$-periodic subspace. Also, it is clear from (6) that a canonical representation of the periodic subspace can be computed in $\mathrm{poly}(k, m, \log q)$ time. □

Combining Lemmas 4.3 and 4.4, we see that one can find an affine space of dimension $(s-1)k$ that contains the coefficients of all polynomials whose encodings differ from the input $(y_1, \ldots, y_n)$ in at most a fraction $\frac{s}{s+1}(1-R)$ of the positions. Note the dimension of the message space of the Reed-Solomon code $\mathsf{RS}^{(q,m)}[n,k]$ over $\mathbb{F}_q$ is km. The above lemma pins down the candidate polynomials to a space of dimension $(s-1)k$. For $s \ll m$, this is a lot smaller. In particular, it implies one can list decode in time sub-linear in the code size (the proof follows by taking $s = \lceil 1/\varepsilon \rceil$ and $m > \frac{s}{\gamma}$).

Corollary 4.5. *For every $R \in (0,1)$, and $\varepsilon, \gamma > 0$, there is a positive integer m such that for all large enough prime powers q, the Reed-Solomon code $C = \mathsf{RS}^{(q,m)}[q, Rq]$ can be list decoded from a fraction $(1 - R - \varepsilon)$ of errors in $|C|^\gamma$ time, outputting a list of size at most $|C|^\gamma$.*

Since the dimension of the subspace guaranteed by Lemma 4.4 grows linearly in k, we still cannot afford to list this subspace as the decoder's output for polynomial time decoding. Instead, we will use a "pre-code" that only allows polynomials with coefficients in a carefully chosen subset that is guaranteed to have small intersection with the space of solutions to any equation of the form (4). Further, this intersection can be found quickly without going over all solutions to (4). In Sections 5 and 6, we will see two approaches to accomplish this based on subspace designs and hierarchical subspace-evasive sets respectively.

4.2 Decoding algebraic-geometric codes

In this section we generalize the Reed-Solomon algorithm to algebraic-geometric codes. The description in this section will be for a general abstract AG code. So we will focus on the algebraic ideas, and not mention complexity estimates. The next subsection will focus on a specific AG code based on Garcia-Stichtenoth function fields, which will require a small change to the setup, and where we will also mention computational aspects. We assume familiarity with the basic setup of algebraic function fields and codes based on function fields, and use standard terminology and notation; the reader is referred to Stichtenoth's book for basic background information [38].

Let $F/\mathbb{F}_q$ be a function field of genus g. Let $P_\infty, P_1, P_2, \ldots, P_N$ be $N+1$ distinct $\mathbb{F}_q$-rational places. Let $\sigma \in \mathrm{Gal}(\mathbb{F}_{q^m}/\mathbb{F}_q)$ be the Frobenius automorphism, i.e, $\alpha^\sigma = \alpha^q$ for all $\alpha \in \mathbb{F}_{q^m}$. Then we can extend σ to an automorphism in $\mathrm{Gal}(F_m/F)$, where F_m is the constant extension $\mathbb{F}_{q^m} \cdot F$. Note that $P^\sigma = P$ for any place of F.

For a place P of F, we denote by ν_P the discrete valuation of P. For an integer l, we consider the Riemann-Roch space over $\mathbb{F}_q$ defined by

$$\mathcal{L}(lP_\infty) := \{h \in F \setminus \{0\} : \ \nu_{P_\infty}(h) \geqslant -l\} \cup \{0\}.$$

Then the dimension $\ell(lP_\infty)$ is at least $l-g+1$ and equality holds if $l \geqslant 2g-1$. Furthermore, we define the Riemann-Roch space over $\mathbb{F}_{q^m}$ by

$$\mathcal{L}_m(lP_\infty) := \{h \in F_m \setminus \{0\} : \ \nu_{P_\infty}(h) \geqslant -l\} \cup \{0\}.$$

Then $\mathcal{L}_m(lP_\infty)$ is the tensor product of $\mathcal{L}(lP_\infty)$ with $\mathbb{F}_{q^m}$. This implies that

$$\dim_{\mathbb{F}_{q^m}}(\mathcal{L}_m(lP_\infty)) = \dim_{\mathbb{F}_q}(\mathcal{L}(lP_\infty))$$

and an $\mathbb{F}_q$-basis of $\mathcal{L}(lP_\infty)$ is also an $\mathbb{F}_{q^m}$-basis of $\mathcal{L}_m(lP_\infty)$.

Consider the Goppa geometric code defined by

$$C(m;l) := \{(f(P_1), f(P_2), \ldots, f(P_N)) : \ f \in \mathcal{L}_m(lP_\infty)\}.$$

The following result is a fundamental fact about algebraic-geometric codes.

Lemma 4.6. *The above code $C(m;l)$ is an $\mathbb{F}_{q^m}$-linear code over $\mathbb{F}_{q^m}$, rate at least $\frac{l-g+1}{N}$, and minimum distance at least $N-l$.*

We now present a list decoding algorithm for the above codes. The algorithm follows the linear-algebraic list decoding algorithm for RS codes. Suppose a codeword encoding $f \in \mathcal{L}_m((k+2g-1)P_\infty)$ is transmitted and received as $\mathbf{y} = (y_1, y_2, \ldots, y_N)$.

Given such a received word, we will interpolate a nonzero linear polynomial over F_m

$$Q(Y_1, Y_2, \ldots, Y_s) = A_0 + A_1 Y_1 + A_2 Y_2 + \cdots + A_s Y_s \tag{7}$$

where $A_i \in \mathcal{L}_m(DP_\infty)$ for $i = 1, 2, \ldots, s$ and $A_0 \in \mathcal{L}_m((D+k+2g-1)P_\infty)$ with the degree parameter D chosen to be

$$D = \left\lfloor \frac{N-k+(s-1)g+1}{s+1} \right\rfloor. \tag{8}$$

If we fix a basis of $\mathcal{L}_m(DP_\infty)$ and extend it to a basis of $\mathcal{L}_m((D+k+2g-1)P_\infty)$, then the number of freedoms of A_0 is at least $D+k+g$ and the number of freedoms of A_i is at least $D-g+1$ for $i \geqslant 1$. Thus, the total number of freedoms in the polynomial Q equals

$$s(D-g+1) + D + k + g = (s+1)(D+1) - (s-1)g - 1 + k > N. \tag{9}$$

for the above choice (8) of D. The interpolation requirements on $Q \in F_m[Y_1, \ldots, Y_s]$ are the following:

$$Q(y_i, y_i^\sigma, \ldots, y_i^{\sigma^{s-1}}) = A_0(P_i) + A_1(P_i) y_i + A_2(P_i) y_i^\sigma + \cdots + A_s(P_i) y_i^{\sigma^{s-1}} = 0 \tag{10}$$

for $i = 1, 2, \ldots, N$. Thus, we have a total of N equations to satisfy. Since this number is less than the number of freedoms in Q, we can conclude that a nonzero linear function $Q \in F_m[Y_1, \ldots, Y_s]$ of the form (7) satisfying the interpolation conditions (10) can be found by solving a homogeneous linear system over $\mathbb{F}_{q^m}$ with at most N constraints and at least $s(D-g+1)+D+k+g$ variables.

The following lemma gives the algebraic condition that the message functions $f \in \mathcal{L}_m((k+2g-1)P_\infty)$ we are interested in list decoding must satisfy.

Lemma 4.7. *If f is a function in $\mathcal{L}_m((k+2g-1)P_\infty)$ whose encoding agrees with the received word $\mathbf{y}$ in at least t positions with $t > D+k+2g-1$, then*

$$Q(f, f^\sigma, \ldots, f^{\sigma^{s-1}}) = A_0 + A_1 f + A_2 f^\sigma + \cdots + A_s f^{\sigma^{s-1}} = 0. \tag{11}$$

Proof. The proof proceeds by comparing the number of zeros of the function $Q(f, f^\sigma, \ldots, f^{\sigma^{s-1}}) = A_0 + A_1 f + A_2 f^\sigma + \cdots + A_s f^{\sigma^{s-1}}$ with $D+k+2g-1$. Note that $Q(f, f^\sigma, \ldots, f^{\sigma^{s-1}})$ is a function in $\mathcal{L}_m((D+k+2g-1)P_\infty)$. If position i of the encoding of f agrees with $\mathbf{y}$, then

$$\begin{aligned}
0 &= A_0(P_i) + A_1(P_i)y_i + A_2(P_i)y_i^\sigma + \cdots + A_s(P_i)y_i^{\sigma^{s-1}} \\
&= A_0(P_i) + A_1(P_i)f(P_i) + A_2(P_i)(f(P_i))^\sigma + \cdots + A_s(P_i)(f(P_i))^{\sigma^{s-1}} \\
&= A_0(P_i) + A_1(P_i)f(P_i) + A_2(P_i)f^\sigma(P_i) + \cdots + A_s(P_i)f^{\sigma^{s-1}}(P_i) \\
&= (A_0 + A_1 f + A_2 f^\sigma + \cdots + A_s f^{\sigma^{s-1}})(P_i)
\end{aligned}$$

i.e., P_i is a zero of $Q(f, f^\sigma, \ldots, f^{\sigma^{s-1}})$. Thus, there are at least t zeros for all the agreeing positions. Hence, $Q(f, f^\sigma, \ldots, f^{\sigma^{s-1}})$ must be the zero function when $t > D+k+2g-1$. □

Let P be a rational place in F and let $T \in F$ be a local parameter of P. Then $T^\sigma = T$. Here, we have two scenarios, i.e., $P = P_\infty$ or $P \neq P_\infty$. In Subsection 4.3, we will consider the case where $P = P_\infty$ for the Garcia-Stichtenoth tower. While in this subsection, we only discuss the case where $P \neq P_\infty$. This was the case with Reed-Solomon codes with message polynomials in $\mathbb{F}_q[X]$ in Subsection 4.1 where P_∞ was the pole of X, and P the zero of X.

Assume that a function $f \in \mathcal{L}_m((k+2g-1)P_\infty)$ has a local expansion at P

$$f = \sum_{j=0}^{\infty} f_j T^j \tag{12}$$

for some $f_j \in \mathbb{F}_{q^m}$. Then f is uniquely determined by $(f_0, f_1, \ldots, f_{k+2g-1})$ since f has the pole degree at most $k+2g-1$.

Lemma 4.8. *The set of solutions $f \in \mathcal{L}_m((k+2g-1)P_\infty)$ to the equation (10)*

$$A_0 + A_1 f + A_2 f^\sigma + \cdots + A_s f^{\sigma^{s-1}} = 0$$

when at least one A_i is nonzero has size at most $q^{(s-1)(k+2g-1)}$. Further, the possible coefficients $(f_0, f_1, \ldots, f_{k+2g-1})$ of f's local expansion at P belong to an $(s-1,m)$-ultra periodic affine subspace of $\mathbb{F}_q^{(k+2g-1)m}$.

Proof. The argument is very similar to Lemma 4.4. Let $u = \min\{\nu_{P_\infty}(A_i) : i = 1, 2, \ldots, s\}$. Then it is clear that $u \geqslant 0$ and $\nu_P(A_0) \geqslant u$. Each A_i has a local expansion at P:

$$A_i = T^u \sum_{j=0}^{\infty} a_{i,j} T^j$$

for $i = 0, 1, \ldots, s$.

Assume that at P, the function f has a local expansion (12). Then f^{σ^i} has a local expansion at P as follows

$$f^{\sigma^i} = \sum_{j=0}^{\infty} f_j^{q^i} T^j .$$

For $l = 0, 1, \ldots$, define the linearized polynomial

$$B_l(X) := a_{1,l}X + a_{2,l}X^q + \cdots + a_{s,l}X^{q^{s-1}}$$

From the definition of u, one knows that $B_0(X)$ is nonzero. Equating the coefficient of T^{d+u} in $A_0 + A_1 f + A_2 f^{\sigma} + \cdots + A_s f^{\sigma^{s-1}}$ to equal 0 gives us the condition

$$a_{0,d} + B_d(f_0) + B_{d-1}(f_1) + \cdots + B_0(f_d) = 0 \; . \tag{13}$$

Let $W = \{\alpha \in \mathbb{F}_{q^m} : \;\; B_0(\alpha) = 0\}$. Then W is an $\mathbb{F}_q$-subspace of $\mathbb{F}_{q^m}$ of dimension at most $s - 1$, since B_0 is a nonzero linearized polynomial of q-degree at most $s - 1$. As in Lemma 4.4, for each fixed $f_0, f_1, \ldots, f_{d-1}$, the coefficient f_d must belong to a coset of the subspace W. This implies that the coefficients $(f_0, f_1, \ldots, f_{k+2g-1})$ belong to an $(s-1, m, k+2g-1)$-periodic subspace of $\mathbb{F}_q^{m(k+2g-1)}$. In particular, there are at most $q^{(s-1)(k+2g-1)}$ solutions $f \in \mathcal{L}_m((k+2g-1)P_\infty)$ to (10).

The equation (13) also shows that each group of ℓ successive coefficients $f_{d-\ell+1}, f_{d-\ell+2}, \cdots, f_d$ belong to cosets of the same underlying $\ell(s-1)$ dimensional subspace of $\mathbb{F}_q^{m\ell}$. This implies that $(f_0, f_1, \ldots, f_{k+2g-1})$ in fact belong to an $(s-1, m)$-ultra periodic subspace.[7] □

Restricting message functions using local expansions. Using Lemma 4.8, we will recover the message in terms of the coefficients of its local expansion at P. In order to prune the subspace of possible solutions, we will pick a subcode that corresponds to restricting the coefficients to a carefully constructed subset of all possibilities. This requires us to index message functions in terms of the local expansion coefficients. However, not all $(k + 2g - 1)$ tuples over $\mathbb{F}_{q^m}$ arise in the local expansion of functions in the k-dimensional subspace $\mathcal{L}_m((k + 2g - 1)P_\infty)$. Below we show that we can find a k-dimensional subspace of $\mathcal{L}_m((k + 2g - 1)P_\infty)$ such that their top k local expansion coefficients give rise to all k-tuples over $\mathbb{F}_{q^m}$.

Lemma 4.9. *There exist a set of functions $\{g_1, g_2, \ldots, g_k\}$ in $\mathcal{L}_m((k + 2g - 1)P_\infty)$ such that the $k \times k$ matrix A formed by taking the ith row of A to be the first k coefficients in the local expansion* (12) *for g_i at P is nonsingular.*

Proof. Let $\{\psi_1, \psi_2, \ldots, \psi_g\}$ be a basis of $\mathcal{L}_m((k + 2g - 1)P_\infty - kP)$. Extend this basis to a basis $\{\psi_1, \psi_2, \ldots, \psi_g, g_1, g_2, \ldots, g_k\}$ of $\mathcal{L}_m((k + 2g - 1)P_\infty)$. We claim that the functions $\{g_1, g_2, \ldots, g_k\}$ are our desired functions.

Suppose that the matrix A is obtained from expansion of functions g_i and it is singular. This implies that there exists elements $\{\lambda_i\}_{i=1}^k$ such that the function $\sum_{i=1}^k \lambda_i g_i$ has expansion $\sum_{i=k}^\infty a_i T^i$ at P for some $a_i \in \mathbb{F}_{q^m}$. Therefore, the function $\sum_{i=1}^k \lambda_i g_i$ belongs to the space $\mathcal{L}_m((k+2g-1)P_\infty - kP)$, i.e., $\sum_{i=1}^k \lambda_i g_i$ is a linear combination of $\psi_1, \psi_2, \ldots, \psi_g$. This forces that all λ_i are equal to 0 since $\{\psi_1, \ldots, \psi_g, g_1, g_2, \ldots, g_k\}$ is linearly independent. This completes the proof. □

[7] This ultra-periodicity was also true for the Reed-Solomon case in Lemma 4.4, but we did not state it there as we will not make use of this extra property for picking a subcode in the case of Reed-Solomon codes.

