Electronic Colloquium on Computational Complexity, Report No. 172 (2020)

OPTIMAL RATE LIST DECODING OVER BOUNDED ALPHABETS USING ALGEBRAIC-GEOMETRIC CODES

VENKATESAN GURUSWAMI AND CHAOPING XING

Abstract. We construct two classes of algebraic code families which are efficiently list decodable with small output list size from a fraction $1-R-\varepsilon$ of adversarial errors where R is the rate of the code, for any desired positive constant ε. The alphabet size depends only on ε and is nearly-optimal.

The first class of codes are obtained by *folding* algebraic-geometric codes using automorphisms of the underlying function field. The second class of codes are obtained by restricting evaluation points of an algebraic-geometric code to rational points from a *subfield*. In both cases, we develop a linear-algebraic approach to perform list decoding, which pins down the candidate messages to a subspace with a nice "periodic" structure.

To prune this subspace and obtain a good bound on the list-size, we pick subcodes of these codes by pre-coding into certain *subspace-evasive* sets which are guaranteed to have small intersection with the sort of periodic subspaces that arise in our list decoding. We develop two approaches for constructing such subspace-evasive sets. The first is a Monte Carlo construction of *hierearchical subspace-evasive* (h.s.e) sets which leads to excellent list-size but is not explicit. The second approach exploits a further *ultra-periodicity* of our subspaces and uses a novel construct called *subspace designs*, which were subsequently constructed explicitly and also found further applications in pseudorandomness.

To get a family of codes over a fixed alphabet size, we instantiate our approach with algebraic-geometric codes based on the Garcia-Stichtenoth tower of function fields. Combining this with pruning via h.s.e sets yields codes list-decodable up to a $1-R-\varepsilon$ error fraction with list size bounded by $O(1/\varepsilon)$, matching the existential bound for random codes up to constant factors. Further, the alphabet size can be made $\exp(\tilde{O}(1/\varepsilon^2))$ which is not much worse than the lower bound of $\exp(\Omega(1/\varepsilon))$. The parameters we achieve are thus quite close to the existential bounds in all three aspects—error-correction radius, alphabet size, and list-size— simultaneously. This construction is, however, Monte Carlo and the claimed list decoding property only holds with high probability. Once the code is (efficiently) sampled, the encoding/decoding algorithms are deterministic with a running time $O_\varepsilon(N^c)$ for an absolute constant c, where N is the code's block length.

Using subspace designs instead for the pruning, our approach yields a 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 $\exp(\tilde{O}(1/\varepsilon^2))$. The list size bound is upper bounded by a very slowly growing function of the block length N; in particular, it is at most $O(\log^{(r)} N)$ (the r'th iterated logarithm) for any fixed integer r. The explicit construction avoids the shortcoming of the Monte Carlo sampling at the expense of a worse list size.

Extended abstracts announcing these results were presented at the 2012 and 2013 ACM Symposia on Theory of Computing (STOC) [18, 19]. This is a merged and significantly revised version of these conference papers, that accounts for the explicit subspace designs that were constructed in [11] subsequent to [19], makes some simplifications and improvements in the construction of h.s.e sets in Section 9 compared to [18], and reorganizes the material and flow substantially.

The research of V. Guruswami was supported in part by a Packard Fellowship and NSF grants CCF-0963975, CCF-1422045, and CCF-1814603. Some of this work was done during visits by the author to Nanyang Technological University.

The research of C. Xing was supported by the National Research Foundation, Prime Minister's Office, Singapore under its Strategic Capability Research Centres Funding Initiative; and the Singapore MoE Tier 1 grants RG25/16 and RG21/18.

ISSN 1433-8092

Contents

1. Introduction

An error-correcting code C of block length N over a finite alphabet Σ maps a set $\mathcal{M}$ of messages into codewords in Σ^N. The rate of the code C, denoted R, equals $\frac{1}{N}\log_{|\Sigma|}|\mathcal{M}|$. In this work, we will be interested in codes for adversarial noise, where the channel can arbitrarily corrupt any subset of up to τN symbols of the codeword. The goal will be to correct such errors and recover the original message/codeword efficiently. It is easy to see that information-theoretically, we need to receive at least RN symbols correctly in order to recover the message (since $|\mathcal{M}| = |\Sigma|^{RN}$), so we must have $\tau \leqslant 1 - R$.

Perhaps surprisingly, in a model called list decoding, recovery up to this information-theoretic limit becomes possible. Let us say that a code $C \subseteq \Sigma^N$ is (τ, ℓ)-list decodable if for every received word $\mathbf{y} \in \Sigma^N$, there are at most ℓ codewords $\mathbf{c} \in C$ such that $\mathbf{y}$ and $\mathbf{c}$ differ in at most τN positions. Such a code allows, in principle, the correction of a fraction τ of errors, outputting at most ℓ candidate codewords one of which is the originally transmitted codeword.

The probabilistic method shows that a random code of rate R over an alphabet of size $\exp(O(1/\varepsilon))$ is with high probability $(1 - R - \varepsilon, O(1/\varepsilon))$-list decodable [4]. However, it is not known how to construct or even randomly sample such a code for which the associated algorithmic task of list decoding (i.e., given $\mathbf{y} \in \Sigma^N$, find the list of codewords within fractional radius $1 - R - \varepsilon$) can be performed efficiently. This work takes a big step in that direction, giving a randomized construction of such efficiently list-decodable codes over a slightly worse alphabet size of $\exp(\tilde{O}(1/\varepsilon^2))$. We note that the alphabet size needs to be at least $\exp(\Omega(1/\varepsilon))$ in order to list decode from a fraction $1 - R - \varepsilon$ of errors, [1] so this is close to optimal. For the list-size needed as a function of ε for decoding a $1 - R - \varepsilon$ fraction of errors, the best lower bound is only $\Omega(\log(1/\varepsilon))$ [12], but as mentioned above, even random coding arguments only achieve a list-size of $O(1/\varepsilon)$, which our construction matches up to constant factors. We also give a fully *deterministic* construction with a list-size that is very slowly growing as a function of the block length.

We now review some of the key results on algebraic list decoding leading up to this work. A more technical comparison with related work appears in Section 1.1. The work of Sudan [31] used bivariate polynomial interpolation to give the first efficient list decoding algorithm for Reed-Solomon codes, which for rates R below $1/3$ corrected a fraction of errors exceeding the $(1 - R)/2$ bound achievable by unique decoding. Guruswami and Sudan [15] introduced multiplicities in the interpolation step and gave an efficient list decoding algorithm that could correct an error-fraction $1 - \sqrt{R}$. The multiplicities also offered an avenue to incorporate "soft" information about varying reliability of different symbols, which was developed by Koetter and Vardy [21] to give an influential algebraic soft-decision decoder for Reed-Solomon codes. The $1 - \sqrt{R}$ bound remained the largest known efficiently list-decodable error-fraction for any value of rate R till Parvaresh and Vardy [26] gave a variant of Reed-Solomon codes list-decodable up to error fraction $1 - O(R\log(1/R))$ which beats the $1 - \sqrt{R}$ bound for low-rates.

Building on the Parvaresh-Vardy work together with further new algebraic ideas, Guruswami and Rudra [13] gave the first construction of codes that achieved the optimal trade-off between rate and list-decoding radius, i.e., enabled list decoding up to a fraction $1 - R - \varepsilon$ of worst-case errors with rate R. They showed that a variant of Reed-Solomon

[1] The best trade-off between rate R and list decoding radius τ is the Gilbert-Varshamov bound, i.e., $R \leqslant 1 - H_q(\tau)$, where $H_q(\tau)$ is the q-ary entropy function $x\log_q(q-1) - x\log_q x - (1-x)\log_q(1-x)$. The function $1 - H_q(\tau)$ is equal to $1 - \tau - \varepsilon$ if the alphabet size is at least $\exp(\Omega(1/\varepsilon))$.

(RS) codes called *folded* RS codes admit such a list decoder. For a decoding radius of $1 - R - \varepsilon$, the code was based on bundling together disjoint windows of $m = \Theta(1/\varepsilon^2)$ consecutive symbols of the RS codeword into a single symbol over a larger alphabet. As a result, the alphabet size of the construction was $N^{\Omega(1/\varepsilon^2)}$. It was also shown in [13] that ideas based on code concatenation and expander codes can be used to bring down the alphabet size to $\exp(\tilde{O}(1/\varepsilon^4))$ which is independent of the block length. However, the resulting codes lose some important and powerful features such as list recovery and soft decoding of the folded RS codes. Also, the decoding time complexity as well as proven bound on worst-case output list size for folded RS codes were $N^{\Omega(1/\varepsilon)}$.[2]

Our main final result statement is the following, offering two constructions, one randomized and one deterministic, of variants of algebraic-geometric (AG) codes that are list-decodable with optimal rate. These appear as Theorems 11.4 and 11.8 in the final section of the paper.

Theorem 1.1 (Main). *For any $R \in (0,1)$ and positive constant $\varepsilon \in (0,1)$, there is*

(i) *a Monte Carlo construction of a family of codes of rate at least R over an alphabet size $\exp(O(\log(1/\varepsilon)/\varepsilon^2))$ that are encodable and $(1-R-\varepsilon, O(1/(R\varepsilon))$-list decodable in $O_\varepsilon(N^c)$ time*[3]*, where N is the block length of the code and c is an absolute positive constant.*

(ii) *a deterministic construction of a family of codes of rate at least R over an alphabet size $\exp(O(\log^2(1/\varepsilon)/\varepsilon^2))$ that are encodable and $(1 - R - \varepsilon, L(N))$-list decodable in $O_\varepsilon(N^c)$ time, for a list size that satisfies $L(N) = o(\log^{(r)} N)$ (the r'th iterated logarithm) for any fixed integer r.*

The first part of Theorem 1.1 is achieved through folded algebraic-geometric codes. To fold algebraic-geometric codes, we first find suitable automorphisms of the ground function field. The list of possible candidate messages output by the list decoder has exponential size, but is contained in a well structured subspace. To prune down the list size, we only encode messages that belong to so-called hierarchical subspace-evasive sets, which are chosen to have small intersection with the structured subspaces arising in the decoding. To make use of subspace-evasive sets efficiently, we have to: (i) give an efficient pseudorandom construction of these sets; and (ii) encode the messages to subspace-evasive sets efficiently. We refer to Section 2 for details.

The second part of Theorem 1.1 is obtained through usual algebraic-geometric codes with evaluation points over subfields. As in the first part, the list of possible candidate messages belongs to a subspace that is well structured, specifically with a property we called ultra-periodicity (Definition 3). The approach based on hierarchical subspace-evasive sets in the first part leads to excellent list size; however, we only know randomized constructions of hierarchical subspace-evasive sets. To obtain a deterministic list decoding, we prune down the list of possible solutions through subspace designs (see Section 2 for details).

We note that our Monte Carlo construction gives codes that are quite close to the existential bounds in three aspects simultaneously — the trade-off between error fraction $1 - R - \varepsilon$ and rate R, the list-size as a function of ε, and the alphabet size of the code family (again as a function of ε). Even though these codes are not fully explicit, they are "functionally explicit" in the sense that once the code is (efficiently) sampled, with high

[2] The list size for decoding folded RS codes was shown to be bounded by a constant depending only on ε in subsequent work [23].

[3] We use the $O_\varepsilon(\cdot)$ notation to hide constant factors that depend on ε.

probability the polynomial time encoding and decoding algorithms deliver the claimed error-correction guarantees for *all* allowed error pattern. The explicit construction avoids this shortcoming at the expense of a slightly worse list size.

1.1. **Prior and related work.** Let us recap a bit more formally the construction of folded RS codes from [13]. One begins with the Reed-Solomon encoding of a polynomial $f \in \mathbb{F}_q[X]$ of degree $< k$ consisting of the evaluation of f on a subset of field elements ordered as $1, \gamma, \dots, \gamma^{N-1}$ for some primitive element $\gamma \in \mathbb{F}_q$ and $N < q$. For an integer "folding" parameter $m \geqslant 1$ that divides N, the folded RS codeword is defined over alphabet $\mathbb{F}_q^m$ and consists of n/m blocks, with the j'th block consisting of the m-tuple $(f(\gamma^{(j-1)m}), f(\gamma^{(j-1)m+1}), \dots, f(\gamma^{jm-1}))$. The algorithm in [13] for list decoding these codes was based on the algebraic identity $\overline{f(\gamma X)} = \overline{f(X)}^q$ in the residue field $\mathbb{F}_q[X]/(X^{q-1} - \gamma)$ where $\overline{f}$ denotes the residue f mod $(X^{q-1} - \gamma)$. This identity is used to solve for f from an equation of the form $Q(X, f(X), f(\gamma X), \dots, f(\gamma^{s-1}X)) = 0$ for some low-degree nonzero multivariate polynomial Q. The high degree $q > n$ of this identity, coupled with $s \approx 1/\varepsilon$, led to the large bounds on list-size and decoding complexity in [13].

One possible approach to reduce q (as a function of the code length) in this construction would be to work with algebraic-geometric codes based on function fields K over $\mathbb{F}_q$ with more rational points. However, an automorphism σ of K that can play the role of the automorphism $f(X) \mapsto f(\gamma X)$ of $\mathbb{F}_q(X)$ is only known (or even possible) for very special function fields. This approach was used in [9] to construct list-decodable codes based on cyclotomic function fields using as σ certain Frobenius automorphisms. These codes improved the alphabet size to polylogarithmic in N, but the bound on list-size and decoding complexity remained $N^{\Omega(1/\varepsilon)}$.

Subsequently, a linear-algebraic approach to list decoding folded RS codes was discovered in [32, 10]. Here, in the interpolation stage, which is common to all list decoding algorithms for algebraic codes [31, 15, 26, 13], one finds a *linear* multivariate polynomial $Q(X, Y_1, \dots, Y_s)$ whose total degree in the Y_i's is 1. The simple but key observation driving the linear-algebraic approach is that the equation $Q(X, f(X), \dots, f(\gamma^{s-1}X)) = 0$ now becomes a linear system in the coefficients of f. Further, it is shown that the solution space has dimension less than s, which again gives a list-size upper bound of q^{s-1}. Finally, since the list of candidate messages fall in an affine space, it was noted in [10] that one can bring down the list size by carefully "pre-coding" the message polynomials so that their k coefficients belong to a "subspace-evasive set" (which has small intersection with every s-dimensional subspace of $\mathbb{F}_q^k$). This idea was used in [16] to give a *randomized* construction of $(1 - R - \varepsilon, O(1/\varepsilon))$-list decodable codes of rate R. However, the alphabet size and runtime of the decoding algorithm both remained $N^{\Omega(1/\varepsilon)}$. Similar results were also shown in [16, 22] for univariate multiplicity codes, where the encoding of a polynomial f consists of the evaluations of f and its first $m - 1$ derivatives at distinct field elements.

Concurrently with the conference version of part of this work reported in [18], Dvir and Lovett [3] gave an elegant construction of explicit subspace evasive sets based on certain algebraic varieties. Furthermore, Ben-Aroya and Shinkar [1] improved the result of [3] slightly by using an elementary construction. Their results yield an explicit version of the codes from [10], albeit with a worse list size bound of $(1/\varepsilon)^{O(1/\varepsilon)}$. This work and [3, 1] are incomparable in terms of results. The big advantage of [3, 1] is the deterministic construction of the code. The benefits in our work are: (i) both constructions in the present paper give codes over an alphabet size that is a constant independent of N, whereas in [3] the $N^{\Omega(1/\varepsilon^2)}$ alphabet size of folded RS codes is inherited; (ii) our first

Monte Carlo construction ensures list-decodability with a list-size of $O(1/\varepsilon)$ that is much better and in fact matches the full random construction up to constant factors,[4] and (iii) our second construction gives a deterministic algorithm as well with almost constant list size (and constant alphabet size). Another important feature is that both our work and [3, 1] achieve a decoding complexity of $O_\varepsilon(N^c)$ with exponent independent of ε.

Our paper presents two class of codes: folded algebraic-geometric codes and usual algebraic-geometric codes with evaluation points over subfields. For both the classes of codes, we can apply hierarchical subspace-evasive sets as well as subspace design to prune down the list size by taking certain subcodes. This is because of the "periodic" structure of the subspace in which the candidate messages are pinned down by the linear-algebraic list decoder is similar in both cases. Thus, we can obtain both randomized and deterministic algorithms from each of the two classes of codes. In total, we have four combinations of constructions. To illustrate both algebraic approaches, we decide to focus on two combinations, i.e., (i) folded algebraic-geometric codes with hierarchical subspace-evasive sets; and (ii) usual algebraic-geometric codes with evaluation points over subfields with subspace designs. These are listed in Figure 1. We note that the other two combinations are also possible, as the pruning of the subspace of solutions is "black-box" with respect to its periodic structure.

In the table presented in Figure 1, we list previous results and those in this paper. The major improvement of this work is to bring down the alphabet size to constant, while at the same time ensuring small list size and low decoding complexity where the exponent of the polynomial run time does not depend on ε. Our folded algebraic-geometric subcodes achieve a list size matching the fully random constructions up to constant factors, together with alphabet size not much worse than the lower bound $\exp(\Omega(1/\varepsilon))$. On the last line, our algebraic-geometric subcodes give a deterministic list decoding with almost constant list size and optimal decoding radius.

1.2. **Subsequent works and open questions.**

The challenge of decoding up to radius approaching the optimal bound $(1-R)$ with rate R along with good list and alphabet size is, for the most part, solved by our work. There are still some goals that have not been met. One is to get a fully deterministic construction with constant list-size and alphabet size (as a function of ε), and construction/decoding complexity $O_\varepsilon(N^c)$. This has been almost achieved by Kopparty, Ron-Zewi, Saraf, and Wootters [23]. They prove that the list-size for list-decoding folded Reed-Solomon codes is itself, without any pruning by subspace evasive sets, bounded by a constant. They then combine it with several other tools from algebraic coding theory and pseudorandomness to construct codes of rate R list-decodable up to a $(1-R-\varepsilon)$ error fraction with constant list and alphabet size (depending only on ε) and decoding complexity $O_\varepsilon(N^c)$ (in fact the exponent c can be made arbitrarily close to 1). An exciting recent result by Guo and Ron-Zewi [8] achieves both constant list-size and alphabet within our framework, via improved subspace evasive sets for the ultra-periodic subspaces output by the list decoder.

Another challenge is to construct a $(1-R-\varepsilon, L)$-list decodable code of rate R (for list size L bounded by a polynomial in the block length), over an alphabet of size $\exp(O(1/\varepsilon))$, which is the asymptotically optimal size. All known constructions over a constant-sized alphabet known so far have alphabet size at least $\exp(\Omega(1/\varepsilon^2))$. Finally, the various algebraic and expander-based techiques that have led to progress on list decoding only

[4] As mentioned above, the bound in [3] is $(1/\varepsilon)^{O(1/\varepsilon)}$ and it seems very difficult to get a sub-exponential dependence on $1/\varepsilon$ with the algebraic approach relying on Bezout's theorem to construct subspace-evasive sets.

Code	Construction	Alphabet size	List size	Decoding time	Reference
Folded RS/derivative	Explicit	$N^{O(1/\varepsilon^2)}$	$N^{O(1/\varepsilon)}$	$N^{O(1/\varepsilon)}$	[13, 16]
Folded RS subcode	Randomized	$N^{O(1/\varepsilon^2)}$	$O(1/\varepsilon)$	$N^{O(1/\varepsilon)}$	[16]
Folded RS subcode	Explicit	$N^{O(1/\varepsilon^2)}$	$(1/\varepsilon)^{O(1/\varepsilon)}$	$N^{O(1)}2^{1/\varepsilon^{O(1)}}$	[3]
Folded cyclotomic	Explicit*	$(\log N)^{O(1/\varepsilon^2)}$	$N^{O(1/\varepsilon^2)}$	$N^{O(1/\varepsilon^2)}$	[9]
Folded AG subcode	Randomized	$\exp(\tilde{O}(1/\varepsilon^2))$	$O(1/\varepsilon)$	$N^{O(1)}2^{1/\varepsilon^{O(1)}}$	Thm. 1.1(i)
AG subcode	Explicit	$\exp(\tilde{O}(1/\varepsilon^2))$	$2^{2^{2^{(\log^* N)^2}}}$	$N^{O(1)}(1/\varepsilon)^{O(1)}$	Thm. 1.1(ii)

FIGURE 1. N in the above table stands for the length of codes. Parameters of various constructions of codes that enable list decoding $(1 - R - \varepsilon)$ fraction of errors, with rate R. The last two lines are from this work. “Explicit” means the code can be constructed in deterministic polynomial time (the * for folded cyclotomic is because of requirement of an irreducible polynomial of high degree, which can be sampled and then checked (for a ”Las Vegas” construction)). The rows with first column in boldface are not dominated by other constructions. The last line gives the first deterministic construction of algebraic codes for efficient optimal rate list decoding over constant-sized alphabets.

work over large alphabets. The challenge of efficient optimal rate list decoding over say the binary alphabet, even for the simpler model of erasures, remains wide open. The best known constructions are obtained via concatenation, and are list-decodable up to the so-called Blokh-Zyablov bound [14].

1.3. **Organization.** The paper is organized as follows. In Section 2, we describe the detailed techniques of our paper including algebraic approaches and pseudorandomness. Following the section on techniques, in Section 3 we introduce periodic and ultra-periodic subspaces, give definitions and basic properties. In Section 4, we recall some basic results on function fields and algebraic-geometric codes. To illustrate our ideas in an algebraically simpler (and perhaps more practical) setting, in Section 6 we give a construction based on a tower of Hermitian field extensions [27]. This is capable of giving a similar result to our best ones based on the Garcia-Stichtenoth tower, albeit with alphabet size and list-size upper bound polylogarithmic in the code length. In Section 9 we first introduce hierarchical subspace-evasive sets, then show that random sets are hierarchical subspace-evasive with high probability. We also present a pseudorandom construction of hierarchical subspace-evasive sets, which also allow for efficient encoding and efficient computation of intersection with periodic subspaces.

Folded algebraic-geometric codes from the Garcia-Stichtenoth tower are studied in Section 8. The list size, decoding radius and decoding algorithm via local expansion are also discussed in this section. Section 11 is devoted to the discussion of pruning down the list size for folded codes from both the Hermitian and the Garcia-Stichtenoth towers using hierarchical subspace-evasive sets. The second class of our codes, namely usual algebraic-geometric codes with evaluation points over subfields is presented in Section 7. In this section, we first discuss list decoding for the simpler Reed-Solomon case, and then generalize it to list decoding of arbitrary algebraic geometric codes and finally instantiate the approach with the codes from the Garcia-Stichtenoth tower. In Section 10, we introduce subspace designs and cascaded subspace designs, and discuss parameters of random and explicit constructions of those. In the last section, the explicit construction of subcodes of RS and AG subcodes based on subspace designs is presented.

