

Cyclotomic function fields, Artin–Frobenius automorphisms, and list error correction with optimal rate

Venkatesan Guruswami

Algebraic error-correcting codes that achieve the optimal trade-off between rate and fraction of errors corrected (in the model of list decoding) were recently constructed by a careful “folding” of the Reed–Solomon code. The “low-degree” nature of this folding operation was crucial to the list decoding algorithm. We show how such folding schemes useful for list decoding arise out of the Artin–Frobenius automorphism at primes in Galois extensions. Using this approach, we construct new folded algebraic-geometric codes for list decoding based on cyclotomic function fields with a cyclic Galois group. Such function fields are obtained by adjoining torsion points of the Carlitz action of an irreducible $M \in \mathbb{F}_q[T]$. The Reed–Solomon case corresponds to the simplest such extension (corresponding to the case $M = T$). In the general case, we need to descend to the fixed field of a suitable Galois subgroup in order to ensure the existence of many degree 1 places that can be used for encoding.

Our methods shed new light on algebraic codes and their list decoding, and lead to new codes with optimal trade-off between rate and error correction radius. Quantitatively, these codes provide list decoding (and list recovery/soft decoding) guarantees similar to folded Reed–Solomon codes but with an alphabet size that is only polylogarithmic in the block length. In comparison, for folded RS codes, the alphabet size is a large polynomial in the block length. This has applications to fully explicit (with no brute-force search) binary concatenated codes for list decoding up to the Zyablov radius.

1. Introduction	434
2. Background on cyclotomic function fields	439
3. Reed–Solomon codes as cyclotomic function field codes	441

MSC2000: primary 11R60; secondary 14Q05, 11G30, 94B27, 12Y05, 68Q30.

Keywords: list decoding, algebraic-geometric codes, Galois extensions, Cyclotomic function fields, Reed–Solomon codes, Frobenius automorphisms.

An extended abstract of this work was presented at the 41st ACM Symposium on Theory of Computing, 2009. This version contains full proofs of the technical results. This material is based upon work supported by the National Science Foundation under Grant Numbers CCF-0343672, CCF-0953155, and ITR-0324906. The author’s research was also supported by a David and Lucile Packard Fellowship.

4. Subfield construction from cyclic cyclotomic function fields	442
5. Code construction from cyclotomic function fields	446
6. List decoding algorithm	452
7. Long codes with optimal rate for list decoding	458
Appendix: List of parameters	460
Acknowledgments	461
References	462

1. Introduction

1A. Background, context, and motivation. Error-correcting codes enable reliable transmission of information over a noisy communication channel (as well as reliable data storage and retrieval from a storage medium). The idea behind error-correcting codes is to *encode* the message to be transmitted (or stored) into a longer, redundant string called a *codeword*, which is then communicated over the noisy channel. This is accompanied by a *decoding* procedure that recovers the correct message even when several symbols in the transmitted codeword are corrupted. In this work, we focus on the worst-case model of errors; here, the channel noise has a single parameter $\rho \in (0, 1)$. We do not assume anything about how the errors are distributed beyond an upper bound of ρ on the total fraction of positions where errors may be caused.

The principal trade-off in this theory is between the redundancy and the fraction ρ of errors that can be corrected. Formally, a code is given by an injective encoding function $E : \Sigma^k \rightarrow \Sigma^n$. The *block length* of the code equals n , and Σ is its *alphabet*. The redundancy is measured by the *rate* R of the code, defined as the ratio k/n of the number of information symbols to the number of codeword symbols. The larger the rate, the less redundant the code. We are interested in an asymptotically good family of codes, that is, an infinite family of codes of increasing block lengths whose rates are lower bounded by R . The goal is to correct a fraction ρ of errors with as high a rate R as possible for the code family. It is simple to see that this rate R cannot exceed $1 - \rho$. Indeed, the channel could corrupt the last ρ fraction of symbols, and the first $(1 - \rho)n$ symbols should thus contain enough information to recover the Rn message symbols, implying $R \leq 1 - \rho$.

Quite remarkably, this simplistic upper bound can in fact be met, via a natural family of algebraic codes together with efficient decoding algorithms. Specifically, recent progress in algebraic coding theory [Parvaresh and Vardy 2005; Guruswami and Rudra 2008] has led to the construction of explicit codes over large alphabets that achieve the optimal rate versus error correction radius trade-off—namely, they admit efficient *list decoding* algorithms to correct close to the optimal fraction $1 - R$ of errors with rate R . List decoding is an error correction model where the

decoder is allowed to output a small list of messages which must include the correct message. Allowing such a list is essential in order to correct more than a fraction $(1 - R)/2$ of errors with rate R . In practice, having more than one codeword on the list is a rare event, and in case of multiple candidates, one can also return the closest codeword (as the one with the highest likelihood). Also, for many applications of codes, pinning down the message to a small list suffices, and some application-specific context information can be used to identify the correct message from the list. See, for instance, [Guruswami 2007, Chapter 1] or [Guruswami 2004] for more detailed background on list decoding.