With the above lemma in place, we now describe our AG code in a manner convenient for pruning the possible local expansion coefficients.

Encoding. Assume that we have found a set of functions $\{g_1, g_2, \ldots, g_k\}$ of $\mathcal{L}_m((k+2g-1)P_\infty)$ as in Lemma 4.9. After elementary row operations on the matrix A defined in Lemma 4.9, we may assume that A is the $k \times k$ identity matrix, i.e., we assume that, for $1 \leqslant i \leqslant k$, the function g_i has expansion $T^{i-1} + \sum_{j=k}^{\infty} \lambda_{ij} T^j$ for some $\lambda_{ij} \in \mathbb{F}_{q^m}$.

Now for any subset $M \subseteq \mathbb{F}_{q^m}^k$, we may assume that our messages belong to M, and encode each message $(a_1, a_2, \ldots, a_k) \in M$ to the codeword $(f(P_1), f(P_2), \ldots, f(P_N))$, where $f = \sum_{i=1}^k a_i g_i$. Thus, our actual code is a subcode of $C(m; k+2g-1)$ given by

$$C(m; k+2g-1 \mid M) \overset{\text{def}}{=} \{(f(P_1), f(P_2), \ldots, f(P_N)) : \ f = \sum_{i=1}^{k} a_i g_i, (a_1, a_2, \ldots, a_k) \in M\} . \quad (14)$$

Decoding. To decode, we first establish the equation (11) and solve this equation to find the subspace of possible first k coefficients $f_0, f_1, \ldots, f_{k-1}$ in the local expansion of the function $f = \sum_{i=1}^k a_i g_i$ at P. The following claim implies that the message tuple $(a_1, a_2, \ldots, a_k)$ belongs to this subspace.

Lemma 4.10. *The first k coefficients $f_0, f_1, \ldots, f_{k-1}$ of the local expansion of $f = \sum_{i=1}^k a_i g_i$ at P equal $a_1, a_2, \ldots, a_k$.*

Proof. Since g_i has local expansion $T^{i-1} + \sum_{j=k}^{\infty} \lambda_{ij} T^j$, it is clear that the local expansion of f is $\sum_{i=0}^{k-1} a_{i+1} T^i + \sum_{j=k}^{\infty} a_{j+1} T^j$ for some $a_{k+1}, a_{k+2}, \ldots$ in $\mathbb{F}_{q^m}$. Thus the first k coefficients of the local expansion of f are $a_1, a_2, \ldots, a_k$. □

Combining Lemmas 4.7, 4.8, and 4.10, and recalling the choice of D in (8), we get the following.

Corollary 4.11. *For the code $C(m; k+2g-1 \mid \mathbb{F}_{q^m}^k)$, we can find an $(s-1, m)$-ultra periodic subspace of $\mathbb{F}_q^{mk}$ that includes all messages whose encoding differs from a received word $\mathbf{y} \in \mathbb{F}_{q^m}^N$ in at most $\frac{s}{s+1}(N-k) - \frac{3s+1}{s+1} g$ positions.*

4.3 Decoding the codes from the Garcia-Stichtenoth tower

Let r be a prime power and let $q = r^2$. The Garcia-Stichtenoth towers that we are going to use for our code construction were discussed in [11, 12]. The reader may refer to [11, 12] for the detailed background on the Garcia-Stichtenoth function tower. There are two optimal Garcia-Stichtenoth towers that are equivalent. For simplicity, we introduce the tower defined by the following recursive equations [12]

$$x_{i+1}^r + x_{i+1} = \frac{x_i^r}{x_i^{r-1} + 1}, \quad i = 1, 2, \ldots, e-1.$$

Put $K_e = \mathbb{F}_q(x_1, x_2, \ldots, x_e)$ for $e \geqslant 2$.

The function field K_e has at least $r^{e-1}(r^2 - r) + 1$ rational places. One of these is the "point at infinity" which is the unique pole P_∞ of x_1 (and is fully ramified). The other $r^{e-1}(r^2 - r)$ come from the rational places lying over the unique zero of $x_1 - \alpha$ for each $\alpha \in \mathbb{F}_q$ with $\alpha^r + \alpha \neq 0$. Note that for every $\alpha \in \mathbb{F}_q$ with $\alpha^r + \alpha \neq 0$, the unique zero of $x_1 - \alpha$ splits completely in K_e, i.e., there are r^{e-1} rational places lying over the zero of $x_1 - \alpha$. Let $\mathbb{P}$ be the set of all the rational places lying over the zero of $x_1 - \alpha$ for all $\alpha \in \mathbb{F}_q$

with $\alpha^r + \alpha \neq 0$. Then, intuitively, one can think of the $r^{e-1}(r^2 - r)$ rational places in $\mathbb{P}$ as being given by e-tuples $(\alpha_1, \alpha_2, \ldots, \alpha_e) \in \mathbb{F}_q^e$ that satisfy $\alpha_{i+1}^r + \alpha_{i+1} = \frac{\alpha_i^r}{\alpha_i^{r-1}+1}$ for $i = 1, 2, \ldots, e-1$ and $\alpha_1^r + \alpha_1 \neq 0$. For each value of $\alpha \in \mathbb{F}_q$, there are precisely r solutions to $\beta \in \mathbb{F}_q$ satisfying $\beta^r + \beta = \frac{\alpha^r}{\alpha^{r-1}+1}$, so the number of such e-tuples is $r^{e-1}(r^2 - r)$ ($r^2 - r$ choices for α_1, and then r choices for each successive α_i, $2 \leqslant i \leqslant e$).

The genus g_e of the function field K_e is given by

$$g_e = \begin{cases} (r^{e/2} - 1)^2 & \text{if } e \text{ is even} \\ (r^{(e-1)/2} - 1)(r^{(e+1)/2} - 1) & \text{if } e \text{ is odd.} \end{cases}$$

Thus the genus g_e is at most r^e. The ratio of g_e to the number of $\mathbb{F}_q$-rational points is at most $1/(r-1) = 1/(\sqrt{q} - 1)$.

Now we put $F = K_e$ and $F_m = \mathbb{F}_{q^m} \cdot K_e$. The encoding and decoding is almost identical to the algebraic geometric code described in the previous section except here we use P_∞ for local expansion.

Encoding. As in Lemma 4.9, we can find a set of functions $\{h_1, h_2, \ldots, h_k\}$ of $\mathcal{L}_m((k + 2g_e - 1)P_\infty)$ such that the $k \times k$ matrix A formed by taking the ith row of A to be the first k coefficients in the expansion (16) for h_i at P_∞ is nonsingular (note that in this case, the local expansion starts from $T^{-(k+2g_e-1)}$, while in the previous subsection the local expansion starts from T^0). Furthermore, after some elementary row operations on A, we may assume that, for $1 \leqslant i \leqslant k$, the function h_i has expansion $T^{-(k+2g_e-1)}\left(T^{i-1} + \sum_{j=k}^{\infty} \lambda_{ij} T^j\right)$ for some $\lambda_{ij} \in \mathbb{F}_{q^m}$.

We encode a message k-tuple $(a_1, a_2, \ldots, a_k) \in \mathbb{F}_{q^m}^k$ by the codeword $(f(P_1), f(P_2), \ldots, f(P_N))$ where $f = \sum_{i=1}^k a_i h_i$, and $P_1, P_2, \ldots, P_N$ are arbitrary $\mathbb{F}_q$-rational points (other than P_∞) in the function field. The block length N can be any integer satisfying $k \leqslant N \leqslant r^{e-1}(r^2 - r)$. As in Section 4.2, for any subset $M \subseteq \mathbb{F}_{q^m}^k$, we can consider the subcode obtained by only encoding tuples in M:

$$C_{\mathrm{GS}}(m; k + 2g_e - 1 \mid M) \stackrel{\text{def}}{=} \{(f(P_1), f(P_2), \ldots, f(P_N)) : \ f = \sum_{i=1}^{k} a_i h_i, (a_1, a_2, \ldots, a_k) \in M\} . \quad (15)$$

Computing the code. Note that an explicit specification of the code simply requires the evaluations of the basis functions $h_1, h_2, \ldots, h_k$ at the N rational points. One can find a basis of $\mathcal{L}_m(lP_\infty)$ along with its evaluations at the rational points using $\mathrm{poly}(N, l, m)$ operations over $\mathbb{F}_q$ [35] (see also [15, Sec. 7]). We can also compute the first l coefficients of the local expansion of the basis functions at P_∞ using $\mathrm{poly}(l, m)$ operations over $\mathbb{F}_q$ as described in [20]. The computation of the h_i's following the method of Lemma 4.9 only requires elementary matrix operations, so we can compute its evaluations at the rational points also in polynomial time.

List decoding. In order to list decode, we can find a functional equation $A_0 + A_1 f + A_2 f^\sigma + \cdots + A_s f^{\sigma^{s-1}}$ exactly as in Lemma 4.7. To solve for f from this equation, we consider the local expansions of the message functions f at P_∞. Let $T \in K_e$ be a local expansion of P_∞ and suppose that a function $f \in \mathcal{L}_m((k + 2g_e - 1)P_\infty)$ has a local expansion at P_∞

$$f = T^{-(k+2g_e-1)} \sum_{j=0}^{\infty} f_j T^j \quad (16)$$

for some $f_j \in \mathbb{F}_{q^m}$. As in Lemma 4.10, if $f = \sum_{i=1}^k a_i h_i$, then the top k coefficients $f_0, f_1, \ldots, f_{k-1}$ in the above local expansion equal $a_1, a_2, \ldots, a_k$. Thus we can determine such f uniquely by finding

$f_0, f_1, \ldots, f_{k-1}$. The following lemma is similar to Lemma 4.8 and shows that the coefficients belong to an ultra-periodic subspace.

Lemma 4.12. *Suppose $f \in \mathcal{L}_m((k+2g_e-1)P_\infty)$ satisfies the equation*

$$A_0 + A_1 f + A_2 f^\sigma + \cdots + A_s f^{\sigma^{s-1}} = 0$$

when at least one A_i is nonzero. Then the possible first k coefficients $(f_0, f_1, \ldots, f_{k-1})$ of f's local expansion (16) *at P_∞ belong to an $(s-1,m)$-ultra periodic affine subspace of $\mathbb{F}_q^{km}$.*

Proof. Let $u = \min\{\nu_{P_\infty}(A_i) : \; i = 1,2,\ldots,s\}$ (so that $-u$ is the maximum number of poles any A_i, $1 \leqslant i \leqslant s$, has at P_∞). Then it is clear that $u \geqslant -D$ and $\nu_{P_\infty}(A_0) \geqslant u - (k+2g_e-1)$. Each A_i has a local expansion at P_∞:

$$A_0 = T^{u-(k+2g_e-1)} \sum_{j=0}^{\infty} a_{0,j} T^j; \quad \text{and } A_i = T^u \sum_{j=0}^{\infty} a_{i,j} T^j \text{ for } i = 1,2,\ldots,s.$$

Assume that at P_∞, the function f has a local expansion (12). Then f^{σ^i} has a local expansion at P as follows

$$f^{\sigma^i} = \sum_{j=0}^{\infty} f_j^{q^i} T^j.$$

For $l = 0,1,\ldots,$ define the linearized polynomial

$$B_l(X) := a_{1,l} X + a_{2,l} X^q + \cdots + a_{s,l} X^{q^{s-1}}$$

From the definition of u, one knows that $B_0(X)$ is nonzero. Equating the coefficient of $T^{d+u-(k+2g_e-1)}$ in $A_0 + A_1 f + A_2 f^\sigma + \cdots + A_s f^{\sigma^{s-1}}$ to equal 0 gives us the condition

$$a_{0,d} + B_d(f_0) + B_{d-1}(f_1) + \cdots + B_0(f_d) = 0 \; .$$

Arguing as in Lemma 4.8 this constrains $(f_0, f_1, \ldots, f_{k-1})$ to belong to an $(s-1,m)$-ultra periodic subspaces of $\mathbb{F}_q^{mk}$. □

Similar to Corollary 4.11, we can now conclude the following:

Corollary 4.13. *The code $C_{\mathrm{GS}}(m; k+2g_e-1 \mid \mathbb{F}_{q^m}^k)$ can be list decoded from up to $\frac{s}{s+1}(N-k) - \frac{3s+1}{s+1} g_e$ errors, pinning down the messages to an $(s-1,m)$-ultra periodic subspace of $\mathbb{F}_q^{mk}$.*

We conclude the section by incorporating the trade-off between g_e and N, and stating the rate vs. list decoding radius trade-off offered by these codes, in a form convenient for improvements to the list size using subspace evasive sets and subspace designs. The claim about the number of possible solution subspaces follows since the subspace is determined by $A_0, A_1, \ldots, A_s$, and for our choice of parameter D, there are at most $q^{O(mN)}$ choices of those.

Theorem 4.14. *Let q be the even power of a prime. Let $1 \leqslant s \leqslant m$ be integers, and let $R \in (0,1)$. Then for infiitely many N (all integers of the form $q^{e/2}(\sqrt{q}-1)$), there is a deterministic polynomial time construction of an $\mathbb{F}_{q^m}$-linear code $\mathrm{GS}^{(q,m)}[N,k]$ of block length N and dimension $k = R \cdot N$ that can be list decoded in* $\mathrm{poly}(N, m, \log q)$ *time from $\frac{s}{s+1}(N-k) - \frac{3N}{\sqrt{q}-1}$ errors, pinning down the messages to one of $q^{O(mN)}$ possible $(s-1,m)$-ultra periodic $\mathbb{F}_q$-affine subspaces of $\mathbb{F}_q^{mk}$.*

5 Subspace designs

The linear-algebraic list decoder discussed in the previous sections pins down the coefficients of the message to a periodic subspace. This subspace has linear dimension, so we need to restrict the coefficients further so that the subspace can be pruned to a small list of solutions. In this section, we will use a special collection of subspaces, which we call a *subspace design* to achieve this.

Definition 7. *Let Λ be a positive integer, and q a prime power. For positive integers $r < \Lambda$ and d, an (r, d)-subspace design in $\mathbb{F}_q^\Lambda$ is a collection of subspaces of $\mathbb{F}_q^\Lambda$ such that for every r-dimensional subspace $W \subset \mathbb{F}_q^\Lambda$, we have*

$$\sum_{H \in \mathcal{H}} \dim(W \cap H) \leqslant d .$$

The cardinality of a subspace design $\mathcal{H}$ is the number of subspaces in its collection, i.e., $|\mathcal{H}|$. If all subspaces in $\mathcal{H}$ have the same dimension t, then we refer to t as the dimension *of the subspace design $\mathcal{H}$.*

The usefulness of subspace designs in the context of pruning periodic subspaces is captured by the following key lemma.