2. Our techniques

We describe some of the main new ingredients that go into our work. We need both new algebraic insights and constructions, as well as ideas in pseudorandomness relating to (variants of) subspace-evasive sets. We describe these in turn below.

2.1. **Algebraic ideas.**

It is shown in [15] that one can list decode the usual algebraic-geometric codes up to the Johnson bound. On the other hand, one has not found list decoding algorithms of the usual algebraic-geometric codes beyond the Johnson bound. Thus, to list decode the algebraic-geometric codes beyond the Johnson bound, it is natural to consider some variants of usual algebraic-geometric codes as one does for Reed-Solomon codes [13]. In this work, we present two new variants of algebraic-geometric codes–folded algebraic geometric codes and usual algebraic geometric codes with evaluation points over subfields. We describe these in turn.

2.1.1. *Folding AG codes.* The first approach is to use suitable automorphisms of function fields to fold the code. This approach was used for Reed-Solomon codes in [13] and for cyclotomic function field in [9], though this was done using the original approach in [13] where the messages to be list decoded were pinned down to the roots of a higher degree polynomial over a large residue field. As mentioned earlier, effecting this "non-linear" approach in [13, 9] with automorphisms of more general function fields seems intricate at best. In this work we employ the linear-algebraic list decoding method of [16]. However, the correct generalization of the linear-algebraic list decoding approach to the function field case is also not obvious. One of the main algebraic insights in this work is noting that a possible way to generalize the linear-algebraic approach to codes based on algebraic function fields is to rely on the *local power series expansion* of functions from the message space at a suitable rational point. (The case for Reed-Solomon codes being the expansion around 0, which is a finite polynomial form.)

Working with a suitable automorphism which has a "diagonal" action on the local expansion lets us extend the linear-algebraic decoding method to AG codes (here by a "diagonal" action, we mean that this action gives rise to equations on coefficients of a polynomial that are diagonal). Implementing this for specific AG codes requires an explicit specification of a basis for an associated message (Riemann-Roch) space, and the efficient computation of the local expansion of the basis elements at a special rational point on the curve. We show how to do this for two towers of function fields: the Hermitian tower [27] and the asymptotically optimal Garcia-Stichtenoth tower [6, 7]. The former tower is quite simple to handle — it has an easily written down explicit basis, and we show how to compute the local expansion of functions around the point with all zero coordinates. However, the Hermitian tower does not have bounded ratio of the genus to number of rational points, and so does not give constant alphabet codes (we can get codes over an alphabet size that is polylogarithmic in the block length though). Explicit basis for Riemann-Roch spaces of the Garcia-Stichtenoth tower were constructed in [28]. Regarding local expansions, one major difference is that we work with local expansion of functions at the point at infinity, which is fully "ramified" in the tower. For both these towers, we find and work with a nice automorphism that acts diagonally on the local expansion, and use it for folding the codes and decoding them by solving a linear system.

2.1.2. *Restricting evaluation points to a subfield.* The second approach is to work with "normal" algebraic-geometric codes, based on evaluating functions from a Riemann-Roch space at some rational places, except we use a constant field extension of the function field for the function space, but restrict to evaluating at rational places over the original base field. Let us give a brief idea why restricting evaluation points to a subfield enables correcting more errors. The idea behind list decoding results for folded RS (or derivative) codes in [13, 16] 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. $A_0(X)+A_1(X)f(X)+A_2(X)f^q(X)+\cdots+A_s(X)f^{q^{s-1}}(X) \equiv 0 \pmod{x^{q-1}-\gamma}$ in the case of folded Reed-Solomon codes [13] (here γ is a primitive element of $\mathbb{F}_q$, 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 Definition 6, 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 $f(X) = \sum_{j=0}^{k-1} f_j X^j$). 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. We remark here that this approach can also be applied effectively to linearized polynomials, and can be used to construct variants of Gabidulin codes that are list-decodable up to the optimal $1-R$ fraction of errors (where R is the rate) in the rank metric [17].

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. 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 [6, 7] 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 8.3), after describing the framework for general AG codes first.

2.2. **Pseudorandomness.** The above algebraic ideas enable us to pin down the messages into a structured subspace of dimension linear in the message length. The specific structure of the subspace is a certain "periodicity" — there is a subspace $W \subset \mathbb{F}_q^m$ such that once $f_0, f_1, \ldots, f_{i-1}$ (the first i coefficients of the message polynomial) are fixed, f_i belongs to a coset of W. We now describe our ideas to prune this list, by restricting (or "pre-coding") the message polynomials to belong to carefully constructed pseudorandom subsets that have small intersection with any periodic subspace.

2.2.1. *Hierarchical subspace-evasive sets.* The first approach follows along the lines of [16] and we only encode messages in a *subspace-evasive set* which has small intersection with low-dimensional subspaces. Implementing this in our case, however, leads to several problems. First, since the subspace we like to avoid intersecting much has large dimension, the list size bound will be linear in the code length and not a constant like in our final result (by "a constant", we mean that the list size is independent of the code length and dependent on ε). More severely, we cannot go over the elements of this subspace to prune the list as that would take exponential time. To solve the latter problem, we observe that the subspace has a special "periodic" structure, and exploit this to show the existence of large "hierarchically subspace evasive" (h.s.e) subsets which have small intersection with the projection of the subspace on certain prefixes. Isolating the periodic property of the subspaces, and formulating the right notion of evasiveness w.r.t to such subspaces, is an important aspect of this work.

We also give a construction of good h.s.e sets using limited wise independent sample spaces, in a manner enabling the efficient iterative computation of the final list of intersecting elements. Further our construction allows for efficient indexing into the h.s.e set which leads to an efficient encoding algorithm for our code). As a further ingredient, we note that the number of possible subspaces that arise in the decoding is much smaller than the total number of possibilities. Using this together with an added trick in the h.s.e set construction, we are able to reduce the list size to a constant.

2.2.2. *Subspace designs.* The approach based on h.s.e sets leads to excellent list size; however, we only know randomized constructions of h.s.e sets with the required properties. Our second approach to prune the subspace of possible solutions is based on *subspace designs* and leads to deterministic subcode constructions. More precisely speaking, the coefficients $f_0, f_1, \ldots, f_{k-1}$ of the message polynomial (which belong to the extension field $\mathbb{F}_{q^m}$) are pinned down by the linear-algebraic list decoder to a periodic subspace with the property that there is an $\mathbb{F}_q$-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. 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 $\{H_i\}_{i=1}^k$ 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 10 as a key contribution in this work. Indeed, subsequent work by Forbes and Guruswami [5] highlighted the central role played by subspace designs in "linear-algebraic pseudorandomness" and in particular how they lead to rank condensers and dimension expanders.

A simple probabilistic argument shows that, with high probability, any $q^{\Omega(\varepsilon m)}$ subspaces of dimension $(1-\varepsilon)m$ that are randomly chosen have small total intersection with every s-dimensional W. This construction can also be derandomized, though the construction complexity of the resulting codes becomes quasi-polynomial with this approach for the parameter choices needed in the construction.

Fortunately, in a follow-on to [19], Guruswami and Kopparty gave explicit constructions of subspace designs with parameters nearly matching the random constructions [11]. One can pre-code with this subspace design to get explicit list-decodable sub-codes of *Reed-Solomon codes* whose evaluation points are in a subfield (Section 11.2.1). However, this construction inherits the large field size of Reed-Solomon codes.

For explicit subcodes of algebraic-geometric codes using subspace designs we need additional ideas. The dimension k in the case of AG codes is much larger than the alphabet size q^m (in fact that is 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 (Section 10.4) — 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 "periodic subspace" (W when we begin) at the corresponding scale. With a careful technical argument and choice of parameters, we are able to obtain the bounds of Theorem 1.1(ii).

3. Periodic subspaces

In this section we formalize a certain "periodic" property of affine subspaces that will arise in our list decoding application.

We begin with some notation. For a vector $\mathbf{y} = (y_1, y_2, \ldots, y_m)^T \in \mathbb{F}_q^m$ 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., $\mathrm{proj}_{[t_1,t_2]}(\mathbf{y}) = (y_{t_1}, y_{t_1+1}, \ldots, y_{t_2})^T$. When $t_1 = 1$, we use $\mathrm{proj}_t(\mathbf{y})$ to denote $\mathrm{proj}_{[1,t]}(\mathbf{y})$. By default, we treat vectors as column vectors. These notions are extended to subsets of strings in the obvious way: $\mathrm{proj}_{[t_1,t_2]}(S) = \{\mathrm{proj}_{[t_1,t_2]}(\mathbf{x}) \mid \mathbf{x} \in S\}$.

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

Definition 1 (Periodic (affine) subspaces). *For positive integers r, b, Δ with $r < \Delta$ and $\kappa := b\Delta$, an affine subspace $H \subset \mathbb{F}_q^\kappa$ is said to be (r, Δ, b)-periodic if there exists a matrix $B \in \mathbb{F}_q^{\Delta\times\Delta}$ whose kernel $\ker(B)$ has dimension at most r, and vectors $\mathbf{a}_\ell \in \mathbb{F}_q^\Delta$ and matrices $A_\ell \in \mathbb{F}_q^{\Delta\times(\ell-1)\Delta}$ for $1 \leqslant \ell \leqslant b$, such that every $\mathbf{x} \in H$ satisfies the following equations for $\ell = 1, 2, \ldots, b$:*

$$\mathbf{a}_\ell + A_\ell \cdot \mathrm{proj}_{(\ell-1)\Delta}(\mathbf{x}) + B \cdot \mathrm{proj}_{[(\ell-1)\Delta+1,\ell\Delta]}(\mathbf{x}) = 0 \ . \tag{1}$$

In other words, the projections of the subspace onto blocks of contiguous Δ symbols, conditioned on any prefix, always belong to an affine shift of the subspace $W := \ker(B)$ of dimension at most r.

For dimensions κ not necessarily divisible by Δ, we say that an affine subspace $H \subseteq \mathbb{F}_q^\kappa$ is (r, Δ)-periodic if there is exists a (r, Δ, b)-periodic subspace $H' \subseteq \mathbb{F}_q^{b\Delta}$ for $b = \lceil \frac{\kappa}{\Delta} \rceil$ such that $H = \mathrm{proj}_{[1,\kappa]}(H')$.

We will call W the recurring subspace *of the periodic subspace H.*

Definition 2 (Representing periodic affine subspaces). *The matrices A_i and vectors $\mathbf{a_i}$, $i = 1, 2, \ldots, b$, and the matrix B, or equivalently the system of equations* (1)*, can be used to specify the (r, Δ, b)-periodic subspace H, and this is the representation of periodic subspaces that will naturally arise in our list decoders.*

The motivation for 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. (Here q^m will be the alphabet size of the code, k its dimension, and s will be a parameter of the algorithm that governs how close the decoding performance approaches the Singleton bound.)

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

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

(1) *the projection of H to the first j blocks of Δ coordinates, $\mathrm{proj}_{j\Delta}(H) = \{\mathrm{proj}_{j\Delta}(\mathbf{x}) \mid \mathbf{x} \in H\}$, has dimension at most jr. (In particular H has dimension at most br.)*

(2) *for each* $\mathbf{a} \in \mathbb{F}_q^{(j-1)\Delta}$*, there are at most* q^r *extensions* $\mathbf{y} \in \mathrm{proj}_{j\Delta}(H)$ *such that* $\mathrm{proj}_{(j-1)\Delta}(\mathbf{y}) = \mathbf{a}$*.*

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 (Ultra-periodic subspace). *For positive integers* r, b, Δ *with* $r < \Delta$*, an affine subspace* H *of* $\mathbb{F}_q^\kappa$ *for* $\kappa = b\Delta$ *is said to be* (r, Δ, b)-ultra periodic *if there exist vectors* $a_\ell \in \mathbb{F}_q^\Delta$ *and matrices* $B_\ell \in \mathbb{F}_q^{\Delta\times\Delta}$ *with* $\dim(\ker(B_\ell)) \leqslant r$ *for* $\ell = 1, 2, \ldots, b$*, such that every* $\mathbf{x} \in H$ *satisfies the following equations for* $\ell = 1, 2, \ldots, b$*:*

$$\mathbf{a}_\ell + \sum_{i=1}^{\ell} B_{\ell-i+1} \cdot \mathrm{proj}_{(i-1)\Delta+1, i\Delta}(\mathbf{x}) = 0 \ . \tag{2}$$

In other words, the space H *is defined by equations that have a lower-triangular "Toeplitz" block-diagonal structure, with the blocks on the diagonal being* B_1*, the blocks on the next lower diagonal being* B_2*, the next diagonal having* B_3*, and so on.*

For ambient dimensions κ *not necessarily divisible by* Δ*, we say that an affine subspace* $H \subseteq \mathbb{F}_q^\kappa$ *is* (r, Δ)*-ultra periodic if there is exists a* (r, Δ, b)*-ultra periodic subspace* $H' \subseteq \mathbb{F}_q^{b\Delta}$ *for* $b = \lceil \frac{\kappa}{\Delta} \rceil$ *such that* $H = \mathrm{proj}_{[1,\kappa]}(H')$*.*

We have the below observation that follows from the definition of ultra-periodicity.

Observation 3.2. *If a subspace* H *of* $\mathbb{F}_q^\kappa$ *is* (r, Δ)-ultra periodic*, then for every integer* ℓ*,* $1 \leqslant \ell \leqslant \frac{\kappa}{\Delta}$*,* H *is* $(\ell r, \ell\Delta)$*-periodic.*

Thus ultra-periodicity captures the fact that the subspace is periodic not only for blocks of size Δ, but also for block sizes that are multiples of Δ. Thus the subspace looks periodic in multiple "scales" simultaneously. As with periodic subspaces, an ultra-periodic subspace is defined by equations of the form (2), and this is how we will specify the subspace.

4. Preliminaries on function fields and algebraic-geometric codes

For convenience of the reader, we start with some background on global function fields over finite fields. The reader may refer to [30, 25] for detailed background on function fields and algebraic-geometric codes.

4.1. General background on function fields.

For a prime power q, let $\mathbb{F}_q$ be the finite field of q elements. An algebraic function field over $\mathbb{F}_q$ in one variable is a field extension $F \supset \mathbb{F}_q$ such that F is a finite algebraic extension of $\mathbb{F}_q(x)$ for some $x \in F$ that is transcendental over $\mathbb{F}_q$. The field $\mathbb{F}_q$ is called the full constant field of F if the algebraic closure of $\mathbb{F}_q$ in F is $\mathbb{F}_q$ itself. Such a function field is also called a global function field. From now on, we always denote by $F/\mathbb{F}_q$ a function field F with the full constant field $\mathbb{F}_q$.

4.1.1. *Valuations, Places, and Divisors.* A discrete valuation of $F/\mathbb{F}_q$ is a map from F to $\mathbb{Z} \cup \{+\infty\}$ satisfying certain properties (see [30, Definition 1.19]). Then each discrete valuation ν from $F/\mathbb{F}_q$ to $\mathbb{Z} \cup \{+\infty\}$ defines a valuation ring $O = \{f \in F : \nu(f) \geqslant 0\}$ that is a local ring [30, Theorem 1.1.13]. The maximal ideal P of O is given by $P = \{f \in F : \nu(f) > 0\}$ and it is called a *place*. We denote the valuation ν and the local ring O corresponding to P by ν_P and O_P, respectively. The residue class field O_P/P, denoted by

F_P, is a finite extension of $\mathbb{F}_q$. The extension degree $[F_P : \mathbb{F}_q]$ is called *degree* of P, denoted by $\deg(P)$. A place of degree one is called a *rational* place. For a nonzero function $z \in F$, the principal divisor of z is defined to be $\operatorname{div}(z) = \sum_{P\in\mathbb{P}_F} \nu_P(z)P$. The zero and pole divisors of z are defined to be $\operatorname{div}(z)_0 = \sum_{\nu_P(z)>0} \nu_P(z)P$ and $\operatorname{div}(z)_\infty = -\sum_{\nu_P(z)<0} \nu_P(z)P$, respectively. Then we have $\deg(\operatorname{div}(z)) = 0$, i.e, $\deg(\operatorname{div}(z)_0) = \deg(\operatorname{div}(z)_\infty)$. For two functions $f, g \in F$ and a place P, we have $\nu_P(f + g) \geqslant \min\{\nu_P(f), \nu_P(g)\}$ and the equality holds if $\nu_p(f) \neq \nu_P(g)$ (note that $\nu_P(0) = +\infty$). This implies that $f + g \neq 0$ if $\nu_P(f) \neq \nu_P(g)$.

If F is the rational function field $\mathbb{F}_q(x)$, then every discrete valuation of $F/\mathbb{F}_q$ is given by either ν_∞ or $\nu_{p(x)}$ for an irreducible polynomial $p(x)$, where ν_∞ is defined by $\nu_\infty(f/g) = \deg(g) - \deg(f)$ and $\nu_{p(x)}(f/g) = a - b$ with $p(x)^a||f$ and $p(x)^b||g$ for two nonzero polynomials $f, g \in \mathbb{F}_q[x]$. It is straightforward to verify that the degrees of places corresponding to ν_∞ and $\nu_{p(x)}$ are 1 and $\deg(p(x))$, respectively.

Let $\mathbb{P}_F$ denote the set of places of F. The divisor group, denoted by $\operatorname{Div}(F)$, is the free abelian group generated by all places in $\mathbb{P}_F$. An element $G = \sum_{P\in\mathbb{P}_F} n_P P$ of $\operatorname{Div}(F)$ is called a divisor of F, where $n_P = 0$ for almost all $P \in \mathbb{P}_F$. We denote n_p by $\nu_P(G)$. The support, denoted by $\operatorname{Supp}(G)$, of G is the set $\{P \in \mathbb{P}_F : n_P \neq 0\}$. Thus, $\operatorname{Supp}(G)$ of a divisor G is always a finite subset of $\mathbb{P}_F$.

4.1.2. *Constant field extension.* One of our code constructions will be based on evaluations of functions at rational points over a *subfield*. For this purpose, we will work with constant field extensions over $\mathbb{F}_{q^m}$ of a function field over a base field $\mathbb{F}_q$. We describe these now.

Let $F/\mathbb{F}_q$ be a function field. Fix an algebraic closure $\bar{F}$ of F. Then $\bar{F}$ contains the algebraic closure $\bar{\mathbb{F}}_q = \cup_{i=1}^{\infty}\mathbb{F}_{q^i}$ as well. Hence, for $m \geqslant 1$, $\bar{F}$ contains the extension field $\mathbb{F}_{q^m}$ of $\mathbb{F}_q$. The composite field $F_m := \mathbb{F}_{q^m} \cdot F$ is defined to be the smallest subfield of $\bar{F}$ that contains both F and $\mathbb{F}_{q^m}$. Then we have the following facts (see [30, Propositions 3.6.1 and 3.6.3]):

(i) the full constant field of F_m is $\mathbb{F}_{q^m}$;

(ii) each subset of F that is linearly independent over $\mathbb{F}_q$ remains so over F_m;

(iii) $[F_m : \mathbb{F}_{q^m}(x)] = [F : \mathbb{F}_q(x)]$ for any $x \in F \setminus \mathbb{F}_q$;

(iv) a place P of F of degree d splits into $\gcd(m, d)$ places of F_m of degree $d/\gcd(m, d)$ (in the case of rational function fields, this means that an irreducible polynomial over $\mathbb{F}_q$ of degree d is factorized into product of $\gcd(m, d)$ irreducible polynomials over $\mathbb{F}_{q^m}$ of degree $d/\gcd(m, d)$);

(v) genus of F_m is equal to genus of F.

A divisor $G = \sum_{P\in\mathbb{P}_F} n_P P$ of F can be viewed as the divisor $\sum_{P\in\mathbb{P}_F}\sum_{P'|P} n_P P'$ of F_m. We still denote this divisor of F_m by G. By (iv) of the above facts, a rational place P of F continues to be a rational place P' of F_m. The valuation ring of $P's$ is the tensor product of O_P with $\mathbb{F}_{q^m}$, i.e, $O_{P'} = O_P \otimes_{\mathbb{F}_q} \mathbb{F}_{q^m}$. If there is no confusion, we still denote P' by P.

4.1.3. *Riemann-Roch spaces.* For a divisor G of $F/\mathbb{F}_q$, we define the *Riemann-Roch space* associated with G by

$$\mathcal{L}(G) := \{f \in F^* : \operatorname{div}(f) + G \geqslant 0\} \cup \{0\},$$

where F^* denotes the set of nonzero elements of F. Then $\mathcal{L}(G)$ is a finite dimensional space over $\mathbb{F}_q$ and its dimension $\ell(G)$ is determined by the Riemann-Roch theorem which gives

$$\ell(G) = \deg(G) + 1 - \mathfrak{g} + \ell(W - G),$$

where $\mathfrak{g}$ is the genus of F and W is a canonical divisor of degree $2\mathfrak{g} - 2$. Therefore, we always have that $\ell(G) \geqslant \deg(G) + 1 - \mathfrak{g}$ and the equality holds if $\deg(G) \geqslant 2\mathfrak{g} - 1$ [30, Theorems 1.5.15 and 1.5.17].

Consider the finite extension $\mathbb{F}_{q^m}$ over $\mathbb{F}_q$ and the constant extension $F_m := \mathbb{F}_{q^m} \cdot F$ over F. As a divisor G of F can be viewed as a divisor of F_m, we can consider the Riemann-Roch space in F_m given by

$$\mathcal{L}_m(G) := \{f \in F_m^* : \operatorname{div}(f) + G \geqslant 0\} \cup \{0\}.$$

Then it is clear that $\mathcal{L}_m(G)$ contains $\mathcal{L}(G)$ and $\mathcal{L}_m(G)$ is a finite dimensional vector space over $\mathbb{F}_{q^m}$. Furthermore, $\mathcal{L}_m(G)$ is the tensor product of $\mathcal{L}(G)$ with $\mathbb{F}_{q^m}$ (see [29, Proposition 5.8 of Chapter II]). This implies that

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

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

4.1.4. *Automorphisms.* The automorphisms of the function field F that fix $\mathbb{F}_q$ are denoted by $\mathrm{Aut}(F/\mathbb{F}_q)$. For an automorphism $\phi \in \mathrm{Aut}(F/\mathbb{F}_q)$ and and a function $f \in F$, we denote by f^ϕ the action of ϕ on f. For a place P, define a map ν_{P^ϕ} from F to $\mathbb{Z} \cup \{+\infty\}$ given by $f \mapsto \nu_P(f^{\phi^{-1}})$. Then one can show that ν_{P^ϕ} indeed satisfies the properties given in [30, Definition 1.19] and hence it is a discrete valuation. The valuation ring O_{P^ϕ} of ν_{P^ϕ} is given by

$$\{h \in F : \nu_{P^\phi}(h) \geqslant 0\} = \{h \in F : \nu_P(h^{\phi^{-1}}) \geqslant 0\} \overset{h=f^\phi}{=} \{f^\phi \in F : \nu_P(f) \geqslant 0\} = \{f^\phi : f \in O_P\}.$$

and the maximal ideal of this valuation ring is $\{\phi(x) : x \in P\}$. Therefore, this maximal ideal is a place of F, denoted by P^ϕ. Moreover, ϕ induces an $\mathbb{F}_q$-isomorphism between the residue fields F_P and F_{P^ϕ}. Hence, we have $\deg(P) = \deg(P^\phi)$.