We now return to the mathematics of optimal rate codes for list decoding. The algebraic codes constructed in [Guruswami and Rudra 2008] are *folded* Reed–Solomon (RS) codes, where the RS encoding $(f(1), f(\gamma), \dots, f(\gamma^{n-1}))$ of a low-degree polynomial $f \in \mathbb{F}_q[T]$ is viewed as a codeword of length $N = n/m$ over the alphabet \mathbb{F}_q^m by identifying successive blocks of m symbols. Here γ is a primitive element of the field \mathbb{F}_q .

Simplifying matters somewhat, the principal algebraic engine behind the list decoding algorithm in [Guruswami and Rudra 2008] was the identity $f(\gamma T) \equiv f(T)^q \pmod{(T^{q-1} - \gamma)}$, and the fact that $(T^{q-1} - \gamma)$ is irreducible over \mathbb{F}_q . This gave a low-degree algebraic relation between $f(T)$ and $f(\gamma T)$ in the residue field $\mathbb{F}_q[T]/(T^{q-1} - \gamma)$. This together with an algebraic relation found by a certain “interpolation step” during decoding enabled us to find the list of all relevant message polynomials $f(T)$ efficiently. Essentially, this gave two algebraically independent low-degree polynomial relations between the residues of $f(T)$ and $f(\gamma T)$ in the extension field $\mathbb{F}_q[T]/(T^{q-1} - \gamma)$. Solving these gives the list of possible values for $f(T) \pmod{(T^{q-1} - \gamma)}$, which also suffices to identify the message polynomial $f(T)$, as its degree is less than $q - 1$.

One of the motivations of this work is to gain a deeper understanding of the general algebraic principles underlying the above folding, with the hope of extending it to more general algebraic-geometric (AG) codes—an interesting algebraic question in its own right, but also important for potentially improving the alphabet size of the codes, as well as the decoding complexity and output list size of the decoding algorithm. (The large complexity and list size of the folded RS decoding algorithm in [Guruswami and Rudra 2008] are a direct consequence of the large degree q in the identity relating $f(\gamma T)$ and $f(T)$.)

The precursor to the folded RS codes were the Parvaresh–Vardy codes [2005]. Here the encoding of a message polynomial $f(T)$ consists of the evaluations of f at distinct elements of \mathbb{F}_q together with the evaluations of a few other algebraically related polynomials $f_1(T), \dots, f_m(T)$ (for some parameter $m \geq 1$) at these points. The algebraic relations between f_i and f are used at the decoder together with a multivariate polynomial relation between $f(T), f_1(T), \dots, f_m(T)$ to solve for

$f(T)$. An extension of the Parvaresh–Vardy codes [2005] to arbitrary AG codes was achieved in [Guruswami and Patthak 2008]. But in these codes, there is a substantial loss in rate since the encoding includes the evaluations of additional function(s) explicitly picked to satisfy a low-degree relation over some residue field. The crucial insight in the construction of folded RS codes was the fact that this additional function could just be the closely related function $f(\gamma T)$ — the image of $f(T)$ under the automorphism $T \mapsto \gamma T$ of $\mathbb{F}_q(T)$. It is a priori not clear for which algebraic function fields one can have a similar algebraic phenomenon and thereby deduce constructions of folded list-decodable codes analogous to folded RS codes.

1B. Summary of our contributions. We explain how folding schemes conducive to list decoding (such as the above relation between $f(\gamma T)$ and $f(T)$) arise out of the *Artin–Frobenius automorphism* at primes in Galois extensions. We then use this approach to construct new list-decodable folded AG codes based on *cyclotomic function fields* with a cyclic Galois group. Cyclotomic function fields [Carlitz 1938; Hayes 1974] are obtained by adjoining torsion points of the Carlitz action of an irreducible $M \in \mathbb{F}_q[T]$. The RS case corresponds to the simplest such extension (corresponding to the case $M = T$). In the general case, we need to descend to the fixed field of a suitable Galois subgroup in order to ensure the existence of many degree 1 places that can be used for encoding. We establish some key algebraic lemmas that characterize the desired subfield in terms of the appropriate generator μ in the algebraic closure of $\mathbb{F}_q(T)$ and its minimal polynomial over $\mathbb{F}_q(T)$. We then tackle the computational algebra challenge of computing a representation of the subfield and its rational places, and the message space, that is conducive for efficient encoding and decoding of the associated AG code.

Our constructions lead to some substantial quantitative improvements in the alphabet size, which we discuss in Section 1D. We also make some simplifications in the list decoding algorithm and avoid the need of a zero-increasing basis at each code place (Lemma 6.2). This, together with several other ideas, lets us implement the list decoding algorithm in polynomial time assuming *only* the natural representation of the code needed for efficient encoding, namely a basis for the message space. Computing such a basis remains an interesting challenge in computational function field theory. Our description and analysis of the list decoding algorithm in this work is *self-contained*, though it builds strongly on the framework of the algorithms in [Sudan 1997; Parvaresh and Vardy 2005; Guruswami and Patthak 2008; Guruswami and Rudra 2008].