Lemma 5.1 (Periodic subspaces intersected with a subspace design)**.** *Suppose $H_1, H_2, \ldots, H_b$ are subspaces in an (r, d)-subspace design in $\mathbb{F}_q^\Lambda$, and T is a (r, Λ, b)-periodic affine subspace of $\mathbb{F}_q^{\Lambda b}$ with underlying subspace S. Then the set*

$$\mathcal{T} = \{(\mathbf{f_1}, \mathbf{f_2}, \ldots, \mathbf{f_b}) \in T \mid \mathbf{f_j} \in H_j \textit{ for } j = 1, 2, \ldots, b\}$$

is an affine subspace of $\mathbb{F}_q^{\Lambda b}$ of dimension at most d. Also, the underlying subspace of $\mathcal{T}$ is contained in $\mathcal{S} \stackrel{\text{def}}{=} S \cap (H_1 \times H_2 \times \cdots \times H_b)$.

Proof. It is clear that $\mathcal{T}$ is an affine subspace, since its elements are restricted by the set of linear constraints defining T and the H_j's. Also, the difference of two elements in $\mathcal{T}$ is contained in both the subspaces S and $(H_1 \times H_2 \times \cdots \times H_b)$, which implies that the underlying subspace of $\mathcal{T}$ is contained in $\mathcal{S}$.

We will prove the bound on dimension by proving that $|\mathcal{T}| \leqslant q^d$. To prove this, we will imagine the elements of $\mathcal{T}$ as the leaves of a tree of depth b, with the nodes at level j representing the possible projections of $\mathcal{T}$ onto the first j blocks. The root of this tree has as children the elements of the affine space $\mathrm{proj}_{[1,\Lambda]}(T) \cap H_1$. Let W be the subspace of $\mathbb{F}_q^\Lambda$ of dimension at most r associated with the periodic subspace T (in the sense of Definition 1). Note that the underlying subspace of the affine space $\mathrm{proj}_{[1,\Lambda]}(T) \cap H_1$ is contained in the subspace $W \cap H_1$.

Continuing this argument, the children of an element $\mathbf{a} \in \mathbb{F}_q^{j\Lambda}$ at level j will be $\mathbf{a}$ followed by the possible extensions of $\mathbf{a}$ to the $(j + 1)$'th block, given by

$$\{\mathrm{proj}_{[j\Lambda+1,(j+1)\Lambda]}(\mathbf{x}) \mid \mathbf{x} \in T \text{ and } \mathrm{proj}_{j\Lambda}(\mathbf{x}) = \mathbf{a}\} \cap H_{j+1} .$$

The periodic property of T and the fact that H_{j+1} is a subspace implies that the possible extensions of $\mathbf{a}$ are given by a coset of a subspace of $W \cap H_{j+1}$. Thus the nodes at level j have degree at most $q^{\dim(W \cap H_{j+1})}$ for $j = 0, 1, \ldots, b - 1$. Since the H_j's belong to an (r, d)-subspace design we have $\sum_{j=1}^{b} \dim(W \cap H_j) \leqslant d$. Therefore, the tree has at most q^d leaves, which is also an upper bound on $|\mathcal{T}|$. □

5.1 Constructing subspace designs

We now turn to the construction of subspace designs of large size and dimension. We first analyze the performance of a random collection of subspaces.

Lemma 5.2. *Let $\eta > 0$ and q be a prime power. Let r, Λ be integers $\Lambda \geqslant 8/\eta$ and $r \leqslant \eta\Lambda/2$. Consider a collection $\mathcal{H}$ of subspaces of $\mathbb{F}_q^\Lambda$ obtained by picking, independently at random, $q^{\eta\Lambda/8}$ subspaces of $\mathbb{F}_q^\Lambda$ of dimension $(1-\eta)\Lambda$ each. Then, with probability at least $1 - q^{-\Lambda r}$, $\mathcal{H}$ is an $(r, 8r/\eta)$-subspace design.*[8]

Proof. Let $\ell = 8r/\eta$, and let $M = q^{\eta\Lambda/8}$ denote the number of randomly chosen subspaces.[9] Let $H_1, H_2, \ldots, H_M$ be the subspaces in the collection $\mathcal{H}$. Fix a subspace W of $\mathbb{F}_q^\Lambda$ of dimension r. Fix a tuple of nonnegative integers $(a_1, a_2, \ldots, a_M)$ summing up to ℓ. For each $j \in \{1, 2, \ldots, M\}$, the probability that $\dim(W \cap H_j) \geqslant a_j$ is at most $q^{ra_j} q^{-\eta\Lambda a_j}$. Since the choice of the different H_j's are independent, the probability that $\dim(W \cap H_j) \geqslant a_j$ for every j is at most $q^{(r-\eta\Lambda)\ell} \leqslant q^{-\eta\Lambda\ell/2}$ (the last step uses $r \leqslant \eta\Lambda/2$).

A union bound over the at most $q^{\Lambda r}$ subspaces $W \subset \mathbb{F}_q^\Lambda$ of dimension r, and the at most $\binom{\ell+M}{\ell} \leqslant (M+\ell)^\ell \leqslant M^{2\ell}$ choices of the tuples $(a_1, a_2, \ldots, a_M)$, we get the probability that $\mathcal{H}$ is *not* an (r, ℓ)-subspace design is at most

$$q^{\Lambda r} \cdot q^{-\eta\Lambda\ell/2} \cdot (q^{\eta\Lambda/8})^{2\ell} = q^{\Lambda r} \cdot q^{-\eta\Lambda\ell/4} \leqslant q^{-\Lambda r}$$

where the last step uses $\ell \geqslant 8r/\eta$. □

Note that given a collection $\mathcal{H}$ of subspaces, one can deterministically check if it is an (r, d)-subspace design in $\mathbb{F}_q^\Lambda$ in $q^{O(\Lambda r)}|\mathcal{H}|$ time by doing a brute-force check of all r-dimensional subspaces W of $\mathbb{F}_q^\Lambda$, and for each computing $\sum_{H\in\mathcal{H}} \dim(W \cap H)$ using $|\mathcal{H}|\Lambda^{O(1)}$ operations over $\mathbb{F}_q$. Thus the above lemma already gives a *Las Vegas* construction of an (r, d)-subspace design with many subspaces each of large dimension $(1-\eta)m$ (recorded formally as a part of Lemma 5.3 below). We next observe that the construction can in fact be derandomized using the method of conditional expectations, thus giving a deterministic construction in similar runtime.

Lemma 5.3. *For parameters η, r, Λ as in Lemma 5.2, for any $b \leqslant q^{\eta\Lambda/8}$, one can compute an $(r, 8r/\eta)$-subspace design in $\mathbb{F}_q^\Lambda$ of dimension $(1-\eta)\Lambda$ and cardinality b deterministically in time polynomial in $q^{\Lambda(\Lambda+r)}(br/\eta)^{r/\eta}$. One can also compute such a subspace in $q^{O(\Lambda r)}$* Las Vegas *time.*

Proof. The claim about the Las Vegas construction follows from the above argument. Let us turn to the deterministic computation. For each subspace W of $\mathbb{F}_q^\Lambda$ of dimension r, and each b-tuple τ_W of subspaces $W_1, W_2, \ldots, W_b$ of W with $\sum_{j=1}^b \dim(W_j) = \ell = 8r/\eta$, define the indicator random variable $I(W, \tau_W)$ for the event

$$\bigwedge_{j=1}^{b} [W_j \subseteq H_j] \, ,$$

for a random choice of $(1-\eta)\Lambda$-dimensional subspaces $H_1, H_2, \ldots, H_b$ of $\mathbb{F}_q^\Lambda$. Let Z be the random variable equal to the sum of $I(W, \tau_W)$ over all choices of W, τ_W. The proof of Lemma 5.2 in fact shows that $\mathbb{E}[Z] \ll 1$ when the H_j's are chosen independently at random. Clearly, $\mathbb{E}[Z]$ can be calculated in $\text{poly}(b, q^{\Lambda r})$ time.

[8] For sake of clarity, we did make any attempt to optimize the constants.

[9] For simplicity, we ignore the floor and ceil signs in defining integers; these can be easily incorporated.

For any $t \in \{0, 1, \dots, b\}$ and any choice of subspaces $L_1, L_2, \dots, L_t$, we can also compute the conditional expectation $\mathbb{E}[Z \mid H_j = L_j, 1 \leqslant j \leqslant t]$ over the random and independent choices of $H_{t+1}, \cdots, H_b$. Indeed, for a fixed r-dimensional subspace $W \subset \mathbb{F}_q^\Lambda$ and b-tuple τ_W of its subspaces $(W_1, \dots, W_b)$ with $\sum_{j=1}^b \dim(W_j) = b$, the conditional expectation of $I(W, \tau_W)$ equals 0 if $W_j \not\subseteq L_j$ for some $j \in \{1, 2, \dots, t\}$, and equals $\prod_{j=t+1}^b q^{-\eta\Lambda \dim(W_j)}$ otherwise. The conditional expectation of Z is the sum of the conditional expectations of $I(W, \tau_W)$ over all W, τ_W, and be computed in time polynomial in the number of pairs W, τ_W, which is at most $q^{\Lambda r} \cdot q^{r\ell}(b+\ell)^\ell$. For our setting of parameters, this quantity is bounded by a polynomial in $q^{\Lambda r}(br/\eta)^{r/\eta}$.

Using the above, we can deterministically compute a sequence of subspaces $L_1, L_2, \dots, L_b$ with $Z < 1$ (and hence equal to 0), by successively picking, for $t = 1, 2, \dots, b$, a subspace $H_t = L_t$ that minimizes the conditional expectation $\mathbb{E}[Z \mid H_j = L_j, 1 \leqslant j \leqslant t]$ over the random and independent choices of $H_{t+1}, \cdots, H_b$. This choice of H_t can be made by doing a brute-force search over all $(1-\eta)\Lambda$-dimensional subspaces of $\mathbb{F}_q^\Lambda$ in $q^{O(\Lambda^2)}$ time. □

Finally, we record the construction of subspaces with low-dimensional intersection with every periodic subspace based on the above subspace designs. This form will be convenient for later use in pre-coding Reed-Solomon codes.

Theorem 5.4. *Let $\eta \in (0,1)$ and q be a prime power, and r, Λ, b be integers such that $\Lambda \geqslant 8/\eta$, $r \leqslant \eta\Lambda/2$ and $b \leqslant q^{\eta\Lambda/8}$. Then, one can construct a subspace V of $\mathbb{F}_q^{b\Lambda}$ of dimension at least $(1-\eta)b\Lambda$ in either deterministic $q^{O(\Lambda^2)}$ time or Las Vegas $q^{O(\Lambda r)}$ time with the following guarantee: For every (r, Λ, b)-periodic subspace $T \subset \mathbb{F}_q^{b\Lambda}$, $V \cap T$ is an $\mathbb{F}_q$-affine subspace of dimension at most $8r/\eta$.*

Proof. We will take $V = H_1 \times H_2 \times \cdots H_b$ where the H_i's belong to a $(r, 8r/\eta)$-subspace design in $\mathbb{F}_q^\Lambda$ of cardinality b and dimension at least $(1-\eta)\Lambda$ as guaranteed by Lemma 5.3. Clearly $\dim(V) \geqslant (1-\eta)b\Lambda$ since each H_i has dimension at least $(1-\eta)\Lambda$. The claim now follows using Lemma 5.1. □

5.2 Cascaded subspace designs

In preparation for our results about algebraic-geometric codes, whose block length $\gg q^m$ is much larger than the possible size of subspace designs in $\mathbb{F}_q^m$, we now formalize a notion that combines several "levels" of subspace designs. The definition might seem somewhat technical, but it has a natural use in our application to list-size reduction for AG codes. Note that there is no "consistency" requirement between subspace designs at different levels other than the lengths and cardinalities matching.

Definition 8 (Subspace designs of increasing length)**.** *Let l be a positive integer. For positive integers $r_0 \leqslant r_1 \leqslant \cdots \leqslant r_l$ and $m_0 \leqslant m_1 \leqslant \cdots \leqslant m_l$ such that $m_{\iota-1} | m_\iota$ for $1 \leqslant \iota \leqslant l$, an $(r_0, r_1, \dots, r_l)$-cascaded subspace design with length-vector $(m_0, m_1, \dots, m_l)$ and dimension vector $(d_0, d_1, \dots, d_{l-1})$ is a collection of l subspace designs, specifically an $(r_{\iota-1}, r_\iota)$-subspace design in $\mathbb{F}_{q^{m_{\iota-1}}}$ of cardinality $m_\iota/m_{\iota-1}$ and dimension $d_{\iota-1}$ for each $\iota = 1, 2, \dots, l$.*

Note that the $l = 1$ case of the above definition corresponds to an (r_0, r_1)-subspace design in $\mathbb{F}_q^{m_0}$ of dimension d_0 and cardinality m_1/m_0. In Lemma 5.1, we used the subspace $H_1 \times H_2 \times \cdots \times H_b$ based on a subspace design consisting of the H_i's to prune a periodic subspace. Generalizing this, we now define a subspace associated with a cascaded subspace design based on the subspace designs comprising it.

Definition 9 (Canonical subspace)**.** *Let $\mathcal{M}$ be a cascaded subspace design with length-vector $(m_0, m_1, \ldots, m_l)$ such that the ι'th subspace design in $\mathcal{M}$ has subspaces $H_1^{(\iota)}, H_2^{(\iota)}, \cdots, H_{m_\iota/m_{\iota-1}}^{(\iota)} \subset \mathbb{F}_q^{m_{\iota-1}}$, for $1 \leqslant \iota \leqslant l$.*

The canonical subspace *associated with such a cascaded subspace design, denoted $U(\mathcal{M})$, is a subspace of $F_q^{m_l}$ defined as follows: A vector $\mathbf{x} \in \mathbb{F}_q^{m_l}$ belongs to $U(\mathcal{M})$ if and only if for every $\iota \in \{1, 2, \ldots, l\}$, each of the m_ι-sized blocks of $\mathbf{x}$ given* $\mathrm{proj}_{[jm_\iota+1,(j+1)m_\iota]}(\mathbf{x})$ *for $0 \leqslant j < m_l/m_\iota$) belongs $H_1^{(\iota)} \times H_2^{(\iota)} \times \cdots \times H_{m_\iota/m_{\iota-1}}^{(\iota)}$.*

The following simple fact gives a lower bound on the dimension of a canonical subspace.

Observation 5.5. *For a cascaded subspace design $\mathcal{M}$ as above, if the ι'th subspace design has dimension at least $(1 - \xi_{\iota-1})m_{\iota-1}$ for $1 \leqslant \iota \leqslant l$, then the dimension of the canonical subspace $U(\mathcal{M})$ is at least $\Big(1 - (\xi_0 + \xi_1 + \cdots + \xi_{l-1})\Big)m_l$.*

The following is the crucial claim about pruning ultra-periodic subspaces using (the canonical subspace of) a cascaded subspace design. It generalizes Lemma 5.1 which corresponds to the $l = 1$ case.