For a function f and a rational place $P \in \mathbb{P}_F$ with $\nu_P(f) \geqslant 0$, we denote by $f(P)$ the residue class of f in the residue class field F_P at P. If $\nu_P(f) \geqslant 0$ and $\nu_{P^\phi}(f) \geqslant 0$, then one has that $\nu_P(f^{\phi^{-1}}) \geqslant 0$. Furthermore, there is an $\mathbb{F}_q$-isomorphism between O_P and O_{P^ϕ} given by $f \mapsto f^\phi$. This induces the identity map between $F_P = \mathbb{F}_q$ and $F_{P^\phi} = \mathbb{F}_q$. Hence, $f(P) = f^\phi(P^\phi)$. Replacing f by $f^{\phi^{-1}}$ gives $f(P^\phi) = f^{\phi^{-1}}(P)$.

For a divisor $G = \sum_{P \in \mathbb{P}_F} m_P P$ we denote by G^ϕ the divisor $\sum_{P \in \mathbb{P}_F} m_P P^\phi$. Therefore, we have

$$\phi(\mathcal{L}(G)) := \{f^\phi : f \in \mathcal{L}(G)\} = \mathcal{L}(G^\phi).$$

Assume that $E/\mathbb{F}_q$ is a subfield of F and ϕ is an automorphism of $\mathrm{Aut}(F/E)$. Then for a divisor G of F that is invariant under ϕ, we have $\phi(\mathcal{L}(G)) = \mathcal{L}(G)$.

Next we consider the constant extension $F_m = \mathbb{F}_{q^m} \cdot F$. Let σ be the Frobenius automorphism $\mathbb{F}_{q^m}/\mathbb{F}_q$, i.e., $\sigma(\alpha) = \alpha^q$ for any $\alpha \in \mathbb{F}_{q^m}$. Then σ can be extended to an automorphism of $\mathrm{Aut}(F_m/F)$ given by $\sigma(f) = f$ for any $f \in F$ and $\sigma(\alpha) = \alpha^q$ for any $\alpha \in \mathbb{F}_{q^m}$. If P is a rational place of F, then P remains to be a rational place P' of F_m and hence $\sigma(O_{P'}) = \sigma(O_P \otimes_{\mathbb{F}_q} \mathbb{F}_{q^m}) = O_P \otimes_{\mathbb{F}_q} \mathbb{F}_{q^m} = O_{P'}$. Thus, we have $(P')^\sigma = P'$.

4.2. **Algebraic-geometric codes.** Let $\mathcal{P} = \{P_1, P_2, \dots, P_N\}$ be a set of N distinct rational places of a function field $F/\mathbb{F}_q$ of genus $\mathfrak{g}$. Let G be a divisor of F with $\mathrm{Supp}(G) \cap \mathcal{P} = \emptyset$. Then the algebraic-geometric code defined by

$$C(\mathcal{P}, G) := \{(f(P_1), f(P_2), \dots, f(P_N)) : \ f \in \mathcal{L}(G)\} \tag{3}$$

is an $\mathbb{F}_q$-linear code of length N. Furthermore, the dimension of $C(\mathcal{P}, G)$ is equal to $\ell(G)$ if $N > \deg(G)$.

The (generalized) Reed-Solomon codes can be realized under the above framework of algebraic-geometric codes. More precisely speaking, the Reed-Solomon codes are algebraic-geometric codes based on rational function fields. Let us give the detail on construction of the Reed-Solomon codes under the framework of algebraic-geometric codes.

Let $F = \mathbb{F}_q(x)$ be a rational function field. Let $\alpha_1, \alpha_2, \dots, \alpha_n$ be n distinct elements of $\mathbb{F}_q$. Denote by P_i the unique zero of $x - \alpha_i$ for $1 \leqslant i \leqslant n$ and put $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$. Let P_∞ be the unique pole of x. Put $G = (k-1)P_\infty$. Then the Riemann-Roch space $\mathcal{L}(G)$ is the $\mathbb{F}_q$-space consisting of polynomials of degree less than k. By definition, we have

$$\begin{aligned} C(\mathcal{P}, G) &= \{(f(P_1), f(P_2), \dots, f(P_N)) : \ f \in \mathcal{L}(G)\} \\ &= \{(f(\alpha_1), f(\alpha_2), \dots, f(\alpha_N)) : \ f \in \mathbb{F}_q[x], \ \deg(f) \leqslant k-1\}. \end{aligned}$$

The codes considered in this paper are variations of the above algebraic-geometric codes, namely, folded algebraic-geometric codes and algebraic-geometric codes with evaluation points in a subfield.

A folded algebraic-geometric code is a code with each coordinate being a column vector $(f(P), f(P^\sigma), \dots, f(P^{\sigma^{m-1}}))^T \in \mathbb{F}_q^m$ for a function $f \in \mathcal{L}(G)$, a rational place P and an automorphism $\sigma \in \mathrm{Aut}(F/\mathbb{F}_q)$, where T stands for transpose. This is a generalization of folded Reed-Solomon codes introduced in [13]. The main reason why a folded algebraic-geometric code is used is that once a position is transmitted correctly, then one gets m correct components $(f(P), f(P^\sigma), \dots, f(P^{\sigma^{m-1}}))$. Consequently, more interpolation equations are increased and list decoding radius is enlarged (see Lemma 6.2, for instance).

Similar to folded algebraic-geometric codes, introducing algebraic-geometric codes with evaluation points in a subfield is for purpose of increasing list decoding radius as well. We choose N rational places $P_1, P_2, \dots, P_N$ of a function field $F/\mathbb{F}_q$ and let σ be the Frobenius automorphism of $\mathbb{F}_{q^m}/\mathbb{F}_q$. Then one has $P_i^\sigma = P_i$ for all $1 \leqslant i \leqslant N$. Thus, once we have a correct position $f(P_i)$ for some function $f \in \mathcal{L}_m(G)$, we get correct information for other $m-1$ elements $f(P_i)^{\sigma^j} = f^{\sigma^j}(P_i)$ for $i = 1, 2, \dots, m-1$. As a result, list decoding radius is enlarged (see Lemma 7.7, for instance).

4.3. **Background on Hermitian tower.** In what follows, let r be a prime power and let $q = r^2$. We denote by $\mathbb{F}_q$ the finite field with q elements. The Hermitian function tower that we are going to use for our code construction was discussed in [27]. The reader may refer to [27] for the detailed background on the Hermitian function tower. The Hermitian tower is defined by the following recursive equations

$$x_{i+1}^r + x_{i+1} = x_i^{r+1}, \quad i = 1, 2, \dots, e-1. \tag{4}$$

Put $F_e = \mathbb{F}_q(x_1, x_2, \dots, x_e)$ for $e \geqslant 2$. We will assume that $r \geqslant 2e$.

4.3.1. *Rational places.* The function field F_e has $r^{e+1} + 1$ 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} come from the rational places lying over the unique zero P_α of $x_1 - \alpha$ for each $\alpha \in \mathbb{F}_q$. Note that for every $\alpha \in \mathbb{F}_q$, P_α splits completely in F_e, i.e., there are r^{e-1}

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

4.3.2. *Riemann-Roch spaces.* For an integer l, we consider the Riemann-Roch space defined by

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

By the Riemann-Roch theorem, its dimension $\ell(lP_\infty)$ is at least $l - \mathfrak{g}_e + 1$ and furthermore,

$$\ell(lP_\infty) = l - \mathfrak{g}_e + 1 \quad \text{if} \quad l \geqslant 2\mathfrak{g}_e - 1 \ ,$$

where $\mathfrak{g}_e$ is the genus of the function field F_e given by (6) below.

A basis over $\mathbb{F}_q$ of $\mathcal{L}(lP_\infty)$ can be explicitly constructed as follows

$$\left\{ x_1^{j_1} \cdots x_e^{j_e} : \ (j_1, \ldots, j_e) \in \mathbb{Z}_{\geqslant 0}^e, \ \sum_{i=1}^{e} j_i r^{e-i} (r+1)^{i-1} \leqslant l \right\}. \tag{5}$$

We stress that evaluating elements of $\mathcal{L}(lP_\infty)$ at the rational places of F_e (other than P_∞) is easy: we simply have to evaluate a linear combination of the monomials allowed in (5) at the tuples $(\alpha_1, \alpha_2, \ldots, \alpha_e) \in \mathbb{F}_q^e$ mentioned above. In other words, it is just evaluating an e-variate polynomial at a specific subset of r^{e+1} points of $\mathbb{F}_q^e$, and can be accomplished in polynomial time.

4.3.3. *Genus.* The genus $\mathfrak{g}_e$ of the function field F_e is given by

$$\mathfrak{g}_e = \frac{1}{2} \left(\sum_{i=1}^{e-1} r^e \left(1 + \frac{1}{r}\right)^{i-1} - (r+1)^{e-1} + 1 \right) \leqslant \frac{r^e}{2} \sum_{i=1}^{e} \binom{e}{i} \frac{1}{r^{i-1}} \leqslant \frac{er^e}{2} \sum_{i=1}^{e} \left(\frac{e}{r}\right)^{i-1} \leqslant er^e \tag{6}$$

where the last step used $r \geqslant 2e$.

4.3.4. *A useful automorphism.* Let γ be a primitive element of $\mathbb{F}_q$. Then for $i \geqslant 1$, one has $\gamma^{r(r+1)^i} = \gamma^{(r^2+r)(r+1)^{i-1}} = \gamma^{(1+r)(r+1)^{i-1}} = \gamma^{(r+1)^i}$. Consider the automorphism $\sigma \in \mathrm{Aut}(F_e/\mathbb{F}_q)$ defined by

$$\sigma : \ x_i \mapsto \gamma^{(r+1)^{i-1}} x_i \quad \text{for } i = 1, 2, \ldots, e.$$

Indeed, σ defines an automorphism $\sigma \in \mathrm{Aut}(F_e/\mathbb{F}_q)$ since after action of σ the equation (4) becomes $(\gamma^{(r+1)^i} x_{i+1})^r + \gamma^{(r+1)^i} x_{i+1} = (\gamma^{(r+1)^{i-1}} x_i)^{r+1}$, i.e., $x_{i+1}^r + x_{i+1} = x_i^{r+1}$ by cancelling $\gamma^{(r+1)^i}$ in both the sides. The order of σ is $q-1$ and furthermore, we have the following facts:

(i) Let P_0 be the unique common zero of $x_1, x_2, \ldots, x_e$ (this corresponds to the e-tuple $(0, 0, \ldots, 0)$), and P_∞ the unique pole of x_1. The automorphism σ keeps P_0 and P_∞ unchanged, i.e., $P_0^\sigma = P_0$ and ${P_\infty}^\sigma = P_\infty$,

(ii) Let $\mathbb{P}$ be the set of all the rational places which are neither P_∞ nor zeros of x_1. Then $|\mathbb{P}| = (q-1)r^{e-1}$. Moreover, σ divides $\mathbb{P}$ into r^{e-1} orbits and each orbit has $q-1$ places. For an integer m with $1 \leqslant m \leqslant q-1$, we can label Nm distinct elements $P_1, P_1^\sigma, \ldots, P_1^{\sigma^{m-1}}, \ldots, P_N, P_N^\sigma, \ldots, P_N^{\sigma^{m-1}}$ in $\mathbb{P}$, as long as $N \leqslant r^{e-1} \left\lfloor \frac{q-1}{m} \right\rfloor$.

4.4. **Background on Garcia-Stichtenoth tower.** Again let r be a prime power and let $q = r^2$. We denote by $\mathbb{F}_q$ the finite field with q elements. The Garcia-Stichtenoth towers that we are going to use for our code construction were discussed in [6, 7]. The reader may refer to [6, 7] 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 [7]

$$x_{i+1}^r + x_{i+1} = \frac{x_i^r}{x_i^{r-1}+1}, \quad i = 1, 2, \ldots, e-1. \tag{7}$$

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

4.4.1. *Rational places.* 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, $2 \leqslant i \leqslant e$).

4.4.2. *Riemann-Roch spaces.* As shown in [28], every function of K_e with a pole only at P_∞ has an expression of the form

$$x_1^a \left(\sum_{i_1=0}^{(e-2)r+1} \sum_{i_2=0}^{r-1} \cdots \sum_{i_e=0}^{r-1} c_{\mathbf{i}} h_1 \frac{x_1^{i_1} x_2^{i_2} \cdots x_e^{i_e}}{\pi_2 \ldots \pi_{e-1}} \right), \tag{8}$$

where $a \geqslant 0, c_{\mathbf{i}} \in \mathbb{F}_q$, and for $1 \leqslant j < e$, $h_j = x_j^{r-1} + 1$ and $\pi_j = h_1 h_2 \ldots h_j$. Moreover, Shum et al. [28] present an algorithm running in time polynomial in l that outputs a basis of over $\mathbb{F}_q$ of $\mathcal{L}(lP_\infty)$ explicitly in the above form.

We stress that evaluating elements of $\mathcal{L}(lP_\infty)$ at the rational places of $\mathbb{P}$ is easy: we simply have to evaluate a linear combination of the monomials allowed in (8) at the tuples $(\alpha_1, \alpha_2, \ldots, \alpha_e) \in \mathbb{P}$ (note that $h_i(P), \pi_j(P) \in \mathbb{F}_q^*$ for every $P \in \mathbb{P}$). In other words, it is just evaluating an e-variate polynomial at a specific subset of $r^{e-1}(r^2 - r)$ points of $\mathbb{F}_q^e$, and can be accomplished in polynomial time.

4.4.3. *Genus.* The genus $\mathfrak{g}_e$ of the function field K_e is given by

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

Thus the genus $\mathfrak{g}_e$ is at most r^e. (Compare this with the er^e bound for the Hermitian tower; this smaller genus is what allows to pick e as large as we want in the Garcia-Stichtenoth tower, while keeping the field size q fixed.)

4.4.4. *A useful automorphism.* Let γ be a primitive element of $\mathbb{F}_q$ and consider the automorphism $\sigma \in \mathrm{Aut}(K_e/\mathbb{F}_q)$ defined by

$$\sigma: \ x_i \mapsto \gamma^{r+1}x_i \quad \text{for } i = 1, 2, \dots, e.$$

Indeed, σ defines an automorphism $\sigma \in \mathrm{Aut}(K_e/\mathbb{F}_q)$ since after action of σ the equation (7) becomes $(\gamma^{r+1}x_{i+1})^r + \gamma^{r+1}x_{i+1} = \frac{(\gamma^{r+1}x_i)^r}{(\gamma^{r+1}x_i)^{r-1}+1}$, i.e, $x_{i+1}^r + x_{i+1} = \frac{x_i^r}{x_i^{r-1}+1}$ by cancelling γ^{r+1} on both the sides (note the fact that $\gamma^{(r+1)r} = \gamma^{r+1}$ and $\gamma^{r^2-1} = 1$). The order of σ is $r-1$ and furthermore, we have the following facts:

(i) σ keeps P_∞ unchanged, i.e., ${P_\infty}^\sigma = P_\infty$;

(ii) Let $\mathbb{P}$ be the set of all the rational places lying over $x_1 - \alpha$ for all $\alpha \in \mathbb{F}_q$ with $\alpha^r + \alpha \neq 0$. Then $|\mathbb{P}| = (r-1)r^e$. Moreover, σ divides $\mathbb{P}$ into r^e orbits and each orbit has $r-1$ places. For an integer m with $1 \leqslant m \leqslant r-1$, we can label Nm distinct elements
$$P_1, P_1^\sigma, \dots, P_1^{\sigma^{m-1}}, \dots, P_N, P_N^\sigma, \dots, P_N^{\sigma^{m-1}}$$
in $\mathbb{P}$, as long as $N \leqslant r^e \left\lfloor \frac{r-1}{m} \right\rfloor$.

5. Local expansions and encoding

Similar to the Laurent series expansion of a complex function $f(z)$ in the neighborhood of a complex number, one can write functions in a function field as a power series (with finitely many negative powers) around a place P, called the *local expansion around P*. Local expansions play an important role in the encoding and decoding of the codes we construct, and we discuss them separately in this section.

5.1. **Local expansion at a place.**

Let $F/\mathbb{F}_q$ be a function field and let P be a rational place. An element t of F is called a local parameter at P if $\nu_p(t) = 1$ (such a local parameter always exists) — intuitively this is a function which has a simple zero at P, similar to how $(z-1)$ has a simple zero at 1. For a nonzero function $f \in F$ with $\nu_P(f) \geqslant v$, we have $\nu_P\left(\frac{f}{t^v}\right) \geqslant 0$. Put $f_v = \left(\frac{f}{t^v}\right)(P)$, i.e., f_v is the value of the function f/t^v at P. Note that the function $f/t^v - f_v$ satisfies $\nu_P\left(\frac{f}{t^v} - f_v\right) \geqslant 1$, hence we know that $\nu_P\left(\frac{f - f_v t^v}{t^{v+1}}\right) \geqslant 0$. Put $f_{v+1} = \left(\frac{f - a_v t^v}{t^{v+1}}\right)(P)$. Then $\nu_P(f - f_v t^v - f_{v+1}t^{v+1}) \geqslant v+2$.

Assume that we have obtained a sequence $\{f_r\}_{r=v}^m$ $(m > v)$ of elements of $\mathbb{F}_q$ such that $\nu_P(f - \sum_{r=v}^k f_r t^r) \geqslant k+1$ for all $v \leqslant k \leqslant m$. Put $f_{m+1} = \left(\frac{f - \sum_{r=v}^m f_r t^r}{t^{m+1}}\right)(P)$. Then $\nu_P(f - \sum_{r=v}^{m+1} f_r t^r) \geqslant m+2$. In this way we continue our construction of f_r. Then we obtain an infinite sequence $\{f_r\}_{r=v}^\infty$ of elements of $\mathbb{F}_q$ such that $\nu_P(f - \sum_{r=v}^m f_r t^r) \geqslant m+1$ for all $m \geqslant v$. We summarize the above construction in the formal expansion

$$f = \sum_{r=v}^{\infty} f_r t^r, \tag{9}$$

which is called the *local expansion* of f at P.

It is clear that the local expansion of a function depends on the choice of the local parameter t. Note that if a power series $\sum_{i=v}^\infty a_i t^i$ satisfies $\nu_P(f - \sum_{i=v}^m a_i t^i) \geqslant m+1$ for all $m \geqslant v$, then it is a local expansion of f. The above procedure shows that finding a local expansion at a rational place is very efficient as long as the computation of evaluations of functions at this place is easy.

If f belongs to a Riemann-Roch space $\mathcal{L}(G)$ with $\deg(G) = d$. Denote $\nu_P(G)$ by v, then the first $d+1$ coefficients $a_v, a_{v+1}, \ldots, a_{v+d}$ in (9) determines the function f. To see this, assume that g is a function of $\mathcal{L}(G)$ with the first $d+1$ coefficients in its local expansion equal to those of f. Then we have $f - g \in \mathcal{L}(G - (d+1)P)$ which is the zero vector space. This implies that $f = g$.

5.2. **Encodings using local expansion.** An algebraic-geometric code as defined in (3) encodes messages which belong to a Riemann-Roch space $\mathcal{L}(G)$. The most common instantiation, which suffices for most purposes, is to take $G = lP_\infty$ for some rational place P_∞ (though of as the place at infinity). For such spaces, one can compute bases for the Riemann-Roch spaces explicitly in many cases, including the Hermitian and Garcia-Stichtenoth towers as mentioned in Sections 4.3 and 4.4. For k linearly independent functions $g_1, g_2, \ldots, g_k$ in $\mathcal{L}(lP_\infty)$, one can interpret a message vector $(a_1, \ldots, a_k) \in \mathbb{F}_q^k$ as the function $f = \sum_{i=1}^k a_i g_i \in \mathcal{L}(lP_\infty)$ and then encode it.

For our decoding, we will actually recover the message $f \in \mathcal{L}(lP_\infty)$ in terms of the coefficients of its local expansion around a rational place P

$$f = x^{-\nu}(f_0 + f_1 x + f_2 x^2 + \cdots) \tag{10}$$

where x is a local parameter at P. The place P may or may not equal P_∞ — when we instantiate the algorithm of this section for the Hermitian tower, we will use a place different than P_∞ for P, whereas for the Garcia-Stichtenoth tower, we will use $P = P_\infty$. The description of the algorithm and its analysis in this section will be general and cover both cases. Let $\nu_P(P_\infty) = \nu$ with $\nu = 0$ if $P_\infty \neq P$ and $\nu = l$ if $P = P_\infty$. Realizing that one must work in this power series representation is one of the key insights in this work behind the extension of the linear-algebraic folded Reed-Solomon list decoding algorithm [16] to the algebraic-geometric setting.

Given this, we will find it convenient to let the message vector consist of $(f_0, f_1, \ldots, f_{k-1}) \in \mathbb{F}^k$ (k being the dimension of the code), which we will then map to a function f in an appropriate Riemann-Roch space. Here we denote the field by $\mathbb{F}$, to capture both $\mathbb{F} = \mathbb{F}_q$ and $\mathbb{F} = \mathbb{F}_{q^m}$ when we work with constant field extensions F_m and seek functions $f \in \mathcal{L}_m(lP_\infty)$. Likewise, we use the common notation $\mathcal{L}_\mathbb{F}(lP_\infty)$ to denote $\mathcal{L}(lP_\infty)$ when $\mathbb{F} = \mathbb{F}_q$, and $\mathcal{L}_m(lP_\infty)$ when $\mathbb{F} = \mathbb{F}_{q^m}$.

If we seek a k-dimensional message space, it is natural to let the message functions belong to $\mathcal{L}((k + \mathfrak{g} - 1)P_\infty)$ which has dimension exactly k by the Riemann-Roch theorem (when k is at least the genus $\mathfrak{g}$, which will always hold for our codes). However, we desire to index the messages of the code instead by the first k coefficients $(f_0, f_1, \ldots, f_{k-1})$ of the local expansion of the function f at P. Therefore we require that for every $(f_0, f_1, \ldots, f_{k-1})$ there is a $f \in \mathcal{L}_\mathbb{F}(lP_\infty)$ whose local expansion at P has the f_i's as the first k coefficients. We can ensure by taking a slightly larger value of l, namely $l = k + 2\mathfrak{g} - 1$ as we argue below. Since the genus will be much smaller than the code length, we can afford the resulting small loss in distance and list-decoding radius.