1C. Galois extensions and Artin automorphisms in list decoding. We will now discuss how and why Artin–Frobenius automorphisms arise in the seemingly distant world of list decoding, and why we make the choice of cyclotomic function

fields for the underlying function field. In order to generalize the folding operation from the RS case, it is natural to look for function fields whose automorphisms we understand reasonably well. Galois extensions are a natural subclass of function fields to consider, with the hope that some automorphism in the Galois group will give a low-degree relation over some residue field. Unfortunately, the explicit constructions of good AG code families are typically based on a tower of function fields [Garcia and Stichtenoth 1995; 1996], where each step is Galois, but the whole extension is not. (Stichtenoth [2006] recently showed the existence of a Galois extension with the optimal trade-off between genus and number of rational places, but this extension is not, and cannot be, cyclic, as we require.)

In Galois extensions K/F , for each place A' in the extension field K , there is a special and important automorphism called the Artin–Frobenius automorphism (see, for example, [Marcus 1977, Chapter 4]) that simply powers the residue of any (regular) function at that place. The exponent or degree of this map is the norm of the place A of F lying below A' . Since the degree dictates the complexity of decoding, we would like this norm to be small. On the other hand, the residue field at A' needs to be large enough so that the message functions can be uniquely identified by their residue modulo A' . The most appealing way to realize this is if the place A is inert, that is, has a unique A' lying above it. However, this condition can only hold if the Galois group is cyclic, a rather strong restriction. For example, it is known [Frey et al. 1992] that even abelian extensions must be *asymptotically bad*.

In order to construct AG codes, we also need to have a good control of how certain primes split in the extension. For cyclotomic function fields, and of course their better-known number-theoretic counterparts $\mathbb{Q}(\omega)$ obtained by adjoining a root of unity ω , this theory is well-developed. As mentioned earlier, the cyclotomic function field we use itself has very few rational places. So we need to descend to an appropriate subfield where many degree 1 places of $\mathbb{F}_q(T)$ split completely, and develop some underlying theory concerning the structure of this subfield that can be exploited for efficient computation with them.

The Artin–Frobenius automorphism¹ is of course a well-known and fundamental notion in algebraic number theory, playing a role in the Chebotarev density theorem and Dirichlet’s theorem on infinitude of primes in arithmetic progressions, as well as quadratic and more general reciprocity laws. We find it rather intriguing that this notion ends up playing an important role in algorithmic coding theory as well.

¹Following [Rosen 2002], we will henceforth refer to the Artin–Frobenius automorphisms as simply Artin automorphisms. Many texts refer to these as Frobenius automorphisms. Since the latter term is most commonly associated with automorphism $x \mapsto x^q$ of \mathbb{F}_{q^m} , we use the term Artin automorphism to refer to the general notion that applies to all Galois extensions. The association of a place with its Artin–Frobenius automorphism is called the Artin map.

1D. Long codes achieving list decoding capacity and explicit binary concatenated codes. Quantitatively, our cyclotomic function field codes achieve list decoding (and list recovery²) guarantees similar to folded RS codes, but with an alphabet size that is only *polylogarithmic* in the block length. In comparison, for folded RS codes, the alphabet size is a large polynomial in the block length. We note that Guruswami and Rudra [2008] also present capacity-achieving codes of rate R for list decoding a fraction $(1 - R - \varepsilon)$ of errors with alphabet size $|\Sigma| = 2^{(1/\varepsilon)^{O(1)}}$, a fixed constant depending only on ε . But these codes do not have the strong “list recovery” (or more generally, soft decoding) property of folded RS codes.

Our codes inherit the powerful list recovery property of folded RS codes, which makes them very useful as outer codes in constructions of concatenated codes.³ In fact, due to their small alphabet size, they are even better in this role. Indeed, they can serve as outer codes for a family of concatenated codes list-decodable up to the so-called Zyablov radius, *with no brute-force search* for the inner codes. This is the first such construction for list decoding. It is similar to the “Justesen-style” explicit constructions for rate versus distance from [Justesen 1972; Shen 1993], except even easier, as one can use the ensemble of *all linear codes* instead of the succinct Wozencraft ensemble at the inner level of the concatenated scheme.

1E. Related work. Codes based on cyclotomic function fields have been considered previously in the literature. Some specific (nonasymptotic) constructions of function fields with many rational places over small fields \mathbb{F}_q ($q \leq 5$) appear in [Niederreiter and Xing 1996; 1997]. Cyclotomic codes based on the action of polynomials T^a for small a appear in [Quebbemann 1988], but decoding algorithms are not discussed for these codes, nor are these extensions cyclic as we require. Our approach is more general and works based on the action of an arbitrary irreducible polynomial. Exploiting the Artin automorphism of cyclotomic fields for