Lemma 5.6. *Let $\mathcal{M}$ be a $(r_0, r_1, \ldots, r_l)$-cascaded subspace design with length-vector $(m_0, m_1, \ldots, m_l)$. Let T be a (r_0, m_0)-ultra periodic affine subspace of $\mathbb{F}_q^{m_l}$. Then the dimension of the affine space $T \cap U(\mathcal{M})$ is at most r_l.*

Proof. The idea will be to apply Lemma 5.1 inductively, for increasing periods $m_0, m_1, \ldots, m_{l-1}$. Since T is (r_0, m_0)-ultra periodic, it is (r_0, m_0)-periodic and $((m_1/m_0)r_0, m_1)$-periodic. Using this together with Lemma 5.1, it follows that

$$T \cap \{\mathbf{x} \in \mathbb{F}_q^{m_l} \mid \mathrm{proj}_{[jm_1+1,(j+1)m_1]}(\mathbf{x}) \in H_1^{(1)} \times H_2^{(1)} \times \cdots \times H_{m_1/m_0}^{(1)} \text{ for } 0 \leqslant j < m_l/m_1\}$$

is an affine subspace that is (r_1, m_1)-periodic. Continuing this argument, the affine subspace of T formed by restricting each m_ι-block to belong to $H_1^{(\iota)} \times H_2^{(\iota)} \times \cdots \times H_{m_\iota/m_{\iota-1}}^{(\iota)}$ for $1 \leqslant \iota \leqslant j$ is (r_j, m_j)-periodic. For $j = l$, we get the intersection $T \cap U(\mathcal{M}) \subset \mathbb{F}_q^{m_l}$ will be (r_l, m_l)-periodic, which simply means that it is an r_l-dimensional affine subspace of $\mathbb{F}_q^{m_l}$. □

We conclude this section by recording a good construction of a canonical subspace that has low-dimensional intersection with ultra-periodic subspaces. This statement will be used later in pre-coding algebraic-geometric codes based on the Garcia-Stichtenoth tower.

Theorem 5.7. *Let $\eta \in (0, 1)$, $q \geqslant 4$ be a prime power, and integers r, Λ satisfy $\Lambda \geqslant \Omega(1/\eta^2)$ and $r \leqslant \eta\Lambda/2$. For all large enough multiples κ of Λ, we can construct a subspace U of $\mathbb{F}_q^\kappa$ of dimension at least $(1 - \eta)\kappa$ such that for every (r, Λ)-ultra periodic affine subspace $T \subset \mathbb{F}_q^\kappa$, the dimension of the affine subspace $U \cap T$ is at most $r \cdot (1/\eta)^{O(\log^* \kappa)} 2^{O((\log^* \kappa)^2)}$. The subspace U can be constructed in deterministically in $\kappa^{O(\log_q^3 \kappa)}$ time.*

Proof. We will take U to the canonical subspace $U(\mathcal{M})$ of an appropriate cascaded subspace design $\mathcal{M}$. To this end, given our work so far, the main remaining task is to pick the parameters of $\mathcal{M}$ carefully. Let $\eta_\iota = \frac{\eta}{4 \cdot 2^\iota}$ for $\iota = 0, 1, 2, \ldots$. Let $m_0 = \Lambda$, and for $\iota \geqslant 0$, $m_{\iota+1} = m_\iota \cdot q^{\lceil \sqrt{m_\iota} \rceil}$. Let $r_0 = r$, and for $\iota \geqslant 0$, $r_{\iota+1} = \lfloor 8r_\iota/\eta_\iota \rfloor$. It is easy to check that $r_\iota \leqslant \eta_\iota m_\iota/2$ for all ι. So by Lemma 5.3, we can construct an

$(r_\iota, r_{\iota+1})$-subspace design in $\mathbb{F}_q^{m_\iota}$ of dimension $(1-\eta_\iota)m_\iota$ and cardinality at least $q^{\Omega(\eta_\iota m_\iota)}$. By our choice of parameters, $\eta_\iota m_\iota \gg \sqrt{m_\iota}$,[10] so we can construct such a subspace design of cardinality $q^{\lceil \sqrt{m_\iota} \rceil} = m_{\iota+1}/m_\iota$.

Pick l to the smallest integer so that $m_{l-1} \geqslant (\log_q \kappa)^2$. Since $m_0 = \Lambda \geqslant 2$ and $m_{\iota+1} \geqslant q^{\sqrt{m_\iota}}$ for $0 \leqslant \iota < l$, it is easy to see that that $l \leqslant O(\log^* \kappa)$ Redefine m_{l-1} to equal m'_{l-1} which is the smallest multiple of m_{l-2} that is at least $(\log_q \kappa)^2$. Since $m_{l-2} < (\log_q \kappa)^2$, we have $(\log_q \kappa)^2 \leqslant m'_{l-1} < 2(\log_q \kappa)^2$. We also redefine m_l to equal the largest multiple m'_l of m'_{l-1} that is at most κ. This implies $\kappa - m'_l < m'_{l-1}$. Note that $m'_{l-1} \leqslant m_{l-2} q^{\lceil \sqrt{m_{l-2}} \rceil}$ and $m'_l \leqslant q^{\sqrt{m'_{l-1}}}$. For notational simplicity, let us redenote m'_{l-1} and m'_l by m_{l-1} and m_l.

Thus for these parameters, we can construct an $(r_0, r_1, \ldots, r_l)$-cascaded subspace design $\mathcal{M}_l$ with length-vector $(m_0, m_1, \ldots, m_l)$ and dimension-vector $(d_0, d_1, \ldots, d_{l-1})$ where $d_\iota \geqslant (1 - \eta/2^{\iota+2})m_\iota$.

The construction time for subspace designs guaranteed by Lemma 5.3 implies that $\mathcal{M}_l$ can be constructed in deterministic $q^{O(m_{l-1}^2)} = q^{O(\log_q \kappa)^4}$ time. We define the desired subspace $U \subset \mathbb{F}_q^\kappa$ as $U(\mathcal{M}_l) \times 0^{\kappa - m_l}$, i.e., U consists of the vectors in the canonical subpsace $U(\mathcal{M}_l) \subset \mathbb{F}_q^{m_l}$ padded with $\kappa - m_l$ zeroes at the end. By Observation 5.5, the dimension of U is at least

$$\Big(1 - \sum_{\iota=0}^{l-1} \frac{\eta}{4 \cdot 2^\iota}\Big) m_l \geqslant (1 - \eta/2) m_l > (1 - \eta/2)(\kappa - m_{l-1}) > (1 - \eta/2)\kappa - 2(\log_q \kappa)^2 > (1 - \eta)\kappa$$

for large enough κ. This proves that the subspace U has dimension at least $(1 - \eta)\kappa$, and can be constructed deterministically in $q^{O(\log_q \kappa)^4}$ time.

It remains to prove the claimed intersection property with ultra-periodic subspaces. Let T be an arbitrary (r, Λ)-ultra periodic affine subspace of $\mathbb{F}_q^\kappa$. By Lemma 5.6, $\mathrm{proj}_{m_l}(T) \cap U(\mathcal{M})$ is an affine subspace of $\mathbb{F}_q^{m_l}$ of dimension at most r_l. Clearly, the same dimension bound also holds for $T \cap U$ since the last $\kappa - m_l$ coordinates for vectors in U are set to 0. The proof is complete by noting that for our choice of parameters, $r_l \leqslant r \cdot (1/\eta)^{O(l)} 2^{O(l^2)}$, and $l \leqslant O(\log^* \kappa)$. □

6 Pseudorandom subspace-evasive sets

For our code constructions, we will need to pre-code the messages into large subsets of $\mathbb{F}_q^\kappa$ that have small intersection with the sort of subspaces of the message space $\mathbb{F}_q^\kappa$ output by the linear-algebraic list decoder. We already saw one approach to accomplish this using subspace designs by exploiting the periodic nature of the subspaces we encounter. In this section, we will develop a different approach, again exploiting periodicity of the subspaces, that will lead to better parameters at the expense of settling for Monte Carlo constructions. We begin by recalling the notion of a *subspace-evasive* set [19], which is the most general object that has small intersection with subspaces of bounded dimension.

Definition 10. *For positive integer parameters r, ℓ, a set $S \subset \mathbb{F}_q^\kappa$ is (r, ℓ)-subspace evasive if for every affine subspace H of $\mathbb{F}_q^\kappa$ of dimension at most r, $|S \cap H| \leqslant \ell$.*

Let $\mathcal{F}$ be a family of affine subspaces of $\mathbb{F}_q^\kappa$ each of dimension at most r. A set $S \subset \mathbb{F}_q^\kappa$ is $(\mathcal{F}, r, \ell)$-subspace evasive if $|S \cap H| \leqslant \ell$ for every affine subspace $H \in \mathcal{F}$.

We now turn to a more specific notion of subspace evasiveness, tailored to periodic subspaces. Exploiting the fact the projections of a periodic subspace grow gradually in dimension, we will ensure that the

[10] We could be more careful in lower-bounding $\eta_\iota m_\iota$ here, but prefer the somewhat crude bound for simplicity.

subspace-evasive set also avoids large intersections with projections of the subspace on certain prefixes of the κ coordinates. This will enable the efficient computation of the intersection of the subspace-evasive set with the candidate periodic subspace.

Hierarchical subspace-evasive sets. We now define the special subspace-evasive sets that are useful for efficient pruning of candidate messages belonging to a (r, Λ, b)-periodic subspace. This notion is the same as the one from our previous work [20].

Definition 11. *Let $\mathcal{F}$ be a family of (r, Λ, b)-periodic subspaces of $\mathbb{F}_q^{b\Lambda}$, and $L \geqslant 1$ an integer. A subset $S \subset \mathbb{F}_q^{b\Lambda}$ is said to be $(\mathcal{F}, r, \Lambda, b, L)$-h.s.e (for hierarchically subspace evasive) if for every affine subspace $H \in \mathcal{F}$, the following bound holds for $j = 1, 2, \ldots, b$:*

$$|\mathrm{proj}_{j\Lambda}(S) \cap \mathrm{proj}_{j\Lambda}(H)| \leqslant L\ . \tag{17}$$

The following result is similar to Theorem 4.6 in [20]. The parameters claimed, most notably the encoding time to compute HSE, are somewhat different from those given in [20]. Therefore we briefly sketch the construction in Appendix B. The main difference is that, following [19, Sec. 4.1], we will define the h.s.e set based on the values of random polynomials rather than zero sets of random polynomials.

Lemma 6.1. *Suppose b, c, Λ, r are positive integers, $\kappa = b\Lambda$, q a prime power, and $\zeta \in (0, 1/4)$ satisfying the conditions $r < \zeta\Lambda/4$ and $c\kappa < q^r$. Let $\mathcal{F}$ be a family of (r, Λ, b)-periodic subspaces of $\mathbb{F}_q^\kappa$ with $|\mathcal{F}| \leqslant q^{c\kappa}$. Then there exists a randomized* $\mathrm{poly}(\kappa, 1/\zeta, \log q)$ *time construction of an injective map* $\mathsf{HSE} : \mathbb{F}_q^{(1-2\zeta)\kappa} \to \mathbb{F}_q^\kappa$ *with the following properties.*

1. *[Structure] Given $\mathbf{x} = (x_1, x_2, \ldots, x_b)$ with each $x_i \in \mathbb{F}_q^{(1-2\zeta)\Lambda}$, $\mathsf{HSE}(x)$ is of the form $x_1 \circ \xi_1 \circ x_2 \circ \xi_2 \circ \cdots \circ x_b \circ \xi_b$ where each $\xi_j \in \mathbb{F}_q^{2\zeta\Lambda}$ for $j = 1, 2, \ldots, b$ and depends only on the prefix $x_1 \circ x_2 \circ \cdots \circ x_j$.*
2. *[Computability] Computing ξ_j from $x_1 \circ x_2 \cdots \circ x_j$ can be done deterministic* $\mathrm{poly}(\kappa, 1/\zeta, \log q)$ *time.*[11]
3. *[Subspace-evasiveness] With high probability, the image of* HSE *is $(\mathcal{F}, r, \Lambda, b, c\kappa)$-h.s.e and $(\mathcal{F}, br, 2c/\zeta)$-subspace evasive as a subset of $\mathbb{F}_q^\kappa$. Further, given a (r, Λ, b)-periodic subspace $H \in \mathcal{F}$, one can compute the set $\{\mathbf{x} \in \mathbb{F}_q^{(1-2\zeta)\kappa} \mid \mathsf{HSE}(\mathbf{x}) \in H\}$ of size at most $2c/\zeta$ in deterministic* $\mathrm{poly}(\kappa, q^r, 1/\zeta)$ *time.*

The algorithm claimed above to compute the intersection of the h.s.e set with a periodic subspace H is iterative, and for each block, involves searching over all q^r extensions in H to check which ones belong in the (projection of) the h.s.e set. This search will lead to a large exponent in the polynomial runtime of our algorithm for Reed-Solomon codes. In the case of our decoding algorithm for Gabidulin codes, the field size will be too large and going over all q^r extensions will in fact take super-polynomial time. Therefore, we will use an additional encoding to ensure that the possible extensions will themselves be restricted to belong a (r, ℓ)-subspace evasive subset of $\mathbb{F}_q^\Lambda$, for which the intersection with an r-dimensional affine space can be computed in $\mathrm{poly}(\ell, \log q)$ time. Specifically, we will use the following result, based on a construction due to Dvir and Lovett [6], with a minor change in some degree parameters to handle efficient encoding for any field $\mathbb{F}_q$. A proof sketch appears in Appendix B.

[11] By construction, this immediately implies that the following tasks can also be done in the same time complexity: computing HSE; checking membership in $\mathsf{Im}(\mathsf{HSE})$ (the image of HSE); computing the inverse of HSE on its image; and for every $j = 1, 2, \ldots, b$, checking membership in $\mathrm{proj}_{j\Lambda}(\mathsf{Im}(\mathsf{HSE}))$ and computing the inverse prefix when it exists.

Lemma 6.2. *Let $\mathbb{F}_q$ be a finite field of characteristic p. Let υ, r be positive integers satisfying $r < \upsilon < q$. There is an explicit map $\psi : \mathbb{F}_q^{\upsilon-r} \to \mathbb{F}_q^{\upsilon}$, computable in deterministic* $\mathrm{poly}(\upsilon, \log q)$ *time, such that*

(i) ψ *is injective and further* $\mathrm{proj}_{\upsilon-r}(\psi(\mathbf{x})) = \mathbf{x}$ *for every* $\mathbf{x} \in \mathbb{F}_q^{(1-\eta)\upsilon}$.

(ii) The image of ψ *is a* $(r, (\max\{p\upsilon, p^r\})^r)$*-subspace evasive subset of* $\mathbb{F}_q^{\upsilon}$*. Further, given an affine subspace* $H \subseteq \mathbb{F}_q^{\upsilon}$ *of dimension at most* r*, we can compute the set* $\{\mathbf{x} \in \mathbb{F}_q^{\upsilon-r} \mid \psi(\mathbf{x}) \in H\}$ *in* $\mathrm{poly}(\upsilon^r, p^{r^2}, \log q)$ *time.*[12]

We will now combine the constructions of Lemmas 6.1 and 6.2 to get a construction of an h.s.e set which has the properties we need for efficient pruning of the subspace of candidate solutions produced by the list decoder. The key difference compared to Lemma 6.1 is the improvement in time needed to compute the intersection with a given periodic subspace H — instead of the q^r term, we now have Λ^r, which is a big savings for large fields (particularly in the case of Gabidulin codes). The proof can be found in Appendix B.