We will recover the message in terms of the coefficients of its local expansion at P.

Restricting message functions using local expansions. 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 $(k + 2\mathfrak{g} - 1)$ tuples over $\mathbb{F}$ arise in the local expansion of functions in the k-dimensional subspace $\mathcal{L}_\mathbb{F}((k + 2\mathfrak{g} - 1)P_\infty)$. Below we show that we can find a k-dimensional subspace of $\mathcal{L}_{\mathbb{F}}((k+2\mathfrak{g}-1)P_\infty)$ such that their top k local expansion coefficients give rise to all k-tuples over $\mathbb{F}$.

Lemma 5.1. *There exist a set of functions $\{g_1, g_2, \dots, g_k\}$ in $\mathcal{L}_{\mathbb{F}}((k+2\mathfrak{g}-1)P_\infty)$ 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* (9) *for g_i at P is non-singular.*

Proof. Let $\{\psi_1, \psi_2, \dots, \psi_{\mathfrak{g}}\}$ be a basis of $\mathcal{L}_{\mathbb{F}}((k+2\mathfrak{g}-1)P_\infty - kP)$. Extend this basis to a basis $\{\psi_1, \psi_2, \dots, \psi_{\mathfrak{g}}, g_1, g_2, \dots, g_k\}$ of $\mathcal{L}_{\mathbb{F}}((k+2\mathfrak{g}-1)P_\infty)$. We claim that the functions $\{g_1, g_2, \dots, g_k\}$ 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 local expansion $\sum_{i=k}^\infty a_i T^i$ at P for some $a_i \in \mathbb{F}$. Therefore, the function $\sum_{i=1}^k \lambda_i g_i$ belongs to the space $\mathcal{L}_{\mathbb{F}}((k+2\mathfrak{g}-1)P_\infty - kP)$, i.e., $\sum_{i=1}^k \lambda_i g_i$ is a linear combination of $\psi_1, \psi_2, \dots, \psi_g$. This forces that all λ_i are equal to 0 since $\{\psi_1, \dots, \psi_g, g_1, g_2, \dots, g_k\}$ is linearly independent. This completes the proof. □

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 $\{g_1, g_2, \dots, g_k\}$ of $\mathcal{L}_{\mathbb{F}}((k+2\mathfrak{g}-1)P_\infty)$ as in Lemma 5.1. After elementary row operations on the matrix A defined in Lemma 5.1, 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 local expansion $T^{i-1} + \sum_{j=k}^\infty \lambda_{ij} T^j$ for some $\lambda_{ij} \in \mathbb{F}$. Now we encode each message $(a_1, a_2, \dots, a_k) \in \mathbb{F}^k$ to the codeword $(f(P_1), f(P_2), \dots, f(P_N))$, where $f = \sum_{i=1}^k a_i g_i$.

Now define the map $\phi_P : \mathbb{F}^k \to \mathcal{L}_{\mathbb{F}}((k+2\mathfrak{g}-1)P_\infty)$ by sending $(a_1, a_2, \dots, a_k) \in \mathbb{F}^k$ to $\sum_{i=1}^k a_i g_i$. We record the above fact for easy reference below.

Claim 5.2. *The map $\phi_P : \mathbb{F}^k \to \mathcal{L}_{\mathbb{F}}((k+2\mathfrak{g}-1)P_\infty)$ is $\mathbb{F}_q$-linear and injective. Furthermore, we can compute a representation of this linear transformation using* $\mathrm{poly}(N, \mathfrak{g})$ *operations over $\mathbb{F}_q$, and the map itself can be evaluated using* $\mathrm{poly}(N, \mathfrak{g})$ *operations over $\mathbb{F}_q$ provided that local expansion of the basis elements of $\mathcal{L}((k+2\mathfrak{g}-1)P_\infty)$ at P can be computed using* $\mathrm{poly}(N, \mathfrak{g})$ *operations over $\mathbb{F}_q$.*

6. Folded algebraic-geometric codes and their list decoding

In this section, we will describe a variation of algebraic-geometric codes, namely, folded algebraic-geometric codes and their list decoding. For convenience, we will focus on one-point algebraic-geometric codes though this is not in any way a necessary restriction for our approach.

6.1. Folded algebraic-geometric codes.

Let $F/\mathbb{F}_q$ be a function field. To construct our folded codes, we assume that there exists a global function field F with the full constant field $\mathbb{F}_q$ having the following property:

(i) There exists an automorphism σ in $\mathrm{Aut}(F/\mathbb{F}_q)$ of order at least m;
(ii) F has mN distinct rational places $P_1, P_1^\sigma, \dots, P_1^{\sigma^{m-1}}, P_2, P_2^\sigma, \dots, P_2^{\sigma^{m-1}}, \dots, P_N, P_N^\sigma, \dots, P_N^{\sigma^{m-1}}$;
(iii) F has a rational place P_∞ such that P_∞ is fixed under σ, i.e., ${P_\infty}^\sigma = P_\infty$; and $P_i^{\sigma^j} \neq P_\infty$ for all $1 \leqslant i \leqslant N$ and $0 \leqslant j \leqslant m-1$.

A folded algebraic geometric code can be defined as follows.

Definition 4 (Folded AG codes)**.** The folded code from F with parameters N, l, q, m, denoted by $\mathrm{F}(N, l, q, m)$, encodes a message function $f \in \mathcal{L}(lP_\infty)$ as

$$\pi: \quad f \mapsto \left(\begin{bmatrix} f(P_1) \\ f(P_1^\sigma) \\ \vdots \\ f(P_1^{\sigma^{m-1}}) \end{bmatrix}, \begin{bmatrix} f(P_2) \\ f(P_2^\sigma) \\ \vdots \\ f(P_2^{\sigma^{m-1}}) \end{bmatrix}, \dots, \begin{bmatrix} f(P_N) \\ f(P_N^\sigma) \\ \vdots \\ f(P_N^{\sigma^{m-1}}) \end{bmatrix} \right) \in \left(\mathbb{F}_q^m\right)^N . \tag{11}$$

We will abuse notation and for clarity refer to the encoding map π also as $\mathrm{F}(N, l, q, m)$.

Note that the folded code $\mathrm{F}(N, l, q, m)$ has the alphabet $\mathbb{F}_q^m$ and it is $\mathbb{F}_q$-linear. Furthermore, $\mathrm{F}(N, l, q, m)$ has the following parameters.

Lemma 6.1. *If $l < mN$, then the above code $\mathrm{F}(N, l, q, m)$ is an $\mathbb{F}_q$-linear code with alphabet size q^m, rate at least $\frac{l-\mathfrak{g}+1}{Nm}$, and minimum distance at least $N - \frac{l}{m}$, where $\mathfrak{g}$ is the genus of F.*

Proof. It is clear that the map π in (11) is $\mathbb{F}_q$-linear and the kernel of π is

$$\mathcal{L}\Big(lP_\infty - \sum_{i=1}^{N} \sum_{j=0}^{m-1} P_i^{\sigma^j}\Big)$$

which is $\{0\}$ under the condition that $l < mN$. Thus, π is injective. Hence, the rate is at least $\frac{l-\mathfrak{g}+1}{Nm}$ by the Riemann-Roch theorem. To see the minimum distance, let f be a nonzero function in $\mathcal{L}(lP_\infty)$ and assume that I is the support of $\pi(f)$. Then the Hamming weight $\mathrm{wt}_H(\pi(f))$ of $\pi(f)$ is $|I|$ and $f \in \mathcal{L}\left(lP_\infty - \sum_{i \notin I} \sum_{j=0}^{m-1} P_i^{\sigma^j}\right)$. Thus, $0 \leqslant \deg\left(lP_\infty - \sum_{i \notin I} \sum_{j=0}^{m-1} P_i^{\sigma^j}\right) = l - m(N - |I|)$, i.e., $\mathrm{wt}_H(\pi(f)) = |I| \geqslant N - \frac{l}{m}$. This completes the proof. □

6.2. **Encoding of code using local expansions.** For our decoding, we will actually recover the message $f \in \mathcal{L}(lP_\infty)$ in terms of the coefficients of its power series expansion around a rational place P

$$f = x^{-\nu}(f_0 + f_1 x + f_2 x^2 + \cdots) \tag{12}$$

where x is a local parameter at P. The place P may or may not equal P_∞ – when we instantiate the algorithm of this section for the Hermitian tower, we will use a place different than P_∞ for P, whereas for the Garcia-Stichtenoth tower, we will use $P = P_\infty$. The reason for different choice of P is that we need an explicit and simple local parameter at P such that this local parameter still has an explicit and simple form after action of automorphism. The description of the algorithm and its analysis in this section will be general and cover both cases by letting $\nu_P(P_\infty) = \nu$ with $\nu = 0$ if $P_\infty \neq P$ and $\nu = l$ if $P = P_\infty$. Realizing that one must work in this power series representation is one of the key insights in this work behind the extension of the linear-algebraic folded Reed-Solomon list decoding algorithm [16] to the algebraic-geometric setting.

As already mentioned in Section 5, one can injectively map the top k coefficients of the above local expansion (12) into functions in $\mathcal{L}(lP_\infty)$ for $l = k + 2\mathfrak{g} - 1$. We will now redefine a version of the folded algebraic-geometric code that maps $\mathbb{F}_q^k$ to $(\mathbb{F}_q^m)^N$ by composing the folded encoding (11) from the original Definition 4 with the map $\phi_P : \mathbb{F}_q^k \to \mathcal{L}((k + 2\mathfrak{g} - 1)P_\infty)$ promised in Claim 5.2.

Definition 5 (Folded algebraic-geometric code using local expansion)**.** *The folded algebraic-geometric code* $\widetilde{\mathrm{F}}(N,k,q,m)$ *maps*

$$\mathbf{f}=(f_0,f_1,\dots,f_{k-1})\in\mathbb{F}_q^k \quad\mapsto\quad \mathrm{F}(N,k+2\mathfrak{g}-1,q,m)(\phi_P(\mathbf{f}))\in(\mathbb{F}_q^m)^N\ ,$$

where $\mathrm{F}(\dots)$ *is the folded AG code from Definition 4.*

The rate of the above code equals $k/(Nm)$ and its distance is at least $N-(k+2\mathfrak{g}-1)/m$.

6.3. **List decoding folded algebraic-geometric codes.** We now present a list decoding algorithm for the above codes. The algorithm follows the linear-algebraic list decoding algorithm for folded Reed-Solomon codes.

Suppose a codeword (11) encoded from $f\in\mathcal{L}((k+2\mathfrak{g}-1)P_\infty)$ was transmitted and received as

$$\mathbf{y}=\begin{pmatrix} y_{1,1} & y_{2,1} & & y_{N,1}\\ y_{1,2} & y_{2,2} & & \vdots\\ & & \ddots & \\ y_{1,m} & \cdots & & y_{N,m}\end{pmatrix}, \tag{13}$$

where some columns are erroneous. Let $s\geqslant 1$ be an integer parameter associated with the decoder.

Lemma 6.2. *Given a received word as in* (13)*, we can find a nonzero linear polynomial in* $F[Y_1,Y_2,\dots,Y_s]$ *of the form*

$$Q(Y_1,Y_2,\dots,Y_s)=A_0+A_1Y_1+A_2Y_2+\cdots+A_sY_s \tag{14}$$

satisfying

$$Q(y_{i,j+1},y_{i,j+2},\cdots,y_{i,j+s})=A_0(P_i^{\sigma^j})+A_1(P_i^{\sigma^j})y_{i,j+1}+\cdots+A_s(P_i^{\sigma^j})y_{i,j+s}=0 \tag{15}$$

for $i=1,2,\dots,N$ *and* $j=0,1,\dots,m-s$*. The coefficients* A_i *of* Q *satisfy* $A_i\in\mathcal{L}(\kappa P_\infty)$ *for* $i=1,2,\dots,s$ *and* $A_0\in\mathcal{L}((\kappa+(k+2\mathfrak{g}-1))P_\infty)$ *for a "degree" parameter* κ *chosen as*

$$\kappa=\left\lceil\frac{N(m-s+1)-(k+2\mathfrak{g}-1)+(s+1)(\mathfrak{g}-1)+1}{s+1}\right\rceil. \tag{16}$$

Proof. Let u and v be dimensions of $\mathcal{L}(\kappa P_\infty)$ and $\mathcal{L}((\kappa+(k+2\mathfrak{g}-1))P_\infty)$, respectively. Let $\{x_1,\dots,x_u\}$ be an $\mathbb{F}_q$-basis of $\mathcal{L}(\kappa P_\infty)$ and extend it to an $\mathbb{F}_q$-basis $\{x_1,\dots,x_v\}$ of $\mathcal{L}((d+k+2\mathfrak{g}-1)P_\infty)$. Then A_i is an $\mathbb{F}_q$-linear combination of $\{x_1,\dots,x_u\}$ for $i=1,2,\dots,s$ and A_0 is an $\mathbb{F}_q$-linear combination of $\{x_1,\dots,x_v\}$. Determining the functions A_i is equivalent to determining the coefficients in the combinations of A_i. Thus, there are in total $su+v$ degrees of freedoms to determine $A_0,A_1,\dots,A_s$. By the Riemann-Roch theorem, the number of degrees of freedoms is at least $s(\kappa-\mathfrak{g}+(k+2\mathfrak{g}-1))+(\kappa+k+2\mathfrak{g}-1)-\mathfrak{g}+1$.

On the other hand, there are in total $N(m-s+1)$ equations in (15). Thus, there must be one nonzero solution by the condition (16), i.e., $Q(Y_1,Y_2,\dots,Y_s)$ is a nonzero polynomial. □

Lemma 6.3. *If* f *is a function in* $\mathcal{L}(lP_\infty)$ *whose encoding (11) agrees with the received word* $\mathbf{y}$ *in at least* t *columns with*

$$t>\frac{\kappa+l}{m-s+1}\ ,$$

then $Q(f,f^{\sigma^{-1}},\dots,f^{\sigma^{-(s-1)}})$ *is the zero function, i.e.,*

$$A_0+A_1f+A_2f^{\sigma^{-1}}+\cdots+A_sf^{\sigma^{-(s-1)}}=0. \tag{17}$$

Proof. Since $P_\infty = {P_\infty}^\sigma$, we have $f^{\sigma^i} \in \mathcal{L}(lP_\infty)$ for all $i \in \mathbb{Z}$. Thus, it is clear that $Q(f, f^{\sigma^{-1}}, \dots, f^{\sigma^{-(s-1)}})$ is a function in $\mathcal{L}((\kappa + l)P_\infty)$.

Let us assume that $I \subseteq \{1, 2, \dots, N\}$ is the index set such that the ith columns of $\mathbf{y}$ and $\pi(f)$ agree if and only if $i \in I$. Then we have $|I| \geqslant t$. For every $i \in I$ and $0 \leqslant j \leqslant m - s$, we have by (15)

$$\begin{aligned} 0 &= A_0(P_i^{\sigma^j}) + A_1(P_i^{\sigma^j})y_{i,j+1} + A_2(P_i^{\sigma^j})y_{i,j+2} + \cdots + A_s(P_i^{\sigma^j})y_{i,j+s} \\ &= A_0(P_i^{\sigma^j}) + A_1(P_i^{\sigma^j})f(P_i^{\sigma^j}) + A_2(P_i^{\sigma^j})f(P_i^{\sigma^{j+1}})) + \cdots + A_s(P_i^{\sigma^j})f(P_i^{\sigma^{j+s-1}}) \\ &= A_0(P_i^{\sigma^j}) + A_1(P_i^{\sigma^j})f(P_i^{\sigma^j}) + A_2(P_i^{\sigma^j})f^{\sigma^{-1}}(P_i^{\sigma^j}) + \cdots + A_s(P_i^{\sigma^j})f^{\sigma^{-s+1}}(P_i^{\sigma^j}) \\ &= \left(A_0 + A_1 f + A_2 f^{\sigma^{-1}} + \cdots + A_s f^{\sigma^{-s+1}}\right)(P_i^{\sigma^j}), \end{aligned}$$

i.e., $P_i^{\sigma^j}$ is a zero of $Q(f, f^\sigma, \dots, f^{\sigma^{s-1}})$. Hence, $Q(f, f^{\sigma^{-1}}, \dots, f^{\sigma^{-(s-1)}})$ is a function in $\mathcal{L}\left((\kappa + l)P_\infty - \sum_{i\in I}\sum_{j=0}^{m-s} P_i^{\sigma^j}\right)$. Our desired result follows from the fact that $\deg\left((\kappa + l)P_\infty - \sum_{i\in I}\sum_{j=0}^{m-s} P_i^{\sigma^j}\right) < 0$. □

By Lemma 6.3, we know that all candidate functions f in our list must satisfy equation (17). In other words, we have to study the solution set of equation (17). The method used in [13] for decoding the Reed-Solomon codes is to construct an irreducible polynomial $h(x)$ of degree $q - 1$ such that every polynomial f satisfies $f^{\sigma^{-1}} \equiv f^q \mod h$. Then the solution set of (11) is the same as the solution set of the equation $A_0 + A_1 f + A_2 f^q + \cdots + A_s f^{q^{s-1}} \equiv 0 \mod h$ since $\deg(f) < q - 1 = \deg(h)$. Thus, there are at most q^{s-1} solutions for equation (11). This method does not work for folded algebraic-geometric codes. To upper bound list size of a folded algebraic-geometric code, we require an automorphisms of $\text{Aut}(F/F_q)$ with order proportional to the genus $\mathfrak{g}$ of F. However, it was proved in [24] that the order of an automorphisms of $\text{Aut}(F/\mathbb{F}_q)$ is upper bounded $O(\mathfrak{g}/\log \mathfrak{g})$.

In this paper, we will analyze the solutions of the equation (11) by considering local expansions at a certain point. This local expansion method guarantees a structured list of exponential size. Through precoding by using the structure in the list, we will be able to obtain an explicit construction of subcodes of these codes with polynomial time list decoding.

Solving the functional equation for f. Recall that our goal is to recover the top k coefficients $(f_0, f_1, \dots, f_{k-1})$ of the local expansion $f = x^{-\nu}\sum_{j=0}^{\infty} f_j x^j$ at P, based on the functional equation (17) that f satisfies.

We now prove that $(f_0, f_1, \dots, f_{k-1})$ for f satisfying Equation (17) belong to a periodic subspace (in the sense of Definition 1) of not too large dimension.

Lemma 6.4. *Let P and P_∞ be two rational places of F (P and P_∞ can be the same) and let $f \in \mathcal{L}((k + 2\mathfrak{g} - 1)P_\infty)$. Assume that $\sigma \in \text{Aut}(F/\mathbb{F}_q)$ is an automorphism satisfying ${P_\infty}^\sigma = P_\infty$. Let $x \in F$ be a local parameter at P satisfying $x^\sigma = \frac{x}{\xi}$ for an element $\xi \in \mathbb{F}_q^*$ of order p. Put $\nu = k + 2\mathfrak{g} - 1$ if $P = P_\infty$ and 0 otherwise.*

Then the set of solutions $(f_0, f_1, \dots, f_{k-1}) \in \mathbb{F}_q^k$ such that $f = x^{-\nu}(f_0 + f_1 x + f_2 x^2 + \cdots) \in \mathcal{L}((k + 2\mathfrak{g} - 1)P_\infty)$ obeys the equation

$$A_0 + A_1 f + A_2 f^{\sigma^{-1}} + \cdots + A_s f^{\sigma^{-(s-1)}} = 0, \tag{18}$$

when the A_i's obey the pole order restrictions of Lemma 6.2 and at least one A_i is nonzero, is an $(s - 1, p)$-ultra periodic subspace of $\mathbb{F}_q^k$.

Further, there are at most q^{Nm+s+1} possible choices of this subspace over varying choices of the A_i's.

Proof. Let $u = \min\{\nu_P(A_i) : i = 1, 2, \ldots, s\}$. Then we have $\nu_P(A_0) = \nu_P(-\sum_{i=1}^{s} A_i f^{-\sigma^{i-1}}) \geqslant \min\{\nu_P(A_i f^{-\sigma^{i-1}})) : i = 1, 2, \ldots, s\} \geqslant \min\{\nu_P(A_i) - \nu : i = 1, 2, \ldots, s\} = u - \nu$. Each A_i has a local expansion at P:

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

for $i = 1, \ldots, s$, and $A_0 = x^{u-\nu} \sum_{j=0}^{\infty} a_{0,j} x^j$ which can be efficiently computed from the basis representation of the A_i's. From the definition of u, one knows that the polynomial

$$B_0(X) := a_{1,0} + a_{2,0}X + \cdots + a_{s,0}X^{s-1}$$

is nonzero. Assume that at P, the function f has a local expansion $x^{-\nu} \sum_{j=0}^{\infty} f_j x^j$. Then $f^{\sigma^{-i}}$ has a local expansion at P as follows

$$f^{\sigma^{-i}} = \xi^{-i\nu} x^{-\nu} \sum_{j=0}^{\infty} \xi^{ij} f_j x^j.$$

By direct inspection, we see that for every $d \geqslant 0$, the coefficient of $x^{d+u-\nu}$ in the local expansion of $A_0 + A_1 f + A_2 f^{\sigma^{-1}} + \cdots + A_s f^{\sigma^{-(s-1)}}$ equals

$$0 = B_0(\xi^{d-\nu}) f_d + \sum_{j=1}^{d} B_j(\xi^{d-j-\mu}) f_{d-j} + a_{0,d}, \tag{19}$$

where similarly to $B_0(X)$, the degree $(s-1)$ polynomials $B_j(X)$, $j \geqslant 1$, are defined as

$$B_j(X) = a_{1,j} + a_{2,j}X + \cdots + a_{s,j}X^{s-1}.$$

Hence, f_d is uniquely determined by $f_0, \ldots, f_{d-1}$ as long as $B_0(\xi^{d-\nu}) \neq 0$.

Let $S := \{0 \leqslant i \leqslant p-1 : B_0(\xi^i) = 0\}$. Then it is clear that $|S| \leqslant s-1$ since the order of ξ is p so the powers ξ^i are distinct for $0 \leqslant i \leqslant p-1$, and $B_0(X)$ has degree at most $s-1$. Thus, $B_0(\xi^{d-\nu}) \neq 0$ if and only if $d - \nu \mod p \notin S$; and in this case f_d is a fixed affine linear combination of f_j for $0 \leqslant j < d$.

Let W be the solution space $(z_0, z_1, \ldots, z_{p-1}) \in \mathbb{F}_q^p$ of the equation system

$$B_0(\xi^{d-\mu}) z_d + \sum_{j=1}^{d} B_j(\xi^{d-\mu-j}) z_j = 0 \text{ for } d = 0, 1, \ldots, p-1 . \tag{20}$$

The above argument shows that W is a subspace of $\mathbb{F}_q^p$ of dimension at most $s-1$.

We now claim that the solutions to (19) for $0 \leqslant d < k$ form an $(s-1, p)$-periodic subspace of $\mathbb{F}_q^k$ with $W \subset \mathbb{F}_q^p$ as the recurring subspace. This is immediate by inspecting the system of equations (19) satisfied by the f_i's and the system (20) defining the subspace W. Indeed, once the values of f_i, $0 \leqslant i < p(j-1)$ are fixed, the possible choices for the j'th block of p coordinates, $f_{p(j-1)}, \cdots, f_{pj-1}$, lie in an affine shift of W. Further, this shift is an explicit affine combination of the f_i's for $0 \leqslant i < p(j-1)$ (i.e., the previous $j-1$ blocks).

A closer inspection of (19) reveals that the subspace is in fact $(s-1, p)$-*ultra periodic*, and are defined by a system of equations with the periodic structure of (2) of Definition 3.

Finally, we record the bound on the number of different possible solution spaces (this will be useful when we prune these via h.s.e sets later). By the choice of κ in (17), the total number of possible $(A_0, A_1, \dots, A_s)$ and hence the number of possible functional equations (17), is at most $q^{N(m-s+1)+s+1} \leqslant q^{Nm+s+1}$. Therefore, the number of possible candidate solution spaces is also at most q^{Nm+s+1}. □

Combining Lemmas 6.2 and 6.3 together with some simple calculations leads to the following statement concerning list decoding folded algebraic-geometric codes. We will later instantiate this with Hermitian and Garcia-Stichtenoth towers, and also combine with appropriate hierarchical subspace evasive sets to prune the periodic subspace of solutions into a small list size.

Theorem 6.5. *Consider the folded algebraic-geometric code from Definition 5 based on a function field $F/\mathbb{F}_q$ and automorphism σ. Let P (possibly equal to P_∞) be a rational place for which $x^\sigma = x/\xi$ for some local parameter $x \in F$ at P and ξ of order $p \geqslant m$ in $\mathbb{F}_q^*$. Assume that local expansions of functions in $\mathcal{L}((k+2\mathfrak{g}-1)P_\infty)$ at P can be computed in polynomial time.*

Then one can find a representation (1) of an (s,p)-periodic subspace of $\mathbb{F}_q^k$ containing all candidate messages $(f_0, f_1, \dots, f_{k-1})$ in polynomial time, when the fraction of errors $\tau = 1 - t/N$ satisfies

$$\tau \leqslant \frac{s}{s+1} - \frac{s}{s+1}\frac{k}{N(m-s+1)} - \frac{3m}{m-s+1}\frac{\mathfrak{g}}{mN} \ . \tag{21}$$

7. List decoding algebraic-geometric codes with subfield evaluation points

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 strategy in this section is similar to that of folded algebraic-geometric codes. For a folded algebraic-geometric code, once a coordinate is received correctly, then we have correct information on $f(P_i), f(P_i^\sigma) = f^{\sigma^{-1}}(P_i), \dots, f(P_i^{\sigma^{m-1}}) = f^{\sigma^{-m+1}}(P_i)$. For algebraic-geometric codes in this section, we has a similar property. Namely, once we receive a coordinate correctly, then we have correct information on $f(P_i), f^\sigma(P_i), \dots, f^{\sigma^{m-1}}(P_i)$, where σ is the Frobenius automorphism of an extension field.

For simplicity, to illustrate the ideas in a self-contained way in the setting of univariate polynomials, we begin with the case of Reed-Solomon codes in Section 7.1 . We then extend it to a general framework for decoding (one-point) algebraic-geometric codes based on constant field extensions in Section 7.2. Later on in the paper, we will instantiate the general framework to codes based on the Garcia-Stichtenoth tower discussed in 4.4.

7.1. **Decoding Reed-Solomon codes.**

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

Definition 6. *[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 $\alpha_1, \alpha_2, \dots, \alpha_n$ 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

$$D = \left\lfloor \frac{n-k+1}{s+1} \right\rfloor . \tag{22}$$

Definition 7 (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*

$$Q(X, Y_1, Y_2, \ldots, Y_s) = A_0(X) + A_1(X)Y_1 + A_2(X)Y_2 + \cdots + A_s(X)Y_s \ , \tag{23}$$

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 (24). We include the easy proof for completeness.

Lemma 7.1. *There exists a nonzero polynomial $Q \in \mathcal{P}$ such that*

$$Q(\alpha_i, y_i, y_i^q, y_i^{q^2}, \cdots, y_i^{q^{s-1}}) = 0 \quad \textit{for} \quad i = 1, 2, \ldots, n \ . \tag{24}$$

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 (22). 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 7.3 below shows that any polynomial Q given by Lemma 7.1 yields an algebraic condition that the message functions f we are interested in list decoding must satisfy.

Definition 8 (Frobenius action on polynomials)**.** *For a polynomial $f \in \mathbb{F}_{q^m}[X]$ with $f(X) = f_0 + f_1 X + \cdots + f_{k-1}X^{k-1}$, define the polynomial $f^\sigma \in \mathbb{F}_{q^m}[X]$ as $f^\sigma(X) = f_0^q + f_1^q X + \cdots + f_{k-1}^q X^{k-1}$.*

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

Lemma 7.3. *Suppose $Q \in \mathcal{P}$ satisfies the interpolation conditions* (24). *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 7.2, 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 $A_0 + A_1 f + A_2 f^{\sigma} + \cdots + A_s f^{\sigma^{s-1}} = 0$. We now prove that these solutions form a well-structured affine space over $\mathbb{F}_q$.

Lemma 7.4. *For integers $1 \leqslant s \leqslant m$, the set of solutions $f = \sum_{i=0}^{k-1} f_i X^i \in \mathbb{F}_{q^m}[X]$ to the equation*

$$A_0(X) + A_1(X)f(X) + A_2(X)f^{\sigma}(X) + \cdots + A_s(X)f^{\sigma^{s-1}}(X) = 0 \tag{25}$$

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 representation of this periodic subspace (in the form (1) *from Definition 2) can be computed in* $\mathrm{poly}(k, m, \log q)$ *time.*

Proof. If f, g are two solutions to (25), then so is $\alpha f + \beta g$ for any $\alpha, \beta \in \mathbb{F}_q$ with $\alpha + \beta = 1$. So the solutions to (25) 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 $A_i(X) = a_{i,0} + a_{i,1}X + a_{i,2}X^2 + \cdots$ for $i = 0, 1, 2, \ldots, s$. For $l = 0, 1, 2, \ldots, D$, define the linearized polynomial

$$B_l(X) = a_{1,l}X + a_{2,l}X^q + a_{3,l}X^{q^2} + \cdots + a_{s,l}X^{q^{s-1}} \ . \tag{26}$$

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 $\mathbb{F}_q$-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 (25) and equating the coefficient of X^i to be 0, we get

$$a_{0,i} + B_i(f_0) + B_{i-1}(f_1) + \cdots + B_1(f_{i-1}) + B_0(f_i) = 0 \ . \tag{27}$$

Therefore, for each $i = 0, 1, \ldots, k-1$, f_i must belong to a coset of the subspace $W + \theta_i$ where θ_i is an affine combination of $f_0, f_1, \ldots, f_{i-1}$. It follows that the solutions $(f_0, f_1, \ldots, f_{k-1})$ to 25 viewed as a vector in $\mathbb{F}_q^{mk}$ (w.r.t any fixed $\mathbb{F}_q$-basis of $\mathbb{F}_{q^m}$) belongs to an form an

$(s-1, m, k)$-periodic subspace. The equations (10) give the desired representation of this periodic subspace. □

Combining Lemmas 7.3 and 7.4, 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 7.5. *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 7.4 grows linearly in k, we cannot afford to list this subspace as the decoder's output for polynomial time decoding. However, using the periodic structure of the subspace, one can prune it by using a "pre-code" that only allows polynomials with coefficients in subspace designs or h.s.e sets as we will see in later sections.

7.2. Decoding algebraic-geometric codes.

We now generalize the Reed-Solomon algorithm from the previous subsection 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. Later, we 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 refer to Subsection 5.2 for encoding and will focus on a decoding algorithm.

7.2.1. *AG codes with evaluation points in a subfield.* Let $F/\mathbb{F}_q$ be a function field of genus $\mathfrak{g}$. 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.

Consider the Goppa geometric code defined by

$$C(m; l) := \{(f(P_1), f(P_2), \ldots, f(P_N)) : \ f \in \mathcal{L}_m(lP_\infty)\} \ . \tag{28}$$

We have the following well-known result on the parameters of the above algebraic-geometric codes.

Lemma 7.6. *The above code $C(m;l)$ is an $\mathbb{F}_{q^m}$-linear code over $\mathbb{F}_{q^m}$, rate at least $\frac{l-\mathfrak{g}+1}{N}$, and minimum distance at least $N-l$.*

7.2.2. *Encoding of code using local expansions.* As with the case of folded AG codes, the decoding algorithm will recover the message function f via the coefficients of its local expansion at some place P. Therefore, we will identify the message symbols with local expansion coefficiently of the function f and encode into a subcode of $C(m;l)$.

Definition 9 (Subfield algebraic-geometric code using local expansion)**.** *The folded algebraic-geometric code $\widetilde{C}(m;k)$ maps*

$$\mathbf{f} = (f_0, f_1, \ldots, f_{k-1}) \in \mathbb{F}_{q^m}^k \quad \mapsto \quad (\phi_P(\mathbf{f})(P_1), \phi_P(\mathbf{f})(P_2), \ldots, \phi_P(\mathbf{f})(P_N)) \in \mathbb{F}_{q^m}^N \ ,$$

where $\phi_P(\cdot)$ is the map converting a local expansion into an associated function guaranteed by Claim 5.2.

7.2.3. *A list decoding algorithm.* We now present a list decoding algorithm for the above codes. The algorithm follows the linear-algebraic list decoding algorithm for RS codes. It is quite similar to that of folded algebraic-geometric codes. Suppose a codeword encoding $f \in \mathcal{L}_m((k+2\mathfrak{g}-1)P_\infty)$ is transmitted and received as $\mathbf{y} = (y_1, y_2, \dots, y_N)$.

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

$$Q(Y_1, Y_2, \dots, Y_s) = A_0 + A_1Y_1 + A_2Y_2 + \cdots + A_sY_s \tag{29}$$

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

$$D = \left\lfloor \frac{N-k+(s-1)\mathfrak{g}+1}{s+1} \right\rfloor. \tag{30}$$

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

$$s(D-\mathfrak{g}+1)+D+k+\mathfrak{g} = (s+1)(D+1)-(s-1)\mathfrak{g}-1+k > N. \tag{31}$$

for the above choice (30) of D. The interpolation requirements on $Q \in F_m[Y_1, \dots, Y_s]$ are the following:

$$Q(y_i, y_i^\sigma, \dots, y_i^{\sigma^{s-1}}) = A_0(P_i) + A_1(P_i)y_i + A_2(P_i)y_i^\sigma + \cdots + A_s(P_i)y_i^{\sigma^{s-1}} = 0 \tag{32}$$

for $i = 1, 2, \dots, N$. 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 $Q \in F_m[Y_1, \dots, Y_s]$ of the form (29) satisfying the interpolation conditions (32) can be found by solving a homogeneous linear system over $\mathbb{F}_{q^m}$ with at most N constraints and at least $s(D-\mathfrak{g}+1)+D+k+\mathfrak{g}$ variables.

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

Lemma 7.7. *If f is a function in $\mathcal{L}_m((k+2\mathfrak{g}-1)P_\infty)$ whose encoding agrees with the received word $\mathbf{y}$ in at least t positions with $t > D+k+2\mathfrak{g}-1$, then*

$$Q(f, f^\sigma, \dots, f^{\sigma^{s-1}}) = A_0 + A_1f + A_2f^\sigma + \cdots + A_sf^{\sigma^{s-1}} = 0. \tag{33}$$

Proof. The proof proceeds by comparing the number of zeros of the function $Q(f, f^\sigma, \dots, f^{\sigma^{s-1}}) = A_0 + A_1f + A_2f^\sigma + \cdots + A_sf^{\sigma^{s-1}}$ with $D+k+2\mathfrak{g}-1$. Note that $Q(f, f^\sigma, \dots, f^{\sigma^{s-1}})$ is a function in $\mathcal{L}_m((D+k+2\mathfrak{g}-1)P_\infty)$. If position i of the encoding of f agrees with $\mathbf{y}$, then

$$\begin{aligned}
0 &= A_0(P_i) + A_1(P_i)y_i + A_2(P_i)y_i^\sigma + \cdots + A_s(P_i)y_i^{\sigma^{s-1}} \\
&= A_0(P_i) + A_1(P_i)f(P_i) + A_2(P_i)(f(P_i))^\sigma + \cdots + A_s(P_i)(f(P_i))^{\sigma^{s-1}} \\
&= A_0(P_i) + A_1(P_i)f(P_i) + A_2(P_i)f^\sigma(P_i) + \cdots + A_s(P_i)f^{\sigma^{s-1}}(P_i) \\
&= (A_0 + A_1f + A_2f^\sigma + \cdots + A_sf^{\sigma^{s-1}})(P_i)
\end{aligned}$$

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

Lemma 7.8. *Let P be a rational place of F with a local parameter $x \in F$ (P may or may not be the same as P_∞). The set of solutions $f \in \mathcal{L}_m((k+2\mathfrak{g}-1)P_\infty)$ to 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 has size at most $q^{(s-1)(k+2\mathfrak{g}-1)}$. Further, the possible first k coefficients $(f_0, f_1, \ldots, f_{k-1})$ of f's local expansion at P belong to an $(s-1,m)$-ultra periodic affine subspace of $\mathbb{F}_q^{mk}$.

Proof. The argument is very similar to Lemma 6.4. Define $\nu = k+2\mathfrak{g}-1$ if $P = P_\infty$ and 0 otherwise. Let $u = \min\{\nu_P(A_i) : i = 1,2,\ldots,s\}$. Then we have $\nu_P(A_0) = \nu_P(-\sum_{i=1}^{s} A_i f^{-\sigma^{i-1}}) \geqslant \min\{\nu_P(A_i f^{-\sigma^{i-1}})) : i = 1,2,\ldots,s\} \geqslant \min\{\nu_P(A_i) - \nu : i = 1,2,\ldots,s\} = u - \nu$. Each A_i has a local expansion at P:

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

for $i = 1,\ldots,s$ and A_0 has a local expansion $A_0 = x^{u-\nu} \sum_{j=0}^{\infty} a_{0,j} x^j$.

Assume that at P, the function f has a local expansion (9). Then f^{σ^i} has a local expansion at P as follows

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

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. For $d \geqslant 0$, equating the coefficient of $x^{d+u-\nu}$ in $A_0 + A_1 f + A_2 f^{\sigma} + \cdots + A_s f^{\sigma^{s-1}}$ to equal 0 gives us the condition

$$a_{0,d} + B_d(f_0) + B_{d-1}(f_1) + \cdots + B_0(f_d) = 0 \ . \tag{34}$$

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 7.4, 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 $(f_0, f_1, \ldots, f_{k+2\mathfrak{g}-1})$ belong to an $(s-1, m, k+2\mathfrak{g}-1)$-periodic subspace of $\mathbb{F}_q^{m(k+2\mathfrak{g}-1)}$. In particular, there are at most $q^{(s-1)(k+2\mathfrak{g}-1)}$ solutions $f \in \mathcal{L}_m((k+2\mathfrak{g}-1)P_\infty)$ to (32).

The equation (34) also shows that each group of ι successive coefficients $f_{d-\iota+1}$, $f_{d-\iota+2}, \cdots, f_d$ belong to cosets of the same underlying $\iota(s-1)$ dimensional subspace of $\mathbb{F}_q^{m\iota}$. This implies that $(f_0, f_1, \ldots, f_k)$ in fact belong to an $(s-1,m)$-ultra periodic subspace.[5] □

Decoding. Recall that the first k coefficients of the local expansion of $f \in \mathcal{L}_m(lP_\infty)$ around P is precisely the message that was encoded in the code $\widetilde{C}(m;k)$ of Definition 9.

Therefore, combining Lemmas 7.7 and 7.8, and recalling the choice of D in (30), we can conclude the following result about list-decodability of our code construction.

[5] This ultra-periodicity was also true for the Reed-Solomon case in Lemma 7.4, 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.

Theorem 7.9. *For the code $\widetilde{C}(m;k)$, 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*

$$\frac{s}{s+1}(N-k)-\frac{3s+1}{s+1}\mathfrak{g}$$

positions.

8. Instantiating with Hermitian and Garcia-Sticthenoth towers

In Sections 6 and 7, we discussed list decoding of folded algebraic-geometric codes and algebraic-geometric codes with subfield evaluation points. In this section, we instantiate the codes and list decoding algorithms described in Sections 6 and 7 with two important and explicit towers, i.e., the Hermitian and Garcia-Sticthenoth towers.

8.1. **Folded Hermitian codes.**

In this subsection, let us instantiate the list decoding algorithm of folded algebraic geometric codes from general algebraic function fields with the Hermtian tower. We refer to Subsection 4.3 for detailed background on the Hermitian tower. Let r be a prime power and let $q=r^2$. We denote by $\mathbb{F}_q$ the finite field with q elements. Let $F_e=\mathbb{F}_q(x_1,x_2,\dots,x_e)$ be the Hermitian tower defined by (4). Let γ be a primitive element of $\mathbb{F}_q$. Consider the automorphism $\sigma\in\mathrm{Aut}(F_e/\mathbb{F}_q)$ defined by

$$\sigma:\ x_i\mapsto\gamma^{(r+1)^{i-1}}x_i\quad\text{for } i=1,2,\dots,e.$$

For an integer m with $1\leqslant m\leqslant q-1$, let P_∞ and $P_i^{\sigma_j}$ for $i=1,2,\dots,N$ and $j=0,1,\dots m-1$ be the same as defined in Subsection 4.3.

Definition 10 (Folded codes from the Hermitian tower)**.** *Assume that m,l,N are positive integers satisfying $1\leqslant m\leqslant q-1$ and $l/m\leqslant N\leqslant r^{e-1}\left\lfloor\frac{q-1}{m}\right\rfloor$. The folded code from F_e with parameters N,l,q,e,m, denoted by $\mathsf{FH}_e(N,l,q,m)$, encodes a message function $f\in\mathcal{L}(lP_\infty)$ a folded codeword given in* (11).

When $e=1$, the folded code $\mathsf{FH}_1(N,l,q,m)$ is in fact a folded Reed-Solomon code introduced in [13].

Lemma 8.1. *The above code $\mathsf{FH}_e(N,l,q,m)$ is an $\mathbb{F}_q$-linear code over alphabet size q^m, rate at least $\frac{l-g_e+1}{Nm}$, and minimum distance at least $N-\frac{l}{m}$.*

Proof. It is clear that the map (11) is an $\mathbb{F}_q$-linear map. The dimension over $\mathbb{F}_q$ of the message space $\mathcal{L}(lP_\infty)$ is at least $l-g_e+1$ by the Riemann-Roch theorem, which gives the claimed lower bound on rate. For the distance property, observe that if the i-th column is zero, then f has m zeros. This implies that the encoding of a nonzero function f can have at most l/m zero columns since $f\in\mathcal{L}(lP_\infty)$. □

Let P_0 be the common zero of $x_1,x_2,\dots,x_e$. For our decoding, we will actually recover the message $f\in\mathcal{L}(lP_\infty)$ in terms of the coefficients of its power series expansion around P_0

$$f=f_0+f_1x+f_2x^2+\cdots\tag{35}$$

where $x:=x_1$ is the local parameter at P_0 (which means that x_1 has exactly one zero at P_0, i.e., $\nu_{P_0}(x_1)=1$).

With this in mind, we now define the encoding into the above folded Hermitian code using the map ϕ_{P_0} from Claim 5.2.

Definition 11 (Folded Hermitian code using local expansion)**.** *The folded Hermitian code $\widetilde{\mathsf{FH}}_e(N,k,q,m)$ maps*

$$\mathbf{f}=(f_0,f_1,\dots,f_{k-1})\in\mathbb{F}_q^k \quad \textit{to} \quad \mathsf{FH}_e(N,k+2g_e-1,q,m)(\phi_{P_0}(\mathbf{f}))\in(\mathbb{F}_q^m)^N \ .$$

Given the local expansions of a basis of $\mathcal{L}(lP_\infty)$ at P_0, computing the map ϕ_{P_0} to convert from local expansion to some representative function in $\mathcal{L}(lP_\infty)$ can be done in polynomial time by simply solving a system of linear equations. We now turn to the task of computing the local expansion at P_0 of a basis for $\mathcal{L}(lP_\infty)$.

Lemma 8.2. *For any n, one can compute the first n terms of the local expansion of the basis elements* (35) *at P_0 using* poly(n) *operations over $\mathbb{F}_q$.*

Proof. By the structure of the basis functions in (5), it is sufficient to find an algorithm of efficiently finding local expansions of x_i at P_0 for every $i=1,2,\dots,e$. We can inductively find the local expansions of x_i at P_0 as follows.

For $i=1$, x_1 is the local parameter x of P_0, so x is the local expansion of x_1 at P_0.

Now assume that we know the local expansion of $x_i=\sum_{j=1}^{\infty}c_{i,j}x^j$ at P_0 for some $c_{i,j}\in\mathbb{F}_q$. Then we have

$$\sum_{j=1}^{\infty}c_{i+1,j}^r x^{jr}+\sum_{j=1}^{\infty}c_{i+1,j}x^j=x_{i+1}^r+x_{i+1}=x_i^{r+1}=\left(\sum_{j=1}^{\infty}c_{i,j}^r x^{jr}\right)\left(\sum_{j=1}^{\infty}c_{i,j}x^j\right).$$

Note that r is a power of the characteristic and hence r can be pushed into infinite sums. By comparing the coefficients of x^j in the above identity, we can easily solve $c_{i+1,j}$'s from $c_{i,j}$'s. More specifically, the coefficient of x^j at the left of the identity is

$$\begin{cases} c_{i+1,j} & \text{if } r \not| j \\ c_{i+1,j}+c_{i+1,j/r}^r & \text{if } r|j. \end{cases}$$

Thus, all $c_{i+1,j}$'s can be easily solved recursively. □

By instantiating Theorem 6.5 with our code $\widetilde{\mathsf{FH}}_e(N,k,q,m)$, we obtain the following result.