Theorem 6.3. *Suppose* b, c, Δ, r *are positive integers,* $k = b\Delta$*,* q *a power of a prime* p*, and* $\zeta \in (0, 1/6)$ *satisfying the conditions* $r < \zeta\Delta/4$ *and* $ck < q^r$*. Let* $\mathcal{F}$ *be a family of* (r, Δ, b)*-periodic subspaces of* $\mathbb{F}_q^k$ *with* $|\mathcal{F}| \leqslant q^{ck}$*. Then there exists a randomized* $\mathrm{poly}(k, 1/\zeta, \log q)$ *time construction of an injective map* $\widetilde{\mathsf{HSE}} : \mathbb{F}_q^{(1-3\zeta)k} \to \mathbb{F}_q^k$ *such that*

1. *One can compute* $\widetilde{\mathsf{HSE}}$ *in deterministic* $\mathrm{poly}(k, 1/\zeta, \log q)$ *time.*
2. *With high probability, the image of* $\widetilde{\mathsf{HSE}}$ *is* $(\mathcal{F}, br, 2c/\zeta)$*-subspace evasive as a subset of* $\mathbb{F}_q^k$*.*

   Further, given a (r, Δ, b)*-periodic subspace* $H \in \mathcal{F}$*, one can compute the set* $\{\mathbf{x} \in \mathbb{F}_q^{(1-3\zeta)k} \mid \widetilde{\mathsf{HSE}}(\mathbf{x}) \in H\}$ *of size at most* $2c/\zeta$ *in deterministic* $\mathrm{poly}(k, p^{r^2}, \Delta^r, 1/\zeta, \log q)$ *time.*

7 Good list decodable subcodes of RS and AG codes

We now combine our code constructions with a pre-coding step that restricts coefficients to belong to either a subspace design or a hierarchical subspace-evasive set, and thereby obtain subcodes that are list decodable with smaller list-size in polynomial time.

7.1 Reed-Solomon codes

We begin with the case of Reed-Solomon codes. For a finite field $\mathbb{F}_q$, constant $\varepsilon > 0$, integers n, k, m, s satisfying $1 \leqslant k < n \leqslant q$ and $1 \leqslant s \leqslant \varepsilon m/12$, we will define subcodes of $\mathrm{RS}^{(q,m)}[n, k]$. Below for a polynomial $f \in \mathbb{F}_{q^m}[X]$ with k coefficients $f_0, f_1, \ldots, f_{k-1}$, we denote by $\mathbf{f_0}, \mathbf{f_1}, \ldots, \mathbf{f_{k-1}}$ the representation of these coefficients as vectors in $\mathbb{F}_q^m$ by fixing some $\mathbb{F}_q$-basis of $\mathbb{F}_{q^m}$.

7.1.1 Subcode construction based on subspace designs

Define the subcode $\widehat{\mathsf{RS}}$ of $\mathrm{RS}^{(q,m)}[n, k]$ consisting of the encodings of $f \in \mathbb{F}_{q^m}[X]$ such that $(\mathbf{f_0}, \mathbf{f_1}, \ldots, \mathbf{f_{k-1}}) \in V$ for a subspace $V \subseteq \mathbb{F}_q^{mk}$ guaranteed by Theorem 5.4, when applied with the parameter choices

$$\Lambda = m; \quad b = k; \quad r = s - 1; \quad \eta = \varepsilon \,.$$

[12] If $\mathrm{char}(\mathbb{F}_q)$ is large, then we can use a different construction that replaces the p^{r^2} term by a $(\log q)^r$ term. For simplicity, we assume a choice of field with small characteristic is made.

Note that $\widehat{\mathsf{RS}}$ is an $\mathbb{F}_q$-linear code over the alphabet $\mathbb{F}_{q^m}$ of rate $(1-\varepsilon)k/n$, and it can be constructed in deterministic $q^{O(m^2)}$ time, or Las Vegas $q^{O(ms)}$ time.[13]

Theorem 7.1. *Given an input string $\mathbf{y} \in \mathbb{F}_{q^m}^n$, a basis of an affine subspace of dimension at most $O(s/\varepsilon)$ that includes all codewords of the above subcode within Hamming distance $\frac{s}{s+1}(n-k)$ from $\mathbf{y}$ can be found in deterministic $\mathrm{poly}(n, \log q, m)$ time.*

Proof. By Lemma 4.4, we can compute the $(s-1, m, k)$-periodic subspace T of messages whose Reed-Solomon encodings can be within Hamming distance $\frac{s}{s+1}(n-k)$ from $\mathbf{y}$. By Theorem 5.4, the intersection $T \cap V$ is is an affine subspace over $\mathbb{F}_q$ of dimension $d = O(s/\varepsilon)$. Since both steps involve only basic linear algebra, they can be accomplished using $\mathrm{poly}(n, m)$ operations over $\mathbb{F}_q$. □

By picking $s = \Theta(1/\varepsilon)$ and $m = \Theta(1/\varepsilon^2)$ in the above construction, we can conclude the following.

Corollary 7.2. *For every $R \in (0,1)$ and $\varepsilon > 0$, and all large enough integers $n < q$ with q a prime power, one can construct a rate R $\mathbb{F}_q$-linear subcode of a Reed-Solomon code of length n over $\mathbb{F}_{q^m}$, such that the code can be (i) encoded in $(n/\varepsilon)^{O(1)}$ time and (ii) list decoded from a fraction $(1-\varepsilon)(1-R)$ of errors in $(n/\varepsilon)^{O(1)}$ time, outputting a subspace over $\mathbb{F}_q$ of dimension $O(1/\varepsilon^2)$ including all closeby codewords. The code can be constructed deterministically in $q^{\varepsilon^{-O(1)}}$ time.*

We note that the above list decoding guarantee is in fact weaker than what is achieved for folded Reed-Solomon codes in [19], where the codewords were pinned down to a dimension $O(1/\varepsilon)$ subspace. We can improve the list size above to $\mathrm{poly}(1/\varepsilon)$ using pseudorandom subspace-evasive sets as in [19], or to $\exp(\varepsilon^{-O(1)})$ using the explicit subspace-evasive sets from [6]. The main point of the above result is not the parameters but that an explicit subcode of RS codes has optimal list decoding radius with polynomial complexity.

7.1.2 Subcode construction based on h.s.e sets

We now pre-code the messages of the RS code with an h.s.e set instead of subspace designs. This gives a much better list size, but we only get a randomized construction. Define the subcode of $\mathrm{RS}^{(q,m)}[n,k]$ consisting of the encodings of $f \in \mathbb{F}_{q^m}[X]$ such that $(\mathbf{f_0}, \mathbf{f_1}, \ldots, \mathbf{f_{k-1}}) = \widetilde{\mathsf{HSE}}(\mathbf{x})$ for some $\mathbf{x} \in \mathbb{F}_q^{(1-\varepsilon)mk}$ where $\widetilde{\mathsf{HSE}}$ is the (randomized) map guaranteed by Theorem 6.3 for parameters

$$\zeta = \varepsilon/3, \;\; \Delta = m, \;\; b = k, \text{ and } r = s-1 \,. \tag{18}$$

By definition, the above is a code of rate $(1-\varepsilon)k/n$ over the alphabet $\mathbb{F}_{q^m}$. It is also encodable in $\mathrm{poly}(n, m, \log q, 1/\varepsilon)$ time, since both $\widetilde{\mathsf{HSE}}$ and the Reed-Solomon encoding can be computed in this time. We now turn to the list decoding.

Theorem 7.3. *Given an input string $\mathbf{y} \in \mathbb{F}_{q^m}^n$, a list of size at most $O(1/(R\varepsilon))$ that includes all codewords of the above subcode within Hamming distance $\frac{s}{s+1}(n-k)$ from $\mathbf{y}$ can be found in deterministic $\mathrm{poly}(n, \log q, 1/\varepsilon, m^s, p^{s^2})$ time, where $p = \mathrm{char}(\mathbb{F}_q)$.*[14]

[13] It can also be constructed in Monte Carlo $(q/\varepsilon)^{O(1)}$ time by randomly picking subspaces for the subspace design used to construct V in Theorem 5.4.

[14] When p is large, the p^{s^2} term in the runtime can be replaced by $(\log q)^r$ using a different construction in Lemma 6.2, but we skip this variant for simplicity.

Proof. We first use Lemma 4.4 to compute the affine subspace $\mathcal{W}$ of messages whose Reed-Solomon encodings can be within Hamming distance $\frac{s}{s+1}(n-k)$ from $\mathbf{y}$. We note two crucial things about this space $\mathcal{W}$ of solutions: (i) it is an $(s-1,m,k)$-periodic subspace of $\mathbb{F}_q^{mk}$, and (ii) it belongs to a family $\mathcal{F}$ consisting of at most $q^{O(mn)}$ subspaces. The latter follows since the space of solutions is determined by the polynomials $A_0, A_1, \ldots, A_s$, and by our choice of degree parameter D, there are at most $q^{O(mn)} = q^{O(mk/R)}$ choices of those.

We now apply Theorem 6.3 with parameter $c = O(1/R)$ (and $b = k$, $\Delta = m$, $r = s-1$, and $\zeta = \varepsilon/3$ as chosen in (18)). Note that the condition $r < \zeta\Delta/4$ is satisfied by virtue of the requirement $s \leqslant \varepsilon m/12$. By Theorem 6.3, the intersection of the h.s.e set with $\mathcal{W}$ will have size at most $O(1/(R\varepsilon))$ and can be found in $\mathrm{poly}(n, \log q, 1/\varepsilon, m^s, p^{s^2})$ time. □

By picking $q \leqslant 2n$ to be a power of 2 and setting $s = \Theta(1/\varepsilon)$ and $m = \Theta(1/\varepsilon^2)$ in the above construction (so that the requirement $s \leqslant \varepsilon m/12$ is met), we can conclude the following.

Corollary 7.4. *For every $R \in (0,1)$ and $\varepsilon > 0$, there is a Monte Carlo construction of a rate R subcode of a Reed-Solomon code of length n over a field of characteristic 2 and size at most $n^{O(1/\varepsilon^2)}$ (with evaluation points in a subfield) that can be encoded in $(n/\varepsilon)^{O(1)}$ time and that with high probability can be list decoded from a fraction $(1-\varepsilon)(1-R)$ of errors in deterministic $\mathrm{poly}(n, \exp(1/\varepsilon^2))$ time, outputting a list of size at most $O(1/(R\varepsilon))$.*

7.2 Subcodes of Garcia-Stichtenoth codes

We now pre-code the codes constructed in Section 4.3. For a finite field $\mathbb{F}_q$, constant $\varepsilon > 0$, and integers s, m satisfying $1 \leqslant s \leqslant \varepsilon m/12$ and $m \geqslant \Omega(1/\varepsilon^2)$, we will define subcodes of $\mathrm{GS}^{(q,m)}[N,k]$ guaranteed by Theorem 4.14. Note that messages space of this code can be identified with $\mathbb{F}_q^{mk}$.

7.2.1 Subcode construction based on cascaded subspace designs

Define the subcode $\widehat{GS}$ of $\mathrm{GS}^{(q,m)}[N,k]$ consisting of the encodings of a subspace $U \subseteq \mathbb{F}_q^{mk}$ guaranteed by Theorem 5.7, when applied with the parameter choices

$$\eta = \varepsilon; \quad r = s-1; \quad \Lambda = m; \quad \kappa = km\,. \tag{19}$$

Note that $\widehat{GS}$ is an $\mathbb{F}_q$-linear code over the alphabet $\mathbb{F}_{q^m}$ of rate $(1-\varepsilon)k/N$. Also, it can be constructed in deterministic $(km)^{O(\log_q^3(km))}$ time by virtue of the construction complexity of U.

Lemma 7.5. *Given an input string $\mathbf{y} \in \mathbb{F}_{q^m}^N$, a basis of an affine subspace of dimension at most*

$$s \cdot (1/\varepsilon)^{O(\log^*(km))} 2^{O((\log^*(km))^2)}$$

that includes all codewords of the above subcode within Hamming distance $\frac{s}{s+1}(N-k) - 3N/(\sqrt{q}-1)$ from $\mathbf{y}$ can be found in deterministic $\mathrm{poly}(n, \log q, m)$ time.

Proof. By Theorem 4.14, we can compute the $(s-1,m)$-ultra periodic subspace T of messages whose encodings can be within Hamming distance $\frac{s}{s+1}(N-k) - 3N/(\sqrt{q}-1)$ from $\mathbf{y}$. By Theorem 5.7, for the above choice of parameters (19), the intersection $T \cap U$ is is an affine subspace over $\mathbb{F}_q$ of dimension $s \cdot (1/\varepsilon)^{O(\log^*(km))} 2^{O((\log^*(km))^2)}$. Since both steps involve only basic linear algebra, they can be accomplished using $\mathrm{poly}(N,m)$ operations over $\mathbb{F}_q$. □

By taking $q = \Theta(1/\varepsilon^2)$, and choosing $s = \Theta(1/\varepsilon)$ and $m = \Theta(1/\varepsilon^2)$ in the above lemma, we conclude the following main result concerning deterministic construction of codes list decodable up to the Singleton bound.

Theorem 7.6 (Main deterministic code construction)**.** *For every $R \in (0,1)$ and $\varepsilon > 0$, for $q = \Theta(1/\varepsilon^2)$, we can construct an $\mathbb{F}_q$-linear family of codes of rate R over an alphabet of size $q^{O(\varepsilon^{-2})}$ such that a code of block length N in the family can be (i) encoded in $(N/\varepsilon)^{O(1)}$ time, and (ii) list decoded from a fraction $(1 - R - \varepsilon)$ of errors in $(N/\varepsilon)^{O(1)}$ time, outputting a subspace over $\mathbb{F}_q$ of dimension at most $2^{O((\log^* N)^2)}$ (for large enough N) that includes all closeby codewords. The code can be constructed* deterministically *in $N^{O(\log_q^3 N)}$ time.*

7.2.2 Subcode construction based on h.s.e. sets

We now pre-code the messages of the Garcia Stichtenoth code with an h.s.e set. This gives a much better list size, but we only get a randomized construction. Define the subcode of $\mathrm{GS}^{(q,m)}[N,k]$ consisting of the encodings of messages in $\mathbb{F}_q^{mk}$ which are of the form $\widetilde{\mathsf{HSE}}(\mathbf{x})$ for some $\mathbf{x} \in \mathbb{F}_q^{(1-\varepsilon)mk}$ where $\widetilde{\mathsf{HSE}}$ is the (randomized) map guaranteed by Theorem 6.3 for parameters

$$\zeta = \varepsilon/3, \;\; \Delta = m, \;\; b = k, \text{ and } r = s - 1 \; . \tag{20}$$

By definition, the above is a code of rate $(1 - \varepsilon)k/n$ over the alphabet $\mathbb{F}_{q^m}$. It is also encodable in $\mathrm{poly}(N, m, \log q, 1/\varepsilon)$ time, since both $\widetilde{\mathsf{HSE}}$ and the encoding into $\mathrm{GS}^{(q,m)}[N,k]$ can be computed in this time. We now turn to the list decoding.

Lemma 7.7. *Given an input string $\mathbf{y} \in \mathbb{F}_{q^m}^N$, a list of size at most $O(1/(R\varepsilon))$ that includes all codewords of the above subcode within Hamming distance $\frac{s}{s+1}(N - k) - 3N/(\sqrt{q} - 1)$ from $\mathbf{y}$ can be found in deterministic* $\mathrm{poly}(N, \log q, 1/\varepsilon, m^s, p^{s^2})$ *time, where $p = \mathrm{char}(\mathbb{F}_q)$.*

Proof. By Theorem 4.14, we can compute an affine subspace $\mathcal{W} \subset \mathbb{F}_q^{mk}$ of messages whose encodings lie within Hamming distance $\frac{s}{s+1}(N-k) - 3N/(\sqrt{q}-1)$ from $\mathbf{y}$. The rest of the proof is identical to Theorem 7.3. We recall two crucial things about this space $\mathcal{W}$ of solutions: (i) it is an $(s-1, m, k)$-periodic subspace of $\mathbb{F}_q^{mk}$, and (ii) it belongs to a family $\mathcal{F}$ consisting of at most $q^{O(mN)}$ subspaces.