Theorem 8.3. *One can find a representation of an $(s,q-1)$-periodic subspace*[6] *of $\mathbb{F}_q^k$ containing all candidate messages $\mathbf{f}=(f_0,f_1,\dots,f_{k-1})$ in polynomial time, when the fraction of errors $\tau=1-t/N$ in its encoding by $\widetilde{\mathsf{FH}}_e(N,k,q,m)$ satisfies*

$$\tau\leqslant\frac{s}{s+1}-\frac{s}{s+1}\frac{k}{N(m-s+1)}-\frac{3m}{m-s+1}\frac{\mathfrak{g}_e}{mN}\ . \tag{36}$$

8.2. **Folded codes from the Garcia-Stichtenoth tower.** Compared with the Hermitian tower of function fields, the Garcia-Stichtenoth tower of function fields yields folded codes with better parameters due to the fact that the Garcia-Stichtenoth tower is an optimal one in the sense that the ratio of number of rational places against genus achieves the maximal possible value. The construction of folded codes from the Garcia-Stichtenoth tower is almost identical to the one from the Hermitian tower except for one major difference: the redefined code from the Garcia-Stichtenoth tower is constructed in terms of the local expansion at point P_∞, while in the Hermitian case local expansion at P_0 is considered. For convenience of the reader, we give a parallel description of folded codes

[6] In fact, this subspace will be $(s,q-1)$-ultra periodic.

from the Garcia-Stichtenoth tower, while only sketching the identical parts. We refer to Subsection 4.4 for background on the Garcia-Stichtenoth tower.

Let r be a prime power and let $q = r^2$. We denote by $\mathbb{F}_q$ the finite field with q elements. Let $K_e = \mathbb{F}_q(x_1, x_2, \dots, x_e)$ be the Garcia-Stichtenoth tower defined by (7). Let γ be a primitive element of $\mathbb{F}_q$ and consider the automorphism $\sigma \in \mathrm{Aut}(K_e/\mathbb{F}_q)$ defined by

$$\sigma : \; x_i \mapsto \gamma^{r+1} x_i \quad \text{for } i = 1, 2, \dots, e.$$

For an integer m with $1 \leqslant m \leqslant q-1$, let P_∞ and $P_i^{\sigma_j}$ for $i = 1, 2, \dots, N$ and $j = 0, 1, \dots m-1$ be the same as defined in Subsection 4.4.

The folded codes from the Garcia-Stichtenoth tower are defined similarly to the Hermitian case.

Definition 12 (Folded codes from the Garcia-Stichtenoth tower)**.** *Assume that m, k, N are positive integers satisfying $1 \leqslant m \leqslant r-1$ and $l/m < N \leqslant r^e \left\lfloor \frac{r-1}{m} \right\rfloor$. The folded code from K_e with parameters N, l, q, e, m, denoted by $\mathsf{FGS}_e(N, l, q, m)$, encodes a message function $f \in \mathcal{L}(lP_\infty)$ a folded codeword given in* (11)*.*

Then we have a similar result on parameters of $\mathsf{FGS}_e(N, l, q, m)$.

Lemma 8.4. *The above code $\mathsf{FGS}_e(N, l, q, m)$ is an $\mathbb{F}_q$-linear code over alphabet size q^m, rate at least $\frac{l - \mathfrak{g}_e + 1}{Nm}$, and minimum distance at least $N - \frac{l}{m}$.*

Similar to the the Hermitian case, we need to redefine the code in terms of local expansion at a point. In the Hermitian case, we use coefficients of its power series expansion around P_0 which has a simple local parameter x_1. However, for the Garcia-Stichtenoth tower we do not have such a nice point P_0. Fortunately, we can use point P_∞ to achieve our mission, i.e., P_∞ has a simple local parameter $\frac{1}{x_e}$.

Definition 13 (Folded Garcia-Stichtenoth code using local expansion)**.** *The folded Garcia-Stichtenoth code (FGS code for short) $\widetilde{\mathsf{FGS}}_e(N, k, q, m)$ maps*

$$\mathbf{f} = (f_0, f_1, \dots, f_{k-1}) \in \mathbb{F}_q^k \quad \textit{to} \quad \mathsf{FGS}_e(N, k + 2\mathfrak{g}_e - 1, q, m)(\phi_{P_\infty}(\mathbf{f})) \in (\mathbb{F}_q^m)^N \; .$$

The rate of the above code equals $k/(Nm)$ and its distance is at least $N - (k + 2\mathfrak{g}_e - 1)/m$.

As in the Hermitian case, we now turn to the task of computing the local expansion around P_∞ of a basis for $\mathcal{L}(lP_\infty)$, which then suffices to compute the map ϕ_{P_∞} efficiently. The local expansion of $f \in \mathcal{L}(lP_\infty)$ around P_∞ is of the form

$$f = T^{-l}(f_0 + f_1 T + f_2 T^2 + \cdots) \tag{37}$$

where $T := \frac{1}{x_e}$ is the local parameter at P_∞ (the function x_e has exactly one pole at P_∞).

Lemma 8.5. *For any n, one can compute the first n terms of the local expansion* (37) *of a basis of $\mathcal{L}(lP_\infty)$ at P_∞ using* $\mathrm{poly}(n)$ *operations over $\mathbb{F}_q$.*

Proof. First let h be a nonzero function in $\mathbb{F}_q(x_1, x_2, \dots, x_e)$ with $\nu_{P_\infty}(h) = v \in \mathbb{Z}$. Assume that the local expansion $h = T^v \sum_{j=0}^{\infty} a_j T^j$ is known. To find the local expansion $\frac{1}{h} = T^{-v} \sum_{j=0}^{\infty} c_j T^j$. Consider the identity

$$1 = \left(\sum_{j=0}^{\infty} c_j T^j \right) \left(\sum_{j=0}^{\infty} a_j T^j \right) .$$

Then by comparing the coefficients of T^i in the above identity, one has $c_0 = a_0^{-1}$ and $c_i = -a_0^{-1}(c_{i-1}a_1 + \cdots + c_0 a_i)$ can be easily computed recursively for all $i \geqslant 1$.

Thus, by the structure of the basis functions in (8), it is sufficient to find an algorithm efficiently finding local expansions of x_i at P_∞ for every $i = 1, 2, \dots, e$. We can inductively find the local expansions of x_i at P_∞ as follows. We note that $\nu_{P_\infty}(x_i) = -r^{e-i}$ for $i = 1, 2, \dots, e$.

For $i = e$, x_e has the local expansion $\frac{1}{T}$ at P_∞.

Now assume that we know the local expansion of x_i. Then we can easily compute the local expansion of $x_i^r + x_i$ and hence the local expansion of $1/(x_i^r + x_i)$. Let us assume that $1/(x_i^r + x_i)$ has local expansion $1/(x_i^r + x_i) = T^{r^{e-i+1}} \sum_{j=0}^{\infty} \alpha_j T^j$ at P_∞ for some $\alpha_i \in \mathbb{F}_q$. Assume that $1/x_{i-1}$ has the local expansion $1/x_{i-1} = T^{r^{e-i+1}} \sum_{j=0}^{\infty} \beta_j T^j$. To find β_j, we consider the identity

$$T^{r^{e-i+1}} \sum_{j=0}^{\infty} \beta_j T^j + T^{r^{e-i+2}} \sum_{j=0}^{\infty} \beta_j^r T^{rj} = \frac{1}{x_{i-1}} + \left(\frac{1}{x_{i-1}}\right)^r = \frac{1}{x_i^r + x_i} = T^{r^{e-i+1}} \sum_{j=0}^{\infty} \alpha_j T^j.$$

By comparing the coefficients of $T^{j+r^{e-i+1}}$ in the above identity, we have that $\beta_0 = \alpha_0$ and β_j can be easily computed recursively by the following formula for all $i \geqslant 1$.

$$\beta_j = \begin{cases} \alpha_j & \text{if } r \nmid j \\ \alpha_j - \beta_{j/r-1}^r & \text{if } r | j. \end{cases}$$

Therefore, the local expansion of x_{i-1} at P_∞ can be easily computed. □

Similar to the Hermitian case, by instantiating Theorem 6.5 with our code $\widetilde{\mathsf{FGS}}_e(N, k, q, m)$, we obtain the following result.

Theorem 8.6. *One can find a representation of the $(s, r-1)$-ultra periodic subspace containing all candidate messages $(f_0, f_1, \dots, f_{k-1})$ in polynomial time, when the fraction of errors $\tau = 1 - t/N$ in its encoding by $\widetilde{\mathsf{FGS}}_e(N, k, q, m)$ satisfies*

$$\tau \leqslant \frac{s}{s+1}\left(1 - \frac{k}{N(m-s+1)}\right) - \frac{3m}{m-s+1}\frac{r^e}{mN}\ . \tag{38}$$

8.3. **Subfield evaluation codes from the Garcia-Stichtenoth tower.** Let r be a prime power and let $q = r^2$. For $e \geqslant 2$, let K_e be the function field $\mathbb{F}_q(x_1, x_2, \dots, x_e)$ given by Garcia-Stichtenoth tower (7), with genus $\mathfrak{g}_e \leqslant r^e$.

Put $F = K_e$ and $F_m = \mathbb{F}_{q^m} \cdot K_e$. Let $P_1, P_2, \dots, P_N$ be the rational points of F besides the place P_∞; we have $N \geqslant r^e(r-1)$. Let k be the desired dimension of the code (where $1 \leqslant k < N - 2\mathfrak{g}_e$) and let $l = k + 2\mathfrak{g}_e - 1$. We will now instantiate the code $C(m; l)$ defined in (28) with the Garcia-Stichtenoth function field F_m and $P_1, P_2, \dots, P_N$ as evaluation points. (We hide the dependence on e and N in the specification of the code $C(m; l)$ as implicit.)

Let us call the resulting code $C_{\mathrm{GS}}(m; l)$. As in the previous sections, we will encode into subcode $\widetilde{C}_{\mathrm{GS}}(m; k)$ using as message vector the first k coefficients of the local expansion around P_∞. The Garcia-Stichtenoth code with subfield evaluation using local expansion, $\widetilde{C}_{\mathrm{GS}}(m; k)$ is defined from $C_{\mathrm{GS}}(m; l)$ as in Definition 9.

BY virtue of Theorem 7.9, we can now conclude the following:

Corollary 8.7. *The code $C_{\mathrm{GS}}(m; k + 2g_e - 1)$ 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 (see Section 10). 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 8.8. *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 infinitely 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 $\mathbb{F}_q^{mk}$.

9. Hierarchical subspace-evasive sets

Let us first recall the notion of "ordinary" subspace-evasive sets from [10].

Definition 14. *A subset $S \subset \mathbb{F}_q^k$ is said to be (d,ℓ)-subspace-evasive if for all d-dimensional affine subspaces H of $\mathbb{F}_q^k$, we have $|S \cap H| \leqslant \ell$.*

We next define the notion of evasiveness w.r.t a collection of subspaces instead of all subspaces of a particular dimension.

Definition 15. *Let $\mathcal{F}$ be a family of (affine) subspaces of $\mathbb{F}_q^k$, each of dimension at most d. A subset $S \subset \mathbb{F}_q^k$ is said to be $(\mathcal{F}, d, \ell)$-evasive if for all $H \in \mathcal{F}$, we have $|S \cap H| \leqslant \ell$.*

The key to pruning the list to a small size is the notion of a *hierarchical subspace-evasive set*, which is defined as a subset of $\mathbb{F}_q^k$ with the property that some of its prefixes are subspace-evasive with respect to (s,Δ,b)-periodic subspaces. We will show how the special subspace-evasive sets help towards pruning the list in our list decoding context in Section 9.3.

Definition 16. *Let $\mathcal{F}$ be a family of (s,Δ,b)-periodic subspaces of $\mathbb{F}_q^k$ with $k = b\Delta$. A subset $S \subset \mathbb{F}_q^k$ is said to be $(\mathcal{F}, s, \Delta, b, L)$-h.s.e (for hierarchically subspace evasive for block size Δ) if for every affine subspace $H \in \mathcal{F}$, the following bound holds for $j = 1, 2, \ldots, b$:*

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

Remark 1. For h.s.e based pruning, a property weaker than (s,Δ)-periodicity of H suffices. Namely, it is enough if for each prefix $a \in \mathbb{F}_q^{j\Delta}$, the extensions of a in H form an affine space of dimension s (it is not necessary that this be a coset of the *same* subspace of $\mathbb{F}_q^{\Delta}$ for every j). However, we stick with the periodicity assumption since it is available to us in the subspaces output by the list decoder, and is also necessary for the subspace design based pruning of the next section.

9.1. Random sets are subspace evasive.

Our goal is to give a randomized construction of large h.s.e sets that works with high probability, with the further properties that one can index into elements of this set efficiently (necessary for efficient encoding), and one can check membership in the set efficiently (which is important for efficient decoding).

An easy probabilistic argument, see [10], shows that a random subset of $\mathbb{F}_q^k$ of size about $q^{(1-\zeta)k}$ is $(d, O(d/\zeta))$-subspace evasive with high probability. As a warmup, let us work out the similar proof for the case when we have only to avoid a not too large family $\mathcal{F}$ of all possible d-dimensional affine subspaces. The advantage is that the guarantee on the intersection size is now $O(1/\zeta)$ and independent of the dimension d of the subspaces one is trying to evade.

Lemma 9.1. *Let $\zeta \in (0,1)$ and k be a large enough positive integer. Let $\mathcal{F}$ be a family of affine subspaces of $\mathbb{F}_q^k$, each of dimension at most $d \leqslant \zeta k/2$, with $|\mathcal{F}| \leqslant q^{ck}$ for some positive constant c.*

Let S be a random subset of $\mathbb{F}_q^k$ chosen by including each $x \in \mathbb{F}_q^k$ in S with probability $q^{-\zeta k}$. Then with probability at least $1 - q^{-ck}$, S satisfies both the following conditions: (i) $|S| \geqslant q^{(1-2\zeta)k}$, and (ii) S is $(\mathcal{F}, d, 4c/\zeta)$-evasive.

Proof. The first part follows by noting that the expected size of S equals $q^{(1-\zeta)k}$ and a standard Chernoff bound calculation. For the second part, fix an affine subspace $H \subseteq \mathcal{F}$ of dimension at most d, and a subset $T \subseteq H$ of size t, for some parameter t to be specified shortly. The probability that $S \supseteq T$ equals $q^{-\zeta k t}$. By a union bound over the at most q^{ck} choices for the affine subspace $H \in \mathcal{F}$, and the at most q^{dt} choices of t-element subsets T of H, we get that the probability that S is not $(\mathcal{F}, d, t)$-evasive is at most $q^{ck+dt} \cdot q^{-\zeta k t} \leqslant q^{ck} q^{-\zeta k t/2}$ since $d \leqslant \zeta k/2$. Choosing $t = \lceil 4c/\zeta \rceil$, this quantity is bounded from above by q^{-ck}. □

9.2. **Pseudorandom construction of large h.s.e subsets.** We next turn to the pseudorandom construction of large h.s.e subsets. Suppose, for some fixed subset $\mathcal{F}$ of (s, Δ, b)-periodic subspaces of $\mathbb{F}_q^k$ with $k = b\Delta$, we are interested in an $(\mathcal{F}, s, \Delta, b, L)$-h.s.e subset of $\mathbb{F}_q^k$ of size $\approx q^{(1-\zeta)k}$ for a constant ζ, $1/\Delta < \zeta < 1/3$. (Bwlow, we will ignore floors and ceilings in the description to avoid notational clutter; those are easy to accommodate and do not affect any of the claims.)

Denote $\Delta' = (1-\zeta)\Delta$, $b' = (1-\zeta)b$, and $k' = b'\Delta = (1-\zeta)k$.

The random part of the construction will consist of mutually independent, random univariate polynomials $P_1, P_2, \ldots, P_{b'}$ and Q, where $P_j \in \mathbb{F}_{q^{j\Delta'}}[T]$ for $1 \leqslant j \leqslant b'$ and $Q \in \mathbb{F}_{q^{k'}}[T]$ are random polynomials of degree λ.[7] The degree parameter will be chosen to be $\lambda = \Theta(k)$.[8]

The key fact we will use about random polynomials is the following, which follows by virtue of the λ-wise independence of the values of a random degree λ polynomial.

Fact 9.2. *Let $P \in \mathbb{K}[T]$ be a polynomial of degree λ whose coefficients are picked uniformly and independently at random from the field $\mathbb{K}$. For a fixed subset $T \subseteq \mathbb{K}$ with $|T| \leqslant \lambda$, the values $\{P(\alpha)\}_{\alpha \in T}$ are independent random values in $\mathbb{K}$.*

We remark that this property of low-degree polynomials was also the basis of the pseudorandom construction of subspace evasive sets in [16]. However, since we require the h.s.e property, and need to exploit the periodicity of the subspaces we are trying to evade

[7] We will assume that representations of the necessary extension fields $\mathbb{F}_q^{i\Delta'}$ are all available. For this purpose, we only need irreducible polynomials over $\mathbb{F}_q$ of appropriate degrees, which can be constructed by picking random polynomials and checking them for irreducibility. Our final construction is anyway randomized, so the randomized nature of this step does not affect the results.

[8] The degree of Q can in fact be just $O(1/\zeta)$, but for uniformity we fix the degree of all polynomials to be the same.

(which can have large dimension), the construction here is more complicated, and needs to use several polynomials P_j's evaluated in a nested fashion, and one further polynomial Q to further bring down the list size to a constant (this final use of Q is similar in spirit to the construction in [16]). We remark that the construction presented here is a bit simpler and cleaner than the one in the conference version [18], and comes with efficient encoding automatically by construction. In contrast, the construction in [18] required some additional work in order to allow for efficient encoding.

In what follows we assume that, for $j = 1, 2, \ldots, b'$, some fixed bases of the fields $\mathbb{F}_{q^{j\Delta'}}$ have been chosen, giving us some canonical $\mathbb{F}_q$-linear injective maps

$$\rho_j : \mathbb{F}_q^{j\Delta'} \to \mathbb{F}_{q^{j\Delta'}} \ .$$

Also, for $j = 1, 2, \ldots, b'$, let

$$\xi_j : \mathbb{F}_{q^{j\Delta'}} \to \mathbb{F}_q^{\zeta\Delta}$$

be some arbitrary $\mathbb{F}_q$-linear surjective map (thus ξ_j just outputs the first $\zeta\Delta$ coordinates of the representation of elements of $\mathbb{F}_{q^{j\Delta'}}$ as vectors in $\mathbb{F}_q^{j\Delta'}$ w.r.t some fixed basis). Finally, let $\rho : \mathbb{F}_q^{k'} \to \mathbb{F}_{q^{k'}}$ be some fixed $\mathbb{F}_q$-linear injective map, and $\xi : \mathbb{F}_{q^{k'}} \to \mathbb{F}_q^{\zeta k}$ be an arbitrary $\mathbb{F}_q$-linear surjective map.

We are now ready to describe our construction of h.s.e set based on the random polynomials $P_1, P_2, \ldots, P_{b'}, Q$.

Definition 17 (h.s.e set construction)**.** *Given the polynomials $P_j \in \mathbb{F}_{q^{j\Delta'}}[T]$ for $i = 1, 2, \ldots, b'$ and $Q \in \mathbb{F}_{q^{k'}}[T]$, define the subset $\Gamma(P_1, P_2, \ldots, P_b; Q)$ by*

$$\Big\{(y_1, z_1, y_2, z_2, \ldots, y_{b'}, z_{b'}; w) \in \mathbb{F}_q^k \;\Big|\; \textit{for } j = 1, 2, \ldots, b' : y_j \in \mathbb{F}_q^{\Delta'}, \\ z_j = \xi_j(P_j(\rho_j(y_1 \circ y_2 \circ \cdots \circ y_j))) \in \mathbb{F}_q^{\zeta\Delta}; \ \textit{and} \\ w = \xi(Q(\rho(y_1, z_1, \ldots, y_{b'}, z_{b'}))) \in \mathbb{F}_q^{\zeta k}\Big\} \ .$$

By construction, once suitable representations of the extension fields are available by pre-processing and the choice of $P_1, \ldots, P_{b'}, Q$ is made, we can efficiently compute a bijective encoding map $\mathsf{HSE} : \mathbb{F}_q^{(1-\zeta)^2 k} \to \Gamma(P_1, P_2, \ldots, P_b; Q)$. Indeed, we can view the input $\mathbf{y} \in \mathbb{F}_q^{b'\Delta'}$ as $(y_1, y_2, \ldots, y_{b'})$ with $y_j \in \mathbb{F}_q^{\Delta'}$ and then compute the z_j's and w efficiently using poly(k) operations over $\mathbb{F}_q$ (recall that the degree of the polynomials is $\lambda = \Theta(k)$).

We now move on to the main claim about the h.s.e property of our construction.

Theorem 9.3. *Let c be a positive constant. Let $\zeta \in (0, 1/3)$ and s be a positive integer satisfying $s < \zeta\Delta/10$. Let $\mathcal{F}$ be a subset of at most q^{ck} (s, Δ, b)-periodic subspaces of $\mathbb{F}_q^k$ for $k = b\Delta$ that is much bigger than $1/\zeta$. Suppose that the parameters satisfy the condition $q^{\zeta\Delta} \geqslant (2q^2ck)^{10/9}$. Then with probability $1 - q^{-\Omega(k)}$ over the choice of random polynomials $\{P_i\}_{1 \leqslant i \leqslant b}$ and Q each of degree $\lambda = \lceil ck \rceil$, the set $\Gamma(P_1, P_2, \ldots, P_b; Q)$ from Definition 17 is*

$$(\mathcal{F}, s, \Delta, b, L)\text{-}\textit{h.s.e} \text{ and } (\mathcal{F}, sb, \ell)\text{-}\textit{evasive}$$

for $L = \lceil 2ck \rceil$ and $\ell = \lceil 4c/\zeta \rceil$ (note that (i) $L \gg \ell$ as $k \gg 1/\zeta$; and (ii) Q trims down the intersection size from L to ℓ).

Proof. Note that the first $k' = (1-\zeta)k$ symbols of vectors in $\Gamma(P_1, \ldots, P_{b'}; Q)$ only depend on the P_j's. We will first prove that with high probability over the choice of the P_j's the following holds (call such a choice of P_j's as *good*):

For every $H \in \mathcal{F}$, $|\mathrm{proj}_{k'}(H) \cap \mathrm{proj}_{k'}(\Gamma)| < L$, where we denote Γ as shorthand for $\Gamma(P_1, \ldots, P_{b'}; Q)$.

Then, conditioned on a good choice of P_j's, we will prove that with high probability over the choice of the random polynomial Q, $|H \cap \Gamma| < \ell$. Together, these steps will imply that the set $\Gamma(P_1, P_2, \ldots, P_{b'}; Q)$ is $(\mathcal{F}, sb, \ell)$-evasive. (Note that every subspace in $\mathcal{F}$ has dimension at most sb by Claim 3.1.) We will return to the $(\mathcal{F}, s, \Delta, b, L)$-h.s.e property at the end of the proof.

Let us first establish the second step. Fix a good choice of $P_1, \ldots, P_{b'}$, and suppose we pick Q randomly. Fix a subspace $H \in \mathcal{F}$. Since $|\mathrm{proj}_{k'}(H) \cap \mathrm{proj}_{k'}(\Gamma)| < L$ (recall that $\mathrm{proj}_{k'}(\Gamma)$ only depends on the P_j's and thus is already determined), the number of elements of H that could possibly belong to Γ (after the choice of Q) is at most $L \cdot q^{s(b-b')} = Lq^{\zeta sb}$; indeed for each prefix belonging to $\mathrm{proj}_{k'}(\Gamma) \cap \mathrm{proj}_{k'}(H)$, there are most $q^{s(b-b')}$ extensions that can fall in H since H is (s, Δ, b)-periodic. Further, the probability over the choice of Q that any such fixed extension belongs to Γ is at most $q^{-\zeta k}$, and any ℓ of these events are independent. (Note that for a fixed prefix, there can be at most one extension that falls in Γ, so for ℓ different strings to fall in Γ, their prefixes must be distinct and are mapped to independent locations by the random polynomial Γ.) Therefore, the probability over the choice of Q that $|H \cap \Gamma| \geqslant \ell$ is at most $(Lq^{\zeta sb})^{\ell} q^{-\zeta k \ell}$. By a union bound over all $H \in \mathcal{F}$, we conclude that $|H \cap \Gamma| < \ell$ for every $H \in \mathcal{F}$ simultaneously, except with probability at most

$$q^{ck} L^{\ell} q^{\zeta(s-\Delta)b\ell} \leqslant q^{ck} (ck)^{\ell} q^{-\zeta \Delta b \ell/2} \leqslant q^{ck} q^{-\zeta k \ell/4}$$

where in the first inequality we used $s \leqslant \Delta/2$ and in the next one $ck \leqslant q^{\zeta k/4}$ both of which hold comfortably. For $\ell \geqslant 4c/\zeta$, the above probability upper bound is at most q^{-ck}.

We now turn to the first step, on the P_j's being good with high probability. Fix some $H \in \mathcal{F}$; we will prove by induction on j that

$$|\mathrm{proj}_{j\Delta}(H) \cap \mathrm{proj}_{j\Delta}(\Gamma)| < L \tag{39}$$

w.h.p over the choice of $P_1, P_2, \ldots, P_j$, for $1 \leqslant j \leqslant b'$ (note that $\mathrm{proj}_{j\Delta}(\Gamma)$ only depends on $P_1, \ldots, P_j$, so this event is well defined). For the base case $j = 1$, $|\mathrm{proj}_{\Delta}(H)| \leqslant q^s$ as H is (s, Δ, b)-periodic, and the probability that some L of these q^s elements belong to $\mathrm{proj}_{\Delta}(\Gamma)$ is at most q^{sL} times the probability that L distinct elements in $\mathbb{F}_{q^{\Delta'}}$ are mapped to specific values in $\mathbb{F}_q^{\zeta\Delta}$ by $\xi_1 \circ P_1$, which is at most $\left(q^{-\zeta\Delta}\right)^L$. So the overall probability that $|\mathrm{proj}_{\Delta}(H) \cap \mathrm{proj}_{\Delta}(\Gamma)| \geqslant L$ is at most $q^{(s-\zeta\Delta)L}$.

Now let $j \geqslant 2$ and assume $|\mathrm{proj}_{(j-1)\Delta}(H) \cap \mathrm{proj}_{(j-1)\Delta}(\Gamma)| < L$. By the (s, Δ, b)-periodicity of H, for each of the (less than L) prefixes in $\mathrm{proj}_{(j-1)\Delta}(H) \cap \mathrm{proj}_{(j-1)\Delta}(\Gamma)$, there are at most q^s extensions that fall in $\mathrm{proj}_{j\Delta}(H)$. Similarly to the argument used for second step above, the probability that some L of these belong to $\mathrm{proj}_{j\Delta}(\Gamma)$ is at most $(Lq^s)^L \cdot q^{-\zeta\Delta L}$. Thus, the probability that $|\mathrm{proj}_{j\Delta}(H) \cap \mathrm{proj}_{j\Delta}(\Gamma)| \geqslant L$ is at most $\left(L \cdot q^{(s-\zeta\Delta)}\right)^L$.

Combining these arguments, we conclude that the probability over the choice of the P_j's that $|\mathrm{proj}_{b'\Delta}(H) \cap \mathrm{proj}_{b'\Delta}(\Gamma)| \geqslant L$ is at most

$$b'(L \cdot q^{(s-\zeta\Delta)})^L \leqslant (2ckq^{-0.9\zeta\Delta})^L \leqslant q^{-2L}$$

where the last step used the assumption that $q^{\zeta\Delta} \geqslant (2q^2ck)^{10/9}$.

Finally, since there are at most q^{ck} subspaces $H \in \mathcal{F}$, by a union bound we have that for all $H \in \mathcal{F}$ simultaneously, $|\mathrm{proj}_{k'}(H) \cap \mathrm{proj}_{k'}(\Gamma)| < L$ with probability at least $1 - q^{ck}q^{-L} \geqslant 1 - q^{-ck}$ over the choice of $P_1, \ldots, P_{b'}$.

To finish the proof, we need to verify the $(\mathcal{F}, s, \Delta, b, L)$-h.s.e property. That is, we need to prove that w.h.p, $|\text{proj}_{j\Delta}(H) \cap \text{proj}_{j\Delta}(\Gamma)| \leqslant L$ for every $H \in \mathcal{F}$ and $j = 1, 2, \ldots, b$. By (39), this holds for $j = 1, 2, \ldots, b'$. By construction, the last ζk symbols of any vector in Γ is a function of the first $(1-\zeta)k = b'\Delta$ symbols, so $|\text{proj}_{j\Delta}(H) \cap \text{proj}_{j\Delta}(\Gamma)| \leqslant L$ also holds for $b' < j \leqslant b$. □

9.3. **Efficient computation of intersection with h.s.e. subsets.** The key aspect which makes h.s.e subsets useful in our context to prune the affine space of candidate messages, and indeed motivated the exact specifics of the definition and aspects of its construction, is the following claim which shows that intersection of a (s, Δ, b)-periodic subspace with our h.s.e set can be found efficiently.

Lemma 9.4. *(h.s.e.-intersection) There is an algorithm running in time* $\text{poly}(k, q^{\zeta\Delta})$ *that provides the following guarantee. Given as input the polynomials* $P_1, \ldots, P_{b'}$ *and* Q *underlying the construction of an* $(\mathcal{F}, s, \Delta, b, L)$*-h.s.e and* $(\mathcal{F}, sb, \ell)$*-evasive set* $\Gamma = \Gamma(P_1, \ldots, P_{b'}; Q)$ *and an* (s, Δ, b)*-periodic subspace* $H \subseteq \mathbb{F}_q^k$ *belonging to* $\mathcal{F}$*, the algorithm computes the at most* ℓ *elements of* $H \cap \Gamma$*.*

Proof. The proof essentially follows from the observations made in the proof of Theorem 9.3. First note that $|H \cap \Gamma(P_1, \ldots, P_{b'}; Q)| \leqslant \ell$ just follows from the $(\mathcal{F}, sb, \ell)$-evasiveness of Γ. To compute $H \cap \Gamma$, the algorithm iteratively computes the intersections $\text{proj}_{j\Delta}(H) \cap \text{proj}_{j\Delta}(\Gamma)$ for $1 \leqslant j \leqslant b'$. As Γ is $(\mathcal{F}, s, \Delta, b, L)$-h.s.e, this intersection has size at most L. To compute $\text{proj}_{j\Delta}(H) \cap \text{proj}_{j\Delta}(\Gamma)$, the algorithm runs over the at most q^s possible extensions of each element of $\text{proj}_{(j-1)\Delta}(H) \cap \text{proj}_{(j-1)\Delta}(\Gamma)$ that can belong to $\text{proj}_{j\Delta}(H)$ (due to the (s, Δ, b)-periodicity of H), and checks which ones also belong to $\text{proj}_{j\Delta}(\Gamma)$. The complexity amounts to $q^{O(s)}$ evaluations of degree $O(k)$ polynomials, and thus takes $q^{O(\zeta\Delta)}\text{poly}(k)$ time. To compute $H \cap \Gamma$ from $\text{proj}_{b'\Delta}(H) \cap \text{proj}_{b'\Delta}(\Gamma)$, we recall the earlier observation that the construction of Γ implies that there is a unique extension of an element in $\text{proj}_{b'\Delta}(\Gamma)$ that belongs to Γ. □

We conclude this section by recording in convenient form all necessary properties of our h.s.e set construction, which follow from Theorem 9.3 and Lemma 9.4. (We can remove the restriction that k is a multiple of Δ by constructing a subspace in $\mathbb{F}_q^{k^\#}$ for $k^\# = \Delta\lceil \frac{k}{\Delta} \rceil$ and dropping the last $k^\# - k$ coordinates, so we remove that restriction in the final statement below on h.s.e sets.)

Theorem 9.5. *Let* c *be a constant. Let* $\zeta \in (0, 1)$*, and* Δ, s, k *be positive integers satisfying* $s < \zeta\Delta/10$ *and* $k \leqslant q^{\zeta\Delta/2}$*. Let* $\mathcal{F}$ *be a family of at most* q^{ck} (s, Δ)*-periodic subspaces of* $\mathbb{F}_q^k$*. Then there is* $\text{poly}(k, \log q)$ *time randomized construction of an injective map* $\mathsf{HSE} : \mathbb{F}_q^{(1-\zeta)^2 k} \to \mathbb{F}_q^k$ *such that:*

(1) *Given* $\mathbf{x} \in \mathbb{F}_q^{(1-\zeta)^2 k}$*,* $\mathsf{HSE}(\mathbf{x})$ *can be computed using* $\text{poly}(k)$ *operations over* $\mathbb{F}_q$*.*
(2) *With probability at least* $1 - q^{-\Omega(k)}$ *over the construction of* HSE*, the following holds: for every* $H \in \mathcal{F}$*, the set* $\{\mathbf{x} \in \mathbb{F}_q^{(1-\zeta)^2 k} \mid \mathsf{HSE}(\mathbf{x}) \in H\}$ *has size at most* $O(c/\zeta)$*, and further can be computed in* $\text{poly}(k, q^{\zeta\Delta})$ *time.*

10. Subspace designs

The linear-algebraic list decoder discussed in the previous sections pins down the coefficients of the message to a periodic subspace. We already saw, in Section 9, an approach using h.s.e. sets to prune the periodic subspace to a small list. In this section, we will develop an alternate approach based a special collection of subspaces, which we call a *subspace design*, for pruning the periodic subspaces. Further, we will extend the construction to a "cascaded" variant that enables more effective pruning of ultra-periodic subspaces. The advantage of using subspace designs is they can be explicitly constructed, a feature which is (currently) lacking for h.s.e sets.

We begin with the definition of the central object of study in this section, subspace designs, introduced in the conference version [19].[9]

Definition 18 (Subspace design)**.** *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 H of subspaces of $\mathbb{F}_q^\Lambda$ such that for every r-dimensional subspace $W \subset \mathbb{F}_q^\Lambda$, we have*

$$\sum_{H \in \mathcal{H}} \dim(W \cap H) \leqslant d \ . \tag{40}$$

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 $\mathcal{H}$.*

Note that the condition (40) in particular implies for every r-dimensional subspace W, at most d of the subspaces in an (r,d)-subspace design non-trivially intersect it. This weaker property was subsequently called a "weak subspace design" in [11] which gave explicit constructions of subspace designs following our original definition in [19]. For our list decoding application, the stronger property (40) is required. Note though that the weak subspace design property implies the stronger (40) with the r.h.s upper bound d replaced by dr.

10.1. **Subspace designs to prune periodic subspaces.**

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

Lemma 10.1 (Periodic subspaces intersected with a subspace design)**.** *Suppose $H_1, H_2, \dots, H_b$ are subspaces in an (r,d)-subspace design in $\mathbb{F}_q^\Lambda$, and T is an (r,Λ,b)-periodic affine subspace of $\mathbb{F}_q^{\Lambda b}$ with recurring subspace $S \subseteq \mathbb{F}_q^\Lambda$. Then the set*

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

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 $\mathcal{S} \overset{\text{def}}{=} S^b \cap (H_1 \times H_2 \times \cdots \times H_b)$.

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^b 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 $\mathbb{F}_q^\Lambda$

[9] While we were not aware of it when we coined this term to refer to our subspace collections, subspace designs were used to denote the q-analogs of combinatorial designs [2]. We apologize for unknowingly repeating this nomenclature in our (very different) context.

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

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

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 $|\mathcal{T}|$. □

10.2. **Existence and probabilistic construction of 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 10.2. *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.*

Proof. Let $\ell = 8r/\eta$, and let $M = q^{\eta\Lambda/8}$ denote the number of randomly chosen subspaces.[10] 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 non-negative 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 gives a $q^{O(\Lambda r)}$ time *Las Vegas* construction of an (r,d)-subspace design with many subspaces each of large dimension $(1-\eta)m$.

Lemma 10.3. *For parameters η, r, Λ as in Lemma 10.2, 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 in $q^{O(\Lambda r)}$* Las Vegas *time.*

As noted in the conference version [19] of this paper, the construction can be derandomized using the method of conditional expectations to successively find good subspaces H_i to add to the subspace design. However, as each step involves searching over all

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

$(1-\eta)\Lambda$-dimensional subspaces of $\mathbb{F}_q^\Lambda$, the construction time would be $q^{O(\Lambda^2)}$ even for constructing subspace designs with few subspaces. For our application to reducing the list size for long algebraic-geometric codes (either folded or with rational points in a subfield), we will need subspace designs for ambient dimension Λ growing at least logarithmically in the code length. The $q^{O(\Lambda^2)}$ complexity will thus lead to a quasi-polynomial code construction time, as claimed in the conference version [19]. In fact, even the Las Vegas construction time of $q^{O(\Lambda r)}$ will be super-polynomial for the parameters used in the construction.

10.3. **Explicit subspace design constructions.** The question of explicit (polynomial time) constructions of subspace designs naturally arose following [19] and was addressed in the follow-up work by Guruswami and Kopparty [11], who proved the following.

Theorem 10.4 (Explicit subspace designs [11])**.** *For every $\eta > 0$, integers r, Λ with $r \leqslant \eta\Lambda/4$, and prime powers q satisfying $q^{\eta\Lambda/(2r)} > 2r/\eta$, for any $b \leqslant q^{\eta\Lambda/(4r)}$, there exists an explicit $(r, r^2/\eta)$-subspace design of cardinality b and dimension $(1-\eta)\Lambda$, that can be constructed deterministically in time* $\mathrm{poly}(b,q)$ *time. In the case when $q > \Lambda$, one can explicitly construct an $(r, 2r/\eta)$-subspace design with the same parameters.*

We note a couple of senses in which the parameters offered by the explicit construction are weaker than those guaranteed by the probabilistic construction. First, the total intersection dimension (40) is r^2/η rather than $O(r/\eta)$ (except when q is large). This is because, for small fields, their construction yields only a weak subspace design, incurring a factor r loss when passing to a subspace design. Second, the number of subspaces in the design is smaller, roughly $q^{\Omega(\eta\Lambda/r)}$ instead of $q^{\Omega(\eta\Lambda)}$. Finally, there is a modest restriction the field size q, and we need to pick r, Λ suitably to allow for fixed q. Fortunately, all these restrictions can be accommodated for our application. We remark that a recent construction of subspace designs based on cyclotomic function fields [20] gives an $(r, O(r \log_q \Lambda/\eta))$-subspace design over an *arbitrary* field $\mathbb{F}_q$; for our application, however, the r^2/η bound is more useful as $r \ll \Lambda$, and we cannot afford the dependence on Λ in the bound.

Let us now record a construction of a subspace that has large dimension and yet has low-dimensional intersection with every periodic subspace. The construction is based on the above subspace designs. This form will be convenient for later use in Section 11.2.1 for pre-coding Reed-Solomon codes with evaluation points in a subfield.

Theorem 10.5. *Let $\eta \in (0,1)$ and q be a prime power, and r, Λ, b be integers such that $r \leqslant \eta\Lambda/4$ and $b < q$. Then, one can construct a subspace V of $\mathbb{F}_q^{b\Lambda}$ of dimension at least $(1-\eta)b\Lambda$ in deterministic $q^{O(\Lambda)}$ time such that 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 $2r/\eta$.*

Proof. We will take $V = H_1 \times H_2 \times \cdots H_b$ where the H_i's belong to a $(r, 2r/\eta)$-subspace design in $\mathbb{F}_q^\Delta$ of cardinality b and dimension at least $(1-\eta)\Lambda$ as guaranteed by Theorem 10.4 when $q > \Lambda$. er 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 10.1. □

10.4. **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 19 (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 $(r_0, r_1, \ldots, r_l)$-cascaded subspace design with length-vector $(m_0, m_1, \ldots, m_l)$ and dimension vector $(d_0, d_1, \ldots, d_{l-1})$ 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 $\iota = 1, 2, \ldots, l$.*

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 10.1, 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 20 (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*

$$H_1^{(\iota)}, H_2^{(\iota)}, \cdots, H_{m_\iota/m_{\iota-1}}^{(\iota)} \subset \mathbb{F}_q^{m_{\iota-1}} \ , \ \textit{for } 1 \leqslant \iota \leqslant l \ .$$

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 $H_1^{(\iota)} \times H_2^{(\iota)} \times \cdots \times H_{m_\iota/m_{\iota-1}}^{(\iota)}$.*

In other words, we apply the construction of Lemma 10.1 for (disjoint) intervals of length m_ι at each level $\iota \in \{1, 2, \ldots, l\}$.

The following simple fact, which follows by counting number of linear constraints imposed, gives a lower bound on the dimension of a canonical subspace.

Observation 10.6. *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 10.1 which corresponds to the $l = 1$ case.

Lemma 10.7. *Suppose $\mathcal{M}$ is 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 10.1 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 10.1, 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 constructing a canonical subspace that has low-dimensional intersection with ultra-periodic subspaces based on the explicit subspace designs of Theorem 10.4. This statement will be used in Section 11.2.2 for pre-coding algebraic-geometric codes based on the Garcia-Stichtenoth tower.

Theorem 10.8. *Let $q \geqslant 4$ be a prime power. Let $\eta \in (0,1)$ and integers $\Lambda, r \geqslant 2$ satisfy $\Lambda \geqslant cr\eta^{-1}\log(r/\eta)$ for a large enough (absolute) constant $c > 0$. 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 $(r/\eta)^{2^{O(\log^* \kappa)}}$. The subspace U can be constructed in deterministically in* $\mathrm{poly}(\kappa, q)$ *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 $m_0 = \Lambda$, $m_1 = m_0 \cdot \lfloor (r/\eta)^{c/4} \rfloor$, and for $\iota \geqslant 0$, $m_{\iota+1} = m_\iota \cdot q^{\lceil \sqrt{m_\iota} \rceil}$. Let $r_0 = r$, and for $\iota \geqslant 0$, $r_{\iota+1} = \lceil r_\iota^2/\eta_\iota \rceil$. For this choice of parameters, one can verify that (i) $r_\iota \leqslant \eta_\iota m_\iota/4$, and (ii) $q^{\eta_\iota m_\iota/(4r_\iota)} \geqslant m_{\iota+1}/m_\iota$ for all $\iota \geqslant 0$. Indeed, to verify the first condition by induction, one only needs to check that $m_{\iota+1} \geqslant m_\iota^2$, which is true for $\iota = 0$ for a large enough choice of c, and for $\iota \geqslant 1$, $m_{\iota+1}$ in fact grows exponentially in $\sqrt{m_\iota}$. For the second condition, for $\iota = 0$ it follows from our assumption that $\Lambda \geqslant cr\eta^{-1}\log(r/\eta)$. For $\iota \geqslant 1$, it is implied by $r_\iota/\eta_\iota \ll \sqrt{m_\iota}/4$, which is true for $\iota = 1$ for large enough c, and for $\iota > 1$ by induction since r_ι/η_ι grows quadratically in each step, whereas m_ι grows exponentially.

We can therefore conclude by Theorem 10.4 that we can construct a $(r_\iota, r_{\iota+1})$-subspace design of cardinality $m_{\iota+1}/m_\iota$ in $\mathbb{F}_q^{m_\iota}$ of dimension $(1-\eta_\iota)m_\iota$.

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 $1 \leqslant \iota < l$, 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 re-denote 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 Theorem 10.4 implies that $\mathcal{M}_l$ can be constructed in $\mathrm{poly}(m_l, q)$ 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 subspace $U(\mathcal{M}_l) \subset \mathbb{F}_q^{m_l}$ padded with $\kappa - m_l$ zeroes at the end. By Observation 10.6, the dimension of U is at least

$$\begin{aligned} \left(1 - \sum_{\iota=0}^{l-1} \frac{\eta}{4 \cdot 2^\iota}\right) m_l &\geqslant (1-\eta/2)m_l > (1-\eta/2)(\kappa - m_{l-1}) \\ &> (1-\eta/2)\kappa - 2(\log_q \kappa)^2 > (1-\eta)\kappa \end{aligned}$$

for large enough κ. This proves that the subspace U has dimension at least $(1-\eta)\kappa$, and can be constructed deterministically in $\mathrm{poly}(q, \kappa)$ 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 10.7, $\mathrm{proj}_{m_l}(T) \cap U(\mathcal{M})$ 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, $r_l \leqslant (2r/\eta)^{2^l}$ and $l \leqslant O(\log^* \kappa)$. □

11. Pre-coding AG codes using h.s.e sets and subspace designs

In this final section, we combine the algebraic list decoding results (for folded AG codes and AG codes with subfield evaluation points) with subspace evasive sets (h.s.e sets and subspace designs) to deduce our main results on optimal rate list-decodable codes (Theorem 1.1). The idea is to pre-code the messages of the algebraic codes to belong to the subspace evasive sets, so that only a small number of candidates fall in the periodic (or ultrae-periodic) subspaces that arise in algebraic decoding and further they can be enumerated efficiently.

We stress that for our final code constructions either h.s.e sets or subspace designs can be used in combination with either the folded variant or the subspace evaluation variant. For concreteness though, below we focus on the following two combinations:

(1) folded codes with h.s.e sets, and
(2) subspace evaluation codes with subspace designs.

We note that the use of h.s.e sets leads to smaller final list size but their construction is randomized. Subspace designs lead to slightly larger list size (which in particular grows, albeit very slowly, with the block length) but the advantage is that they can be explicitly constructed.

11.1. **Pruning with h.s.e sets.**

We begin with pruning via h.s.e sets, applied to the folded Hermitian and folded Garcia-Stichtenoth codes from Sections 8.1 and 8.2 respectively. In particular, the combination of folded Garcia-Stichtenoth codes with h.s.e sets will give us our final main Monte Carlo code construction (part (i) of Theorem 1.1). We start with the folded Hermitian case as a warmup.