We now apply Theorem 6.3 with parameter $c = O(1/R)$ (and $b = k$, $\Delta = m$, $r = s - 1$, and $\zeta = \varepsilon/3$ as chosen in (20)). Note that the condition $r < \zeta\Delta/4$ is satisfied by virtue of the requirement $s \leqslant \varepsilon m/12$. By Theorem 6.3, the intersection of the h.s.e set with $\mathcal{W}$ will have size at most $O(1/(R\varepsilon))$ and can be found in $\mathrm{poly}(N, \log q, 1/\varepsilon, m^s, p^{s^2})$ time. □

By picking q to be a power of 2 and setting $s = \Theta(1/\varepsilon)$ and $m = \Theta(1/\varepsilon^2)$ in the above construction, we can conclude the following. The parameters of this corollary match our earlier result for folded algebraic-geometric codes from [20].

Theorem 7.8. *For every $R \in (0,1)$ and $\varepsilon > 0$, there is a Monte Carlo construction of a rate R subcode of an algebraic-geometric code of length N over a field of characteristic 2 and size at most $(1/\varepsilon)^{O(1/\varepsilon^2)}$ that can be encoded in $(N/\varepsilon)^{O(1)}$ time and that with high probability can be list decoded from a fraction $(1 - R - \varepsilon)$ of errors in deterministic* $\mathrm{poly}(N, \exp(1/\varepsilon^2))$ *time, outputting a list of size at most $O(1/(R\varepsilon))$.*

8 Linear-algebraic list decoding of Gabidulin codes

First we briefly review rank-metric codes that were introduced by Gabidulin in [8]. Let h be a prime power and denote by $\mathbb{M}_{n\times t}(\mathbb{F}_h)$ the set of $n \times t$ matrices over $\mathbb{F}_h$. For two matrices $A, B \in \mathbb{M}_{n\times t}(\mathbb{F}_h)$, we define the rank distance between A and B, denoted by $d(A, B)$, by the rank of $A-B$, i.e., $d(A, B) = \mathrm{rank}(A-B)$. A rank-metric code $\mathcal{C}$ is a subset of $\mathbb{M}_{n\times t}(\mathbb{F}_h)$. The minimum distance and the rate of a rank-metric code $\mathcal{C}$ are defined as

$$d(\mathcal{C}) := \min\{d(A, B) : \ A \neq B \in \mathcal{C}\}, \quad \text{and} \quad R(\mathcal{C}) := \frac{\log_h |\mathcal{C}|}{nt},$$

respectively. Such a rank-metric code always satisfies the Singleton bound $d \leqslant n - R(\mathcal{C})n + 1$. A rank-metric code achieving the Singleton bound is called a maximum-rank-distance code (or MRD code for short). Linearized polynomials can be used to construct a class of MRD codes, i.e., the *Gabidulin codes*. Let us briefly recall this construction.

For a nonnegative integer i, we denote X^{h^i} by $X^{[i]}$. A h-linearized polynomial over $\mathbb{F}_{h^t}$ is a polynomial of the form $\sum_{i=0}^{\ell} a_i X^{[i]}$, where $a_i \in \mathbb{F}_{h^t}$. The integer ℓ is called the h-degree of this polynomial, denoted by $\deg_h(f)$, if $a_\ell \neq 0$. It is easy to verify that $(X^{[i]})^{[j]} = X^{[i+j]}$. We denote by $\mathcal{L}_h(t)$ the set of h-linearized polynomials over $\mathbb{F}_{h^t}$.

Let $0 < k \leqslant n \leqslant t$ be three integers and choose n elements $\alpha_1, \alpha_2, \ldots, \alpha_n \in \mathbb{F}_{h^t}$ such that they are linearly independent over $\mathbb{F}_h$. For every h-linearized polynomial $f \in \mathbb{F}_{h^t}[X]$ with the h-degree at most $k-1$, we get a column vector $M_f := (f(\alpha_1), f(\alpha_2), \ldots, f(\alpha_n))^T$ over $\mathbb{F}_{h^t}$. Since every coordinate in M_f is $\mathbb{F}_h$-linearly mapped to a row vector in $\mathbb{F}_h^t$ under a fixed $\mathbb{F}_h$-basis of $\mathbb{F}_{h^t}$, the vector M_f can be viewed as an $n \times t$ matrix over $\mathbb{F}_h$. Thus, we get a rank-metric code

$$\mathcal{C}_G(h; n, t, k) := \{M_f \in \mathbb{M}_{n\times t}(\mathbb{F}_h) \mid f \in \mathcal{L}_h(t), \deg_h(f) \leqslant k-1\}\,.$$

It can be shown that this code is in fact the Gabidulin code [8, 22]. Moreover, it is easy to verify that $\mathcal{C}_G(h; n, t, k)$ is an MRD code meeting the Singleton bound for rank-metric codes.

Before list decoding rank-metric codes, let us define the notion of “errors” first. Assume that a codeword matrix X is transmitted and a matrix Y is received. We say that there are e errors from X to Y if $\mathrm{rank}(X - Y) = e$. We denote by $\langle X\rangle$ and $\langle Y\rangle$ the $\mathbb{F}_h$-spaces spanned by the rows of X and Y, respectively. Then we have the following fact.

Lemma 8.1. *Let* $X, Y \in \mathbb{M}_{n\times t}(\mathbb{F}_h)$ *with* $\mathrm{rank}(X - Y) \leqslant e$. *Then* $\dim_{\mathbb{F}_h}(\langle X\rangle \cap \langle Y\rangle) \geqslant \dim_{\mathbb{F}_h}(\langle X\rangle) - e$.

Proof. First we observe that the two $\mathbb{F}_h$-spaces $\langle X\rangle + \langle Y\rangle$ and $\langle X - Y\rangle + \langle Y\rangle$ are equal. Thus,

$$\dim_{\mathbb{F}_h}(\langle X\rangle)+\dim_{\mathbb{F}_h}(\langle Y\rangle)-\dim_{\mathbb{F}_h}(\langle X\rangle\cap\langle Y\rangle) = \dim_{\mathbb{F}_h}(\langle X-Y\rangle)+\dim_{\mathbb{F}_h}(\langle Y\rangle)-\dim_{\mathbb{F}_h}(\langle X-Y\rangle\cap\langle Y\rangle).$$

This gives

$$\dim_{\mathbb{F}_h}(\langle X\rangle \cap \langle Y\rangle) = \dim_{\mathbb{F}_h}(\langle X\rangle) - \dim_{\mathbb{F}_h}(\langle X - Y\rangle) + \dim_{\mathbb{F}_h}(\langle X - Y\rangle \cap \langle Y\rangle) \geqslant \dim_{\mathbb{F}_h}(\langle X\rangle) - e\,,$$

completing the proof. □

To list decode the above Gabidulin code $\mathcal{C}_G(h; n, t, k)$ in a manner similar to the linear-algebraic algorithm for Reed-Solomon codes from Section 4, we need to impose one condition on the parameters, namely, $n|t$. This will ensure that $\mathbb{F}_{h^t}$ has a subfield $\mathbb{F}_{h^n}$. We can now choose n evaluations points $\alpha_1, \alpha_2, \ldots, \alpha_n$

that are linearly independent over $\mathbb{F}_h$ from the subfield $\mathbb{F}_{h^n}$. Put $q = h^n$ and $m = \frac{t}{n}$, so that $\mathbb{F}_{h^n} = \mathbb{F}_q$ and $\mathbb{F}_{h^t} = \mathbb{F}_{q^m}$. In this case, we also denote $\mathcal{C}_G(h; n, t, k)$ by $\mathcal{C}_G(q; n, m, k)$. Suppose that a codeword $M_f = (f(\alpha_1), f(\alpha_2), \dots, f(\alpha_n))^T$ is transmitted and $Y = (y_1, y_2, \dots, y_n)^T$ is received with at most e errors (note that we identify every row vector in Y with an element y_i in $\mathbb{F}_{h^t} = \mathbb{F}_{q^m}$).

Interpolation step. Let $1 \leqslant s \leqslant m$ be an integer parameter of the algorithm. Choose the "degree parameter" $D = \left\lfloor \frac{n-k+1}{s+1} \right\rfloor$.

Definition 12 (Space of interpolation linearized polynomials)**.** *Let $\mathcal{L}$ be the space of polynomials $Q \in \mathbb{F}_{q^m}[X, Y_1, Y_2, \dots, Y_s]$ of the form $Q(X, Y_1, Y_2, \dots, Y_s) = A_0(X) + A_1(Y_1) + A_2(Y_2) + \cdots + A_s(Y_s)$, with each $A_i \in \mathbb{F}_{q^m}[X]$ being an h-linearized polynomial and $\deg_h(A_0) \leqslant D + k - 1$ and $\deg_h(A_i) \leqslant D$ for $i = 1, 2, \dots, s$.*

The following lemma is similar to Lemma 4.1.

Lemma 8.2. *There exists a nonzero polynomial $Q \in \mathcal{L}$ such that $Q(\alpha_i, y_i, y_i^q, y_i^{q^2}, \cdots, y_i^{q^{s-1}}) = 0$ for $i = 1, 2, \dots, n$. Further such a Q can be found using $O(n^3)$ operations over $\mathbb{F}_{q^m}$.*

Proof. Note that $\mathcal{L}$ is an $\mathbb{F}_{q^m}$-vector space of dimension $(D + k) + s(D + 1) = (D + 1)(s + 1) + k - 1$. This dimension is bigger than n by our choice of D. The conditions imposed by the Lemma amount to n homogeneous linear conditions on Q. Since this is smaller than the $\mathbb{F}_{q^m}$-dimension of $\mathcal{L}$, there must exist a nonzero $Q \in \mathcal{L}$ that meets the interpolation conditions $Q(\alpha_i, y_i, y_i^q, y_i^{q^2}, \cdots, y_i^{q^{s-1}}) = 0$ for $i = 1, 2, \dots, n$. Finding such a Q amounts to solving a homogeneous linear system over $\mathbb{F}_{q^m}$ with n constraints and at most $\dim_{\mathbb{F}_{q^m}}(\mathcal{L}) \leqslant n + s + 2$ unknowns, which can be done in $O(n^3)$ time. □

Lemma 8.3 below shows that any polynomial Q given by Lemma 8.2 yields an algebraic condition that the message functions f we are interested in list decoding must satisfy. Recall that for a polynomial $f(X) = f_0 + f_1 X + \cdots + f_{k-1}X^{k-1}$, f^σ denotes the polynomial $f^\sigma(X) = f_0^q + f_1^q X + \cdots + f_{k-1}^q X^{k-1}$.

Lemma 8.3. *Let $f \in \mathbb{F}_{q^m}[X]$ be a h-linearized polynomial with h-degree at most $k - 1$. Suppose that a codeword $M_f = (f(\alpha_1), f(\alpha_2), \dots, f(\alpha_n))^T$ is transmitted and $Y = (y_1, y_2, \dots, y_n)^T$ is received with at most e errors. If $e \leqslant s(n - k)/(s + 1)$, then $Q(X, f(X), f^\sigma(X), f^{\sigma^2}(X), \cdots, f^{\sigma^{s-1}}(X)) = 0$.*

Proof. The polynomial $f(X)$ defines an $\mathbb{F}_h$-linear map from $\mathbb{F}_{h^n} = \mathbb{F}_q$ to $\mathbb{F}_{q^m}$. The kernel $\ker(f)$ has dimension $\dim_{\mathbb{F}_h}(\ker(f)) = \dim_{\mathbb{F}_h}(\mathbb{F}_q) - \operatorname{rank}(M_f) = n - \operatorname{rank}(M_f)$.
Consider the $\mathbb{F}_h$-vector space

$$f^{-1}(\langle Y \rangle) = \{\alpha \in \mathbb{F}_q : \ f(\alpha) \in \langle Y \rangle\}.$$

We claim that the $\mathbb{F}_h$-dimension of $f^{-1}(\langle Y \rangle)$ is at least $n - e$. To see this, we make use of the identity

$$\dim_{\mathbb{F}_h}(f^{-1}(\langle Y \rangle)) = \dim_{\mathbb{F}_h}(\ker(f)) + \dim_{\mathbb{F}_h}(\operatorname{Im}(f) \cap \langle Y \rangle).$$

Now our claim follows from the fact that $\operatorname{Im}(f) = \langle M_f \rangle$ and Lemma 8.1.
It is clear that for all $\alpha \in f^{-1}(\langle Y \rangle)$, we have

$$0 = Q(\alpha, f(\alpha), f(\alpha)^q, \dots, f(\alpha)^{q^{s-1}}) = Q(X, f, f^\sigma, \dots, f^{\sigma^{s-1}})(\alpha) \ .$$

As the h-degree of $Q(X, f, f^\sigma, \dots, f^{\sigma^{s-1}})$ is at most $D + k - 1$, under the condition

$$D + k - 1 < n - e, \tag{21}$$

we have

$$Q(X, f, f^{\sigma}, \dots, f^{\sigma^{s-1}}) = A_0(X) + A_1(f(X)) + A_2(f^{\sigma}(X)) + \cdots + A_s(f^{\sigma^{s-1}}(X)) = 0. \tag{22}$$

This completes the proof since (21) is indeed satisfied by our given condition on e and choice of $D = \lfloor \frac{n-k+1}{s+1} \rfloor$. □

Finding candidate solutions. As in the case of Reed-Solomon codes, we want to study the structure of a linearized polynomial f satisfying the condition of Lemma 8.3.

Lemma 8.4. *The set of solutions $f = \sum_{i=0}^{k-1} f_i X^{h^i} \in \mathbb{F}_{q^m}[X]$ to the equation*

$$A_0(X) + A_1(f(X)) + A_2(f^{\sigma}(X)) + \cdots + A_s(f^{\sigma^{s-1}}(X)) = 0 \tag{23}$$

when at least one of $\{A_0, A_1, \dots, A_s\}$ is nonzero is an $(s-1, m, k)$-periodic subspace. A canonical representation of this periodic subspace (in the sense of Definition 2) can be computed in $\mathrm{poly}(k, m, \log q)$ time.

We can mimic the proof of Lemma 4.4 to prove Lemma 8.4. We skip the details.

Combining Lemmas 8.3 and 8.4, we see that, as in the Reed-Solomon case, one can find an $\mathbb{F}_q$-affine space of dimension $(s-1)k$ that contains the coefficients of all polynomials whose encodings differ from the input $(y_1, \dots, y_n)$ by a matrix of rank at most $\frac{s}{s+1}(1-R)n$ (where R is the rate). When $s \ll m$, this dimension is much smaller than the dimension of the message space of the Gabidulin code $\mathcal{C}_G(q; n, m, k)$ over $\mathbb{F}_q$ which is km.

8.1 List-size reduction using h.s.s sets

The list decoding guarantee proved for Gabidulin codes in Lemma 8.4 is identical to the Reed-Solomon case (Lemma 4.4). Therefore, we can similarly use the h.s.e construction of Theorem 6.3 to randomly construct an efficiently list decodable subcode of the Gabidulin code. Recall that in our Gabidulin construction $\mathcal{C}_G(h; n, t, k)$, $q = h^n$ and $m = n/t$. Also $\mathrm{char}(\mathbb{F}_q) = \mathrm{char}(\mathbb{F}_h)$. We can thus state the following result paralleling Corollary 7.4 (we omit the analogous statement to Theorem 7.3 for brevity).

Theorem 8.5. *Let $\mathbb{F}_h$ be a finite field, $R \in (0,1)$, $\varepsilon > 0$, and n a large enough integer. For $k = \frac{n}{R(1-\varepsilon)}$ and $t = \Theta(n/\varepsilon^2)$, there is a Monte Carlo construction of a rate R subcode of the Gabidulin code $\mathcal{C}_G(h; n, t, k)$ (whose elements are $n \times t$ matrices over $\mathbb{F}_h$) such that the subcode*

(i) admits encoding in $\mathrm{poly}(n, \log h, 1/\varepsilon)$ time, and

(ii) offers the following list decoding guarantee with high probability: given a matrix $M \in \mathbb{M}_{n\times t}(\mathbb{F}_h)$, the set of codeword matrices within rank distance $(1-\varepsilon)(1-R)n$ from M has cardinality at most $O(1/(R\varepsilon))$ and can be found in deterministic $\mathrm{poly}(n, \log h, \mathrm{char}(\mathbb{F}_h)^{1/\varepsilon^2})$ time.

References

[1] A. Ben-Aroya and I. Shinkar. A note on subspace evasive sets. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:95, 2012. 1

[2] E. Ben-Sasson, S. Kopparty, and J. Radhakrishnan. Subspace polynomials and limits to list decoding of Reed-Solomon codes. *IEEE Transactions on Information Theory*, 56(1):113–120, 2010. 4

[3] D. Bleichenbacher, A. Kiayias, and M. Yung. Decoding interleaved reed-solomon codes over noisy channels. *Theor. Comput. Sci.*, 379(3):348–360, 2007. 3

[4] D. Coppersmith and M. Sudan. Reconstructing curves in three (and higher) dimensional spaces from noisy data. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 136–142, June 2003. 3

[5] P. Delsarte. Bilinear forms over a finite field, with applications to coding theory. *J. Comb. Theory, Ser. A*, 25(3):226–241, 1978. 4

[6] Z. Dvir and S. Lovett. Subspace evasive sets. In *Proceedings of the 44th ACM Symposium on Theory of Computing*, pages 351–358, 2012. 1, 2, 6, 23, 25, 36, 37

[7] J.-C. Faugère, P. M. Gianni, D. Lazard, and T. Mora. Efficient computation of zero-dimensional gröbner bases by change of ordering. *J. Symb. Comput.*, 16(4):329–344, 1993. 37

[8] E. M. Gabidulin. Theory of codes with maximal rank distance. *Problems of Information Transmission*, 21(7):1–12, 1985. 3, 4, 28

[9] E. M. Gabidulin. A fast matrix decoding algorithm for rank-error-correcting codes. In G. D. Cohen, S. Litsyn, A. Lobstein, and G. Zémor, editors, *Algebraic Coding*, volume 573 of *Lecture Notes in Computer Science*, pages 126–133. Springer, 1991. 4

[10] E. M. Gabidulin, A. V. Paramonov, and O. V. Tretjakov. Ideals over a non-commutative ring and their applications in cryptology. In D. W. Davies, editor, *EUROCRYPT*, volume 547 of *Lecture Notes in Computer Science*, pages 482–489. Springer, 1991. 4

[11] A. Garcia and H. Stichtenoth. A tower of Artin-Schreier extensions of function fields attaining the Drinfeld-Vlădut bound. *Inventiones Mathematicae*, 121:211–222, 1995. 6, 15

[12] A. Garcia and H. Stichtenoth. On the asymptotic behavior of some towers of function fields over finite fields. *Journal of Number Theory*, 61(2):248–273, 1996. 6, 15

[13] V. Guruswami. Cyclotomic function fields, Artin-Frobenius automorphisms, and list error-correction with optimal rate. *Algebra and Number Theory*, 4(4):433–463, 2010. 2

[14] V. Guruswami, S. Narayanan, and C. Wang. List decoding subspace codes from insertions and deletions. In *Proceedings of Innovations in Theoretical Computer Science (ITCS 2012)*, pages 183–189, January 2012. 4, 34, 35

[15] V. Guruswami and A. Patthak. Correlated Algebraic-Geometric codes: Improved list decoding over bounded alphabets. *Mathematics of Computation*, 77(261):447–473, 2008. 16

[16] V. Guruswami and A. Rudra. Limits to list decoding Reed-Solomon codes. *IEEE Transactions on Information Theory*, 52(8):3642–3649, August 2006. 1

[17] V. Guruswami and A. Rudra. Explicit codes achieving list decoding capacity: Error-correction with optimal redundancy. *IEEE Transactions on Information Theory*, 54(1):135–150, 2008. 1, 2, 3, 4, 5

[18] V. Guruswami and M. Sudan. Improved decoding of Reed-Solomon and Algebraic-geometric codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, 1999. 1

[19] V. Guruswami and C. Wang. Linear-algebraic list decoding for variants of Reed-Solomon codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:73, 2012. 1, 2, 4, 5, 6, 22, 23, 25, 35

[20] V. Guruswami and C. Xing. Folded codes from function field towers and improved optimal rate list decoding. *CoRR*, abs/1204.4209, 2012. Extended abstract appeared in the Proceedings of the 44th ACM Symposium on Theory of Computing (STOC'12). 1, 2, 3, 6, 7, 16, 23, 27, 35, 36

[21] J. W. P. Hirschfield. *Projective Geometries over Finite Fields*. Oxford Univ. Press, 1979. 7

[22] T. Johansson. Authentication codes for nontrusting parties obtained from rank metric codes. *Des., Codes Cryptogr*, 6:205–218, 1995. 28

[23] R. Koetter and F. R. Kschischang. Coding for errors and erasures in random network coding. *IEEE Transactions on Information Theory*, 54(8):3579–3591, 2008. 4, 33

[24] S. Kopparty. List-decoding multiplicity codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:44, 2012. 1, 2

[25] D. Lazard. Solving zero-dimensional algebraic systems. *J. Symb. Comput.*, 13(2):117–132, 1992. 37

[26] P. Loidreau. A Welch-Berlekamp like algorithm for decoding Gabidulin codes. In Ø. Ytrehus, editor, *WCC*, volume 3969 of *Lecture Notes in Computer Science*, pages 36–45. Springer, 2005. 4

[27] P. Loidreau. Designing a rank metric based McEliece cryptosystem. In N. Sendrier, editor, *PQCrypto*, volume 6061 of *Lecture Notes in Computer Science*, pages 142–152. Springer, 2010. 4

[28] P. Lusina, E. M. Gabidulin, and M. Bossert. Maximum rank distance codes as space-time codes. *IEEE Transactions on Information Theory*, 49(10):2757–2760, 2003. 4

[29] H. Mahdavifar and A. Vardy. Algebraic list-decoding on the operator channel. In *Proceedings of the IEEE International Symposium on Information Theory*, pages 1193–1197, 2010. 4, 34

[30] H. Mahdavifar and A. Vardy. List-decoding of subspace codes and rank-metric codes up to Singleton bound. *CoRR*, abs/1202.0866, 2012. 4, 33, 35

[31] F. Parvaresh and A. Vardy. Correcting errors beyond the Guruswami-Sudan radius in polynomial time. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 285–294, 2005. 4

[32] W. W. Peterson. Encoding and error-correction procedures for Bose-Chaudhuri codes. *IEEE Transactions on Information Theory*, 6:459–470, 1960. 1

[33] G. Richter and S. Plass. Error and erasure decoding of rank codes with modified Berlekamp-Massey algorithm. In *5th Int. ITG Conference on Source and Channel Coding (SCC'04)*, 2004. 4

[34] R. M. Roth. Maximum-rank array codes and their application to crisscross error correction. *IEEE Transactions on Information Theory*, 37(2):328–336, 1991. 4

[35] K. Shum, I. Aleshnikov, P. V. Kumar, H. Stichtenoth, and V. Deolalikar. A low-complexity algorithm for the construction of algebraic-geometric codes better than the Gilbert-Varshamov bound. *IEEE Transactions on Information Theory*, 47(6):2225–2241, 2001. 16

[36] D. Silva and F. R. Kschischang. Fast encoding and decoding of Gabidulin codes. In *Proceedings of the IEEE International Symposium on Information Theory*, 2009. Available at http://arxiv.org/abs/0901.2483. 4

[37] D. Silva, F. R. Kschischang, and R. Koetter. A rank-metric approach to error control in random network coding. *IEEE Transactions on Information Theory*, 54(9):3951–3967, 2008. 4, 33

[38] H. Stichtenoth. *Algebraic Function Fields and Codes*. Universitext, Springer-Verlag, Berlin, 1993. 12

[39] M. Sudan. Decoding of Reed-Solomon codes beyond the error-correction bound. *Journal of Complexity*, 13(1):180–193, 1997. 1

[40] A. Wachter-Zeh. Bounds on list decoding Gabidulin codes. *CoRR*, abs/1205.0345, 2012. 4

[41] H. Wang, C. Xing, and R. Safavi-Naini. Linear authentication codes: Bounds and constructions. *IEEE Transactions on Information Theory*, 49(4):866–872, 2003. 4, 33

A List decoding of subspace codes up to Singleton bound

We now briefly sketch the extension of our results from rank-metric codes to the closely related *subspace codes* (also called finite field Grassmanian codes). A subspace code $\mathcal{S}$ over finite field $\mathbb{F}_h$ is a collection of n-dimensional subspaces of $\mathbb{F}_q^N$ for integer parameters $n < N$. The distance property of the code is that any two subspaces $U, V \in \mathcal{S}$ have a low-dimensional intersection (or equivalently, the distance between U and V in the Grassmanian metric, defined as $\dim(U) + \dim(V) - 2\dim(U \cap V)$, is large). The rate of such a subspace code $\mathcal{S}$ is defined to be

$$\frac{\log_h |\mathcal{S}|}{nN}$$

(the rationale is that there are about h^{nN} subspaces of dimension n in $\mathbb{F}_h^N$ when n is much smaller than N).

A.1 Subspace code construction

The construction of subspace codes by Koetter and Kschischang [23] is equivalent to the one given earlier in [41]. We call these subspace codes as KK codes for short. The relation between the KK subspace codes and Gabidulin codes have been discussed in several papers [23, 37, 30]. In this section, we briefly look at the list decoding of the KK codes by employing the same idea that we used for list decoding of Reed-Solomon and Gabidulin codes.

As in the case of the Gabidulin codes, we assume that n is a divisor of t and let $q = h^n$ and $m = t/n$. Let $\alpha_1, \dots, \alpha_n \in \mathbb{F}_q$ be $\mathbb{F}_h$-linearly independent. Let Z be the $\mathbb{F}_h$-vector space $\mathbb{F}_q \times \mathbb{F}_{q^m}$. It is clear that $\dim_{\mathbb{F}_h}(Z) = n + mn = n + t$.

Choose an integer k with $0 < k < n$ and for each h-linearized polynomial $f \in \mathbb{F}_{q^m}[X]$ of h-degree at most $k - 1$, define an $\mathbb{F}_h$-subspace of Z by

$$V_f = \mathrm{span}\{(\alpha_i, f(\alpha_i)) :\ i = 1, 2, \dots, n\}.$$

The KK-code is defined as the collection of V_f for all h-linearized polynomials $f \in \mathbb{F}_{q^m}[X]$ of h-degree at most $k-1$. Note that the rate of this code equals

$$\frac{\log_h q^{mk}}{n(n+t)} = \frac{k}{n}\left(\frac{1}{1+n/t}\right) \approx k/n \quad (\text{when } n \ll t)\,. \tag{24}$$

A.2 List decoding KK-codes

We now present a list decoding algorithm for the above codes. The algorithm follows the earlier linear-algebraic list decoding algorithm for Gabidulin codes.

Before describing the algorithm, we first need to define the error model. The error level will be quantified by two integer parameters: (i) ρ, the maximum number of *insertions* allowed, and (ii) μ, the maximum number of *deletions* allowed.[15]

Suppose a codeword V_f encoded from f is transmitted. In the above error model, the subspace V_f is received as $U = W \oplus E$, where $\dim_{\mathbb{F}_h}(E) \leqslant \rho$ and W is a subspace of V_f with $\dim_{\mathbb{F}_h}(V_f) - \dim_{\mathbb{F}_h}(W) := \nu \leqslant \mu$. Assume that $\dim_{\mathbb{F}_q}(U) = d$.
Consider a nonzero polynomial in $\mathbb{F}_{q^m}[X, Y_1, Y_2, \ldots, Y_s]$ with $1 \leqslant s \leqslant m$

$$Q(X, Y_1, Y_2, \ldots, Y_s) = A_0(X) + A_1(Y_1) + A_2(Y_2) + \cdots + A_s(Y_s), \tag{25}$$

where every $A_i \in \mathbb{F}_{q^m}[X]$ is a h-linearized polynomial with $\deg_h(A_0) \leqslant D + k - 1$ and $\deg_h(A_i) \leqslant D$ for $i = 1, \ldots, s$; and D is chosen to be

$$D = \left\lfloor \frac{d-k-s+1}{s+1} \right\rfloor . \tag{26}$$

Choose an $\mathbb{F}_h$-basis $\{(a_i, b_i)\}_{i=1}^{d}$ of U (where $a_i \in \mathbb{F}_q$ and $b_i \in \mathbb{F}_{q^m}$) and we interpolate a polynomial Q of the above form satisfying

$$Q(a_i, b_i, b_i^q, \ldots, b_i^{q^{s-1}}) = 0 \quad \text{for } i = 1, 2, \ldots, d\,.$$

There are d equations, but $D + k + s(D+1) = (s+1)D + k + s$ freedoms in Q. Hence, such a nonzero polynomial Q exists since $d < (s+1)D + k + s$. It is clear that for all $(\alpha, f(\alpha)) \in W$, we have

$$0 = Q(\alpha, f(\alpha), f(\alpha)^q, \ldots, f(\alpha)^{q^{s-1}}) = Q(X, f, f^{\sigma}, \ldots, f^{\sigma^{s-1}})(\alpha),$$

where σ is the Frobenius automorphism of $\mathbb{F}_{q^m}$ over $\mathbb{F}_q$, i.e., σ sends every element α in $\mathbb{F}_{q^m}$ to α^q.

As the h-degree of $Q(X, f, f^{\sigma}, \ldots, f^{\sigma^{s-1}})$ is at most $D + k - 1$, under the condition

$$D + k - 1 < n - \nu, \tag{27}$$

we have

$$Q(X, f, f^{\sigma}, \ldots, f^{\sigma^{s-1}}) = A_0(X) + A_1(f(X)) + A_2(f^{\sigma}(X)) + \cdots + A_s(f^{\sigma^{s-1}}(X)) = 0. \tag{28}$$

[15] Here we follow the terminology of [14] which gave the first list decoding algorithm for subspace codes that could handle both insertions and deletions. Earlier results in [29] gave list decoding algorithms when only insertions were allowed.