11.1.1. *Combining folded Hermitian codes and h.s.e sets.* Instead of encoding arbitrary $\mathbf{f} \in \mathbb{F}_q^k$ by the folded Hermitian code (Definition 11), we can restrict the messages $\mathbf{f}$ to belong to the range of our h.s.e set, so that the affine space of solutions guaranteed by Lemma 7.4 can be efficiently pruned to a small list. The formal claim is below.

Theorem 11.1. *Let $e \geqslant 2$ be an integer, $r \geqslant 2e$ be a large enough prime power, $q = r^2$, and $\zeta \in (1/q, 1)$. Let $k \leqslant q^{\zeta q/2}$ be a positive integer. Let s, m be positive integers satisfying $1 \leqslant s \leqslant m \leqslant q-1$ and $s < \zeta q/12$. Finally let N be an integer satisfying $k + 2er^e \leqslant Nm \leqslant (q-1)r^e$.*

Consider the code C_1 with encoding $E_1 : \mathbb{F}_q^{(1-\zeta)^2 k} \to (\mathbb{F}_q^m)^N$ defined as

$$E_1(\mathbf{x}) = \widetilde{FH}_e(N, k, q, m)(\mathsf{HSE}(\mathbf{x})) \ ,$$

for a random map $\mathsf{HSE} : \mathbb{F}_q^{(1-\zeta)^2 k} \to \mathbb{F}_q^k$ as constructed in Theorem 9.5 for a period size $\Delta = q-1$ and $b = \lceil \frac{k}{q-1} \rceil$.

Then, the code C_1 code has rate $R = (1-\zeta)^2 k/(Nm)$, can be encoded in $\mathrm{poly}(Nmq^{\zeta q})$ *time, and with high probability over the choice of HSE, it is (τ, ℓ)-list decodable in time*

$\mathrm{poly}(Nmq^{\zeta q})$ *for* $\ell \leqslant O(1/(R\zeta))$ *and*

$$\tau = \frac{s}{s+1}\left(1 - \frac{k}{N(m-s+1)}\right) - \frac{3m}{m-s+1}\frac{er^e}{mN} \ .$$

Proof. The claim about the rate is clear, and the encoding time follows from the time to compute HSE recorded in Theorem 9.5.

By (6), the genus $\mathfrak{g}_e \leqslant er^e$, and so the condition on N, m meets the requirement for the construction of the folded Hermitian tower based code in Definition 10, and the claimed value of the error fraction τ satisfies (36). By Lemma 6.4, we know that the candidate messages found by the decoder lie in one of at most q^{2Nm} possible $(s, q-1)$-periodic subspaces of $\mathbb{F}_q^k$.

One can check that the conditions of Theorem 9.5 are met for our choice of ζ, s, q, k, Δ. Appealing to Theorem 9.5 with the choice $c = 2Nm/k = O(1/R)$, we conclude that, with high probability over the choice of HSE, there is a decoding algorithm running in time $\mathrm{poly}(Nmq^{\zeta q})$ to list decode C_1 from a fraction τ of errors, outputting at most $O(1/(R\zeta))$ messages in the worst-case. □

Choosing parameters. Let $\varepsilon > 0$ be a small positive constant, and a family of codes of length N (assumed large enough) and rate $R \in (0,1)$ is sought. Pick n to be a growing parameter.

By picking $s = \Theta(1/\varepsilon)$, $m = \Theta(1/\varepsilon^2)$, $r = \lfloor \log n \rfloor$, $e = \lceil \frac{\log n}{\log\log n} \rceil$, $\zeta = (\log n \log\log n)^{-1}$, $N = \lfloor \frac{(r^2-1)r^e}{m} \rfloor$, and k proportional to Nm in Theorem 11.1, we can conclude the following.

Corollary 11.2. *For any* $R \in (0,1)$ *and positive constant* $\varepsilon \in (0,1)$*, there is a randomizedconstruction of a family of codes of rate at least R over an alphabet size* $(\log N)^{O(1/\varepsilon^2)}$ *that are encodable and* $(1 - R - \varepsilon, O(R^{-1}\log N \log\log N))$*-list decodable in* $\mathrm{poly}(N, 1/\varepsilon)$ *time, where N is the block length of the code.*

Our promised main result (Theorem 1.1) achieves better parameters than the above, namely an alphabet size of $\exp(\tilde{O}(1/\varepsilon^2))$ and list-size of $O(1/(R\varepsilon))$. This is based on the Garcia-Stichtenoth tower and is described next.

11.1.2. *Combining folded Garcia-Stichtenoth codes and h.s.e sets.* Similarly to Section 11.1.1, we now show how to pre-code the messages of the FGS code with a h.s.e subset. Here we will work with a base field $\mathbb{F}_q$ whose size is fixed and independent of the code dimension k, which will lead both to constant alphabet size and constant list size. To accommodate the requirement that $k \leqslant q^{O(\zeta\Delta)}$, we work with a larger "period" size for the h.s.e sets to evade. Recall Observation 3.2 which implies that if H is an (s, Δ, b)-ultra periodic subspace of $\mathbb{F}_q^k$ for $k = b\Delta$, then H is also $(su, \Delta u, b/u)$-periodic for every integer $u \geqslant 1$ with $u|b$. Thus we can scale up the period size using the ultra-periodicity of the subspace guaranteed by the decoder of Theorem 8.6. Following Remark 1, we actually only need a weaker property to prune via h.s.e sets which can also be ensured without ultra-periodicity. But since we have the stronger property available, we make use of it (this is also for sake of uniformity with the pruning based on subspace designs of Section 11.2 which can also be applied to the FGS code).

As in the Hermitian case, instead of encoding arbitrary $\mathbf{f} \in \mathbb{F}_q^k$ by the folded Garcia-Stichtenoth code, we will restrict the messages $\mathbf{f}$ to belong to the range of our h.s.e set. This will ensure that the affine space of solutions can be efficiently pruned to a small list.

Theorem 11.3. *Let r be a prime power, $q = r^2$, and $e \geqslant 2$ be an integer, and $\zeta \in (0,1)$. Let $\Delta \leqslant k$ be a multiple of $(r-1)$, say $\Delta = u(r-1)$ for a positive integer u. Let $k \leqslant q^{\zeta\Delta/2}$ be a positive integer.*

Let s, m be positive integers satisfying $1 \leqslant s \leqslant m \leqslant r-1$ and $s < \zeta r/12$. Finally let N be an integer satisfying $k + 2r^e \leqslant Nm \leqslant (r-1)r^e$.

Consider the code C_2 with encoding $E_2 : \mathbb{F}_q^{(1-\zeta)^2 k} \to (\mathbb{F}_q^m)^N$ defined as

$$E_2(\mathbf{x}) = \widetilde{FGS}_e(N,k,q,m)(\mathsf{HSE}(\mathbf{x})) \ ,$$

for a random map $\mathsf{HSE} : \mathbb{F}_q^{(1-\zeta)^2 k} \to \mathbb{F}_q^k$ as constructed in Theorem 9.5 for a period size Δ and $b = \lceil \frac{k}{\Delta} \rceil$.

The code C_2 has rate $R = (1-\zeta)^2 k/(Nm)$, can be encoded in $\mathrm{poly}(Nmq^{\zeta\Delta})$ time, and w.h.p over the choice of HSE, it is (τ, ℓ)-list decodable in time $\mathrm{poly}(Nmq^{\zeta\Delta})$ for $\ell \leqslant O(1/(R\zeta))$ and

$$\tau = \frac{s}{s+1}\left(1 - \frac{k}{N(m-s+1)}\right) - \frac{3m}{m-s+1}\frac{r^e}{mN} \ . \tag{41}$$

Proof. The proof is very similar to that of Theorem 11.1. The claim about the rate is clear, and the encoding time follows from the time to compute HSE recorded in Theorem 9.5.

The genus $\mathfrak{g}_e$ is now upper bounded by r^e, and so the condition on N, m meets the requirement for the construction of the folded the Garcia-Stichtenoth tower based code in Definition 12, and the claimed value of the error fraction τ satisfies (38). By Lemma 6.4, we know that the candidate messages found by the decoder lie in one of at most q^{2Nm} possible $(s, r-1, \lceil \frac{k}{r-1} \rceil)$-periodic subspaces.[11] Now by Observation **??**, each of these subspaces is also $(su, \Delta, \lceil \frac{k}{\Delta} \rceil)$-periodic. One can check that the conditions of Theorem 9.5 are met for our choice of ζ, s, q, k, Δ and taking su to play the role of s (since $s < \zeta r/12$, we have $su < \zeta\Delta/10$).

Appealing to Theorem 9.5 with the choice $c = 2Nm/k = O(1/R)$, we conclude that there is a decoding algorithm running in time $\mathrm{poly}(Nmq^{\zeta\Delta})$ to list decode C_2 from a fraction τ of errors, outputting at most $O(1/(R\zeta))$ messages in the worst-case. □

Choosing parameters. Finally, all that is left to be done is to pick parameters to show how the above can lead to optimal rate list-decodable codes over a constant-sized alphabet which further achieve very good list-size.

Let $\varepsilon > 0$ be a small positive constant, and a family of codes of length N (assumed large enough) and rate $R \in (0,1)$ is sought. Pick n to be a growing parameter.

Let us pick $s = \Theta(1/\varepsilon)$, $m = \Theta(1/\varepsilon^2)$, $\zeta = \varepsilon/12$, $r = \Theta(1/\varepsilon)$, $q = r^2$, and $e = \lceil \frac{\log n}{\log r} \rceil$, $N = \lfloor \frac{(r-1)r^e}{m} \rfloor$, and $k = RNm(1+\varepsilon)$. This ensures that (i) there are at least $n = Nm$ rational places and so we get a code of length at least $n/m = N$, (ii) the rate of the code C_2 is at least R, and (iii) the error fraction (41) is at least $1 - R - \varepsilon$.

The remaining part is to pick a multiple Δ of $(r-1)$ so that the $k \leqslant q^{\zeta\Delta/2}$ condition is met. This can be achieved by choosing $u = \lceil \frac{\log n}{\log(1/\varepsilon)} \rceil$ and $\Delta = (r-1)u$. With these

[11] Technically, it will belong to $\mathrm{proj}_k(W)$ of such a periodic subspace W, but we may pretend that there are $(r-1)\lceil k/(r-1) \rceil - k$ extra dummy coordinates which we decode. Or we can just assume for convenience that $r-1$ divides k.

choices, we can conclude the following, which is our main randomized code construction promised in part (i) of Theorem 1.1.

Theorem 11.4 (Main; Corollary to Theorem 11.3 with above choice of parameters)**.** *For any $R \in (0,1)$ and positive constant $\varepsilon \in (0,1)$, there is a Monte Carlo construction of a family of codes of rate at least R over an alphabet size $\exp(O(\log(1/\varepsilon)/\varepsilon^2))$ that are encodable and $(1-R-\varepsilon, O(1/(R\varepsilon))$-list decodable in $\mathrm{poly}(N)$ time, where N is the block length of the code.*

It may be instructive to recap why the Hermitian tower could not give a result like the above one. In the Hermitian case, the ratio $\mathfrak{g}_e/n$ of the genus to the number of rational places was about $e/r = e/\sqrt{q}$, and thus we needed $q > e^2$. Since the period Δ was about q, the running time of the decoder was bigger than $q^{\Omega(\zeta q)}$, whereas the length of the code was at most $q^{O(\sqrt{q})}$. This dictated the choice of $q \approx \log^2 n$, and then to keep the running time polynomial, we had to take $\zeta \approx (\log n \log\log n)^{-1}$.

11.2. **Pruning using subspace designs.** We now combine our the constructions of Reed-Solomon and Garcia-Stichtenoth codes with evaluation points in a subfield (from Section 7.1 and Section 8.3 respectively) with a pre-coding step that restricts the message coefficients to (respectively) the subspaces constructed in Theorem 10.5 (using subspace designs) and Theorem 10.8 (using cascaded subspace designs). These subcodes will then be list decodable with smaller list-size in polynomial time.

In particular, the combination of Garcia-Stichtenoth codes with cascaded subspace design will give us our final main deterministic code construction (part (ii) of Theorem 1.1).

11.2.1. *Subcodes of Reed-Solomon codes.* We begin with the case of Reed-Solomon codes as considered in Section 7.1. 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 $f_0, f_1, \dots, f_{k-1}$, we denote by $\mathbf{f_0}, \mathbf{f_1}, \dots, \mathbf{f_{k-1}}$ the representation of these coefficients as vectors in $\mathbb{F}_q^m$ by fixing some $\mathbb{F}_q$-basis of $\mathbb{F}_{q^m}$.

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 $(\mathbf{f_0}, \mathbf{f_1}, \dots, \mathbf{f_{k-1}}) \in V$ for a subspace $V \subseteq \mathbb{F}_q^{mk}$ guaranteed by Theorem 10.5, when applied with the parameter choices

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

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.[12]

Theorem 11.5. *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 $\widehat{\mathsf{RS}}$ that lie 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 7.4, 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 10.5, 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 $\mathbb{F}_q$. □

[12] 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 10.5.

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

Corollary 11.6. *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 close-by codewords. The code can be constructed deterministically in* $\operatorname{poly}(q)$ *time.*

We note that the above list decoding guarantee is in fact weaker than what is achieved for folded Reed-Solomon codes in [16], where the codewords were pinned down to a dimension $O(1/\varepsilon)$ subspace. We can improve the list size above to $\operatorname{poly}(1/\varepsilon)$ using pseudorandom subspace-evasive sets as in [16], or to $\exp(\varepsilon^{-O(1)})$ using the explicit subspace-evasive sets from [3]. 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.

11.2.2. *Subcodes of Garcia-Stichtenoth codes.* We now pre-code the codes constructed in Section 8.3. For a finite field $\mathbb{F}_q$, constant $\varepsilon > 0$, and integers s, m satisfying $1 \leqslant s \leqslant O(\varepsilon m/\log(1/\varepsilon))$ and $m \geqslant \Omega(1/\varepsilon^2)$, we will define subcodes of $\mathrm{GS}^{(q,m)}[N,k]$ guaranteed by Theorem 8.8. Note that messages space of this code can be identified with $\mathbb{F}_q^{mk}$.

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 10.8, when applied with the parameter choices

$$\eta = \varepsilon; \quad r = s-1; \quad \Lambda = m; \quad \kappa = km\ . \tag{42}$$

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 $\operatorname{poly}(k,m,q)$ time by virtue of the construction complexity of U.

Lemma 11.7. *Given an input string $\mathbf{y} \in \mathbb{F}_{q^m}^N$, a basis of an affine subspace of dimension at most*

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

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 $\operatorname{poly}(n, \log q, m)$ *time.*

Proof. By Theorem 8.8, we can compute the $(s-1,m)$-ultra periodic subspace T of messages whose encodings are within Hamming distance $\frac{s}{s+1}(N-k) - 3N/(\sqrt{q}-1)$ from $\mathbf{y}$. By Theorem 10.8, for the above choice of parameters (42), the intersection $T \cap U$ is an affine subspace over $\mathbb{F}_q$ of dimension $(s/\varepsilon)^{2^{O(\log^*(km)}}$. Since both steps involve only basic linear algebra, they can be accomplished using $\operatorname{poly}(N,m)$ operations over $\mathbb{F}_q$. □

By taking $q = \Theta(1/\varepsilon^2)$, and choosing $s = \Theta(1/\varepsilon)$ and $m = \Theta(\varepsilon^{-2}\log(1/\varepsilon))$ in the above lemma, we conclude our main result (stated informally as part (ii) of Theorem 1.1) concerning the explicit construction of codes list decodable up to the Singleton bound over fixed alphabets and very slowly growing list-size.

Theorem 11.8 (Main deterministic code construction)**.** *For every $R \in (0,1)$ and $\varepsilon > 0$, there is a deterministic polynomial time constructible family of error-correcting codes of rate R over an alphabet of size* $\exp(O(\varepsilon^{-2}\log^2(1/\varepsilon)))$ *that can be list decoded in polynomial time from a fraction $(1-R-\varepsilon)$ of errors, outputting a list of size at most $\exp_{1/\varepsilon}\Big(\exp_{1/\varepsilon}(\exp(O(\log^* N)))\Big)$, where N is block length of the code.*

Acknowledgments. We are grateful to the anonymous reviewers for detailed and perceptive comments which led to significant improvements in the organization and presentation of the paper.

References

1. Avraham Ben-Aroya and Igor Shinkar, *A note on subspace evasive sets*, Chicago Journal of Theoretical Computer Science (2014), 1–11.
2. Michael Braun, Michael Kiermaier, and Alfred Wassermann, *q-analogs of designs: Subspace designs*, Network Coding and Subspace Designs (Silberstein N. Vazquez-Castro M. Greferath M., Pavcevic M., ed.), Signals and Communication Technology, Springer, Cham, 2018.
3. Zeev Dvir and Shachar Lovett, *Subspace evasive sets*, Proceedings of the 44th Symposium on Theory of Computing Conference (STOC), 2012, pp. 351–358.
4. Peter Elias, *Error-correcting codes for list decoding*, IEEE Transactions on Information Theory **37** (1991), 5–12.
5. Michael A. Forbes and Venkatesan Guruswami, *Dimension expanders via rank condensers*, 19th International Workshop on Randomization and Computation (RANDOM), 2015, pp. 800–814.
6. Arnaldo Garcia and Henning Stichtenoth, *A tower of Artin-Schreier extensions of function fields attaining the Drinfeld-Vlădut bound*, Inventiones Mathematicae **121** (1995), 211–222.
7. ———, *On the asymptotic behavior of some towers of function fields over finite fields*, Journal of Number Theory **61** (1996), no. 2, 248–273.
8. Zeyu Guo and Noga Ron-Zewi, *Efficient list-decoding with constant alphabet and list sizes*, CoRR **abs/2011.05884** (2020).
9. Venkatesan Guruswami, *Cyclotomic function fields, Artin-Frobenius automorphisms, and list error-correction with optimal rate*, Algebra and Number Theory **4** (2010), no. 4, 433–463.
10. ———, *Linear-algebraic list decoding of folded Reed-Solomon codes*, Proceedings of the 26th IEEE Conference on Computational Complexity, June 2011.
11. Venkatesan Guruswami and Swastik Kopparty, *Explicit subspace designs*, Combinatorica **36** (2016), no. 2, 161–185, Preliminary version in FOCS 2013.
12. Venkatesan Guruswami and Srivatsan Narayanan, *Combinatorial limitations of average-radius list-decoding*, IEEE Transactions on Information Theory **60** (2014), no. 10, 5827–5842.
13. Venkatesan Guruswami and Atri Rudra, *Explicit codes achieving list decoding capacity: Error-correction with optimal redundancy*, IEEE Transactions on Information Theory **54** (2008), no. 1, 135–150.
14. ———, *Better binary list decodable codes via multilevel concatenation*, IEEE Trans. Information Theory **55** (2009), no. 1, 19–26.
15. Venkatesan Guruswami and Madhu Sudan, *Improved decoding of Reed-Solomon and Algebraic-geometric codes*, IEEE Transactions on Information Theory **45** (1999), no. 6, 1757–1767.
16. Venkatesan Guruswami and Carol Wang, *Linear-algebraic list decoding for variants of Reed-Solomon codes*, IEEE Transactions on Information Theory **59** (2013), no. 6, 3257–3268.
17. Venkatesan Guruswami, Carol Wang, and Chaoping Xing, *Explicit list-decodable rank-metric and subspace codes via subspace designs*, IEEE Trans. Information Theory **62** (2016), no. 5, 2707–2718.
18. Venkatesan Guruswami and Chaoping Xing, *Folded codes from function field towers and improved optimal rate list decoding*, Proceedings of the 44th Symposium on Theory of Computing Conference (STOC), 2012, pp. 339–350.
19. ———, *List decoding Reed-Solomon, algebraic-geometric, and Gabidulin subcodes up to the Singleton bound*, Proceedings of the ACM Symposium on Theory of Computing Conference (STOC), 2013, pp. 843–852.
20. Venkatesan Guruswami, Chaoping Xing, and Chen Yuan, *Constructions of subspace designs via algebraic function fields*, Trans. Amer. Math. Soc. **370** (2018), 8757–8775.
21. Ralf Koetter and Alexander Vardy, *Algebraic soft-decision decoding of Reed-Solomon codes*, IEEE Transactions on Information Theory **49** (2003), no. 11, 2809–2825.
22. Swastik Kopparty, *List-decoding multiplicity codes*, Theory Comput. **11** (2015), 149–182.
23. Swastik Kopparty, Noga Ron-Zewi, Shubhangi Saraf, and Mary Wootters, *Improved decoding of folded Reed-Solomon and multiplicity codes*, Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science, 2018, pp. 212–223.
24. Liming Ma and Chaoping Xing, *The asymptotic behavior of automorphism groups of function fields over finite fields*, Transactions of the Amercan Mathematical Society **372** (2019), 35–52.

25. Harald Niederreiter and Chaoping Xing, *Rational points on curves over finite fields–theory and applications*, Cambridge University Press, 2001.
26. Farzad Parvaresh and Alexander Vardy, *Correcting errors beyond the Guruswami-Sudan radius in polynomial time*, Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, 2005, pp. 285–294.
27. Ba-Zhong Shen, *A Justesen construction of binary concatanated codes that asymptotically meet the Zyablov bound for low rate*, IEEE Transactions on Information Theory **39** (1993), 239–242.
28. Kenneth Shum, Ilia Aleshnikov, P. Vijay Kumar, Henning Stichtenoth, and Vinay Deolalikar, *A low-complexity algorithm for the construction of algebraic-geometric codes better than the Gilbert-Varshamov bound*, IEEE Transactions on Information Theory **47** (2001), no. 6, 2225–2241.
29. Joseph H. Silverman, *The arithmetic of elliptic curves*, Springer, New York, 1985.
30. Henning Stichtenoth, *Algebraic function fields and codes*, Universitext, Springer-Verlag, Berlin, 1993.
31. Madhu Sudan, *Decoding of Reed-Solomon codes beyond the error-correction bound*, Journal of Complexity **13** (1997), no. 1, 180–193.
32. Salil Vadhan, *Pseudorandomness*, Foundations and Trends in Theoretical Computer Science (FnT-TCS), NOW publishers, 2010, To appear. Draft available at `http://people.seas.harvard.edu/~salil/pseudorandomness/`.

Computer Science Department, Carnegie Mellon University, Pittsburgh, USA.

E-mail address: `venkatg@cs.cmu.edu`

School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University; and, Division of Mathematical Sciences, School of Physical & Mathematical Sciences, Nanyang Technological University, Singapore.

E-mail address: `xingcp@ntu.edu.sg`

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il