Note that we have

$$\rho < s(n - \mu - k + 1) \Rightarrow \rho < s(n - \nu - k + 1) \Rightarrow d - n + \nu < s(n - \nu - k + 1) \Rightarrow D + k - 1 < n - \nu.$$

Thus, Condition (27) is met if

$$s\mu + \rho < s(n - k + 1). \tag{29}$$

The above analysis shows that we can list decode up to ρ insertions and μ deletions as long as ρ and ν satisfy (29).

The equation (28) satisfied by f is identical to (23), and therefore one can pin down f to an affine space of solutions exactly as in Lemma 8.4. After pre-coding by $\widetilde{\mathsf{HSE}}$, this space can then be efficiently pruned leading to a claim similar to Theorem 8.5. Since the details are identical to the Gabidulin case, we do not repeat them.

We conclude by commenting on the quality of our condition (29) for successful decoding, which is implied by the condition

$$\mu + \frac{\rho}{s} < n\left(1 - R - \frac{t}{n}\right) , \tag{30}$$

where R is the rate (24) of the code. For comparison, the condition for successful decoding for the folded KK codes in [14] is

$$\mu + \frac{\rho}{s} < n(1 - RN)$$

which necessitates a sub-constant rate for the code (a similar situation holds for [30]).

In [14], a random coding argument is used to show the existence of subspace codes that can be list decoded with list size L when

$$\mu + \frac{\rho}{L+1} < n\left(1 - R - \frac{1}{L+1}\right) , \tag{31}$$

where R is the rate of the code. To compare this with our result for KK codes, we note that after the combination with h.s.e sets, the analog of Theorem 8.5 will imply that we can take $s = \Theta(1/\varepsilon)$, $m = n/t = \Theta(1/\varepsilon^2)$, and have a list decodable subspace code that can correct μ deletions and ρ insertions provided

$$\mu + \varepsilon\rho < n(1 - R - \varepsilon) ,$$

with a worst-case output list size of $O(1/(R\varepsilon))$. This essentially matches the existential trade-off (31) up to constant factors in the list size.

B Omitted proofs concerning h.s.e sets

We now give the proofs that were omitted in Section 6.

B.1 Proof of Lemma 6.1

Proof. (Sketch) The parameters claimed, most notably the encoding time to compute HSE, are somewhat different than in [20], so let us briefly sketch the construction. The main difference is that, following [19, Sec. 4.1], we will define the h.s.e set based on the values of random polynomials rather than zero sets of random polynomials. The claim about the the structure of the map and its efficient computability will be clear from the construction. The proof of the h.s.e property is based on t-wise independence of values of random degree t polynomials and is more or less identical to the argument in [20, Sec. 4.3], and is therefore omitted.

Let $\kappa_1 = (1-\zeta)\kappa$, $\kappa_2 = (1-2\zeta)\kappa$, and for $j = 1, 2, \ldots, b$, let $n_j = (1-\zeta)j\Lambda$ and $m_j = (1-2\zeta)j\Lambda$, and let $\mathbb{K}_j$ be the extension field $\mathbb{F}_{q^{n_j}}$, and $\mathbb{L}_j$ the extension field $\mathbb{F}_{q^{m_j}}$. Fix some basis B_j (resp. B'_j) of $\mathbb{K}_j$ (resp. $\mathbb{L}_j$) over $\mathbb{F}_q$.

The construction of HSE will be based on polynomials $P_j \in \mathbb{K}_j[X]$ and $Q_j \in \mathbb{L}_j[X]$, $j = 1, 2, \ldots, b$, each chosen uniformly and independently among polynomials of degree at most $t := ck$. The polynomials P_j will be used to define a map $\varsigma_j : \mathbb{F}_q^{n_j} \to \mathbb{F}_q^{\zeta\Lambda}$ where to compute $\varsigma_j(\alpha)$, we treat α as an element of $\mathbb{K}_j$ using the basis B_j, compute $P_j(\alpha)$, and output the first $\zeta\Lambda$ coordinates of $P_j(\alpha)$ when viewed as a vector in $\mathbb{F}_q^{n_j}$ w.r.t basis B_j. We can similarly define maps $\vartheta_j : \mathbb{F}_q^{m_j} \to \mathbb{F}_q^{\zeta\Lambda}$ based on polynomials Q_j.

Using the maps ς_j, we will define $\mathsf{HSE}_1 : \mathbb{F}_q^{\kappa_1} \to \mathbb{F}_q^{\kappa}$ as follows: To compute $\mathsf{HSE}_1(\mathbf{x})$, we (i) break $\mathbf{x}$ into b blocks $x_1, x_2, \ldots, x_b$ with each $x_j \in \mathbb{F}_q^{(1-\zeta)\Lambda}$, (ii) compute $v_j = \varsigma_j(x_1 \circ x_2 \cdots \circ x_j) \in \mathbb{F}_q^{\zeta\Lambda}$ for $j = 1, 2, \ldots, b$ (here $\circ$ denotes concatenation of vectors), and (iii) output $x_1 \circ v_1 \circ x_2 \circ v_2 \circ \cdots \circ x_b \circ v_b$. We can similarly define $\mathsf{HSE}_2 : \mathbb{F}_q^{\kappa_2} \to \mathbb{F}_q^{\kappa_1}$ using the maps ϑ_j. The final map $\mathsf{HSE} : \mathbb{F}_q^{\kappa_2} \to \mathbb{F}_q^{\kappa}$ will be defined by the composition $\mathsf{HSE}(\mathbf{y}) = \mathsf{HSE}_1(\mathsf{HSE}_2(\mathbf{y}))$. By construction, if $\mathbf{y} = y_1 \circ \cdots \circ y_b$, then $\mathsf{HSE}(\mathbf{y})$ is of the form $y_1 \circ \xi_1 \circ y_2 \circ \xi_2 \circ \cdots \circ y_b \circ \xi_b$ where, for $1 \leqslant j \leqslant b$, each $\xi_j \in \mathbb{F}_q^{2\zeta\Lambda}$, depends only on the prefix $y_1 \circ y_2 \circ \cdots \circ y_j$, and can be computed in $\mathrm{poly}(\kappa, 1/\zeta, \log q)$ time.

Let $S = \mathsf{Im}(\mathsf{HSE})$. Given a subspace $H \in \mathcal{F}$, one can compute $H \cap S$ by computing $\mathrm{proj}_{j\Lambda}(H) \cap \mathrm{proj}_{j\Lambda}(S)$ for $j = 1, 2, \ldots, b$ iteratively, following the approach taken in [20]. At each step, to extend the intersection from $j-1$ blocks to j blocks, for each prefix $\mathbf{a} \in \mathrm{proj}_{(j-1)\Lambda}(H) \cap \mathrm{proj}_{(j-1)\Lambda}(S)$ (there will be at most $c\kappa$ such prefixes by the guarantee (17)), we can try each of the at most q^r possible extensions of $\mathbf{a}$ in $\mathrm{proj}_{j\Lambda}(H)$ and check which ones fall also belong to $\mathrm{proj}_{j\Lambda}(S)$ (note that this step will prune back the number of candidates to $c\kappa$ due to the h.s.e property of S). Since the latter task can be done efficiently for our construction, the overall runtime will be polynomial in $\kappa, q^r, 1/\zeta$. □

B.2 Proof of Lemma 6.2

Proof. (Sketch) The construction is due to Dvir and Lovett [6], with a minor change in some degree parameters to handle efficient encoding for any field $\mathbb{F}_q$. Choose $\gamma_1, \gamma_2, \ldots, \gamma_v$ to be arbitrary distinct nonzero elements in $\mathbb{F}_q$ (this is possible since $q > v$). Let $d_1, d_2, \ldots, d_v$ be distinct positive integers with $d_j = p^{j-1}$ for $j = 1, 2, \ldots, r$, and $d_j = p(j-r)+1$ for $r < j \leqslant v$. Note that $d_{\max} := \max_j d_j < \max\{p^r, pv\}$.

For $1 \leqslant i \leqslant r$ and $1 \leqslant j \leqslant v$, define $A_{i,j} = \gamma_j^i$ and $f_i \in \mathbb{F}_q[X_1, X_2, \ldots, X_v]$ as

$$f_i(X_1, X_2, \ldots, X_v) = \sum_{j=1}^{v} A_{i,j} X_j^{d_j} \ .$$

The subspace-evasive set will be the variety $V_{\mathbb{F}_q}(f_1, f_2, \ldots, f_r)$ of common zeroes of $f_1, f_2, \ldots, f_r$ in $\mathbb{F}_q^v$. The key property of the $r \times v$ matrix A is its Vandermonde nature: every r columns of A are linearly independent.

Since raising to a power of p is an automorphism of $\mathbb{F}_q$, the argument of [6] can be applied to show that for each setting of $\mathbb{F}_q$-values to $\{X_j \mid r < j \leqslant v\}$ (say $\mathbf{a} = (\alpha_{r+1}, \ldots, \alpha_v)$), there is a unique setting of $\mathbb{F}_q$-values to $X_1, X_2, \ldots, X_r$, say $\mathbf{b} = (\alpha_1, \ldots, \alpha_r)$ such $(\alpha_1, \alpha_2, \ldots, \alpha_v) \in V_{\mathbb{F}_q}(f_1, f_2, \ldots, f_r)$. Further, one can compute $\mathbf{b}$ given $\mathbf{a}$ efficiently (by solving a linear system and some exponentiation operations in $\mathbb{F}_q$) in $\mathrm{poly}(v, \log q)$ time. We use this to define the injective map $\psi : \mathbb{F}_q^{v-r} \to \mathbb{F}_q^v$ claimed in (i) in the natural way: $\psi(\mathbf{a}) = \mathbf{a} \circ \mathbf{b}$.

Let $\overline{\mathbb{F}}$ be the algebraic closure of $\mathbb{F}_q$. The main result in [6] shows that $V_{\overline{\mathbb{F}}}(f_1, f_2, \ldots, f_r)$ is an *everywhere-finite variety* that intersects every r-dimensional affine space over $\overline{\mathbb{F}}$ in at most $d_{\max}^r$ points in $\overline{\mathbb{F}}^v$. This implies that the image of ψ, which equals $V = V_{\mathbb{F}_q}(f_1, f_2, \ldots, f_r)$, is $(r, d_{\max}^r)$-subspace evasive. Known algorithms for solving systems of polynomial equations (see for instance [7, 25]) can be used to compute the intersection $H \cap V$ for an r-dimensional subspace $H \subseteq \mathbb{F}_q^v$ using a number of $\mathbb{F}_q$-operations that is at most polynomial in v and the intersection size.

Let us briefly sketch the high level approach behind this. Suppose H is given by affine constraints $l_1(\mathbf{x}) = l_2(\mathbf{x}) = \cdots = l_{v-r}(\mathbf{x}) = 0$. The approach is to compute the Gröbner basis of the zero-dimensional ideal $(f_1, f_2, \ldots, f_r, l_1, l_2, \ldots, l_{v-r})$ w.r.t. the lexicographic ordering $X_1 < X_2 < \cdots < X_n$ — this can be done using $v^{O(1)} d_{\max}^{O(r)}$ operations over $\mathbb{F}_q$ [7]. One can then use the elimination ideals to solve for all possible values to X_1, extend those to all possible values of (X_1, X_2), and so on, akin to solving a triangular linear system. The finiteness theorem for zero-dimensional ideals and the extension theorem of Gröbner basis theory imply that the number of operations needed by the iterative solution finding procedure is polynomial in the number of solutions. □

B.3 Proof of Theorem 6.3

Proof. The idea will be to use the construction of Lemma 6.1 at the "outer" level, and further encode blocks using Lemma 6.2 at the inner level. Let $\kappa = (1-\zeta)k$, $\Lambda = (1-\zeta)\Delta$. Let $\mathsf{HSE} : \mathbb{F}_q^{(1-2\zeta)\kappa} \to \mathbb{F}_q^{\kappa}$ be the map guaranteed by Lemma 6.1 for this choice of κ, Λ (and b, c, r, q as in the Lemma hypothesis); by padding the input with 0s, we can assume HSE maps $\mathbb{F}_q^{(1-3\zeta)k} \to \mathbb{F}_q^{(1-\zeta)\Delta b}$. Let ψ be the map guaranteed by Lemma 6.2 for the choice $v = \Delta$, $\eta = \zeta$, and r, q as in the hypothesis. Since $r < \zeta\Delta$, again we may pad the input with 0s as needed and assume that $\psi : \mathbb{F}_q^{(1-\zeta)\Delta} \to \mathbb{F}_q^{\Delta}$.

Our final map $\widetilde{\mathsf{HSE}} : \mathbb{F}_q^{(1-3\zeta)k} \to \mathbb{F}_q^k$ will be defined as follows. Given input $\mathbf{x} = x_1 \circ x_2 \circ \cdots \circ x_b \in \mathbb{F}_q^{(1-3\zeta)\Delta b}$, first compute $\mathsf{HSE}(\mathbf{x}) = y_1 \circ y_2 \circ \cdots \circ y_b$ with each $y_j \in \mathbb{F}_q^{(1-\zeta)\Delta}$. we define $\widetilde{\mathsf{HSE}}(\mathbf{x}) = \psi(y_1) \circ \psi(y_2) \circ \cdots \psi(y_b)$. The efficient computability of $\widetilde{\mathsf{HSE}}$ follows by construction and the properties of HSE and ψ. For $j = 1, 2, \ldots, b$, define $\widetilde{\mathsf{HSE}}_j(\mathbf{x})$ to be $\mathrm{proj}_{j\Delta}(\widetilde{\mathsf{HSE}}(\mathbf{x}))$; by construction $\widetilde{\mathsf{HSE}}_j$ is only a function of the the first j blocks $x_1 \circ x_2 \circ \cdots \circ x_j$ of $\mathbf{x}$.

To compute the intersection $H \cap S$, where $S = \mathsf{Im}(\widetilde{\mathsf{HSE}})$ and H is an (r, Δ, b)-periodic subspace, we use the same approach as in Lemma 6.1. The main difference is that, given some prefix $\mathbf{x_{j-1}} = x_1 \circ x_2 \circ \cdots \circ x_{j-1}$ where each $x_l \in \mathbb{F}_q^{(1-3\zeta)\Delta}$ that satisfies

$$\mathbf{a} := \widetilde{\mathsf{HSE}}_{j-1}(\mathbf{x_{j-1}}) \in \mathrm{proj}_{(j-1)\Delta}(H) \cap \mathrm{proj}_{(j-1)\Delta}(S) ,$$

we only need to try $\max\{p^{r^2}, (p\Delta)^r\}$ possible extensions x_j of $\mathbf{x_{j-1}}$ (instead of the earlier q^r bound). These extensions are the ones in the set $\Psi \subseteq \mathbb{F}_q^{(1-3\zeta)\Delta}$ defined as

$$\Psi := \{\mathrm{proj}_{(1-3\zeta)\Delta}(z) \mid z \in \mathbb{F}_q^{(1-\zeta)\Delta} \text{ and } \psi(z) \in H_j\}$$

where H_j is the affine subspace of extensions $\mathbf{h} \in \mathbb{F}_q^{\Delta}$ such that $\mathbf{a} \circ \mathbf{h} \in \mathrm{proj}_{j\Delta}(H)$. Note that the dimension of H_j is at most r as H is an (r, Δ, b)-periodic subspace. The size of Ψ is therefore at most $\max\{p^{r^2}, (p\Delta)^r\}$ by Lemma 6.2 and it can be enumerated in time polynomial in its size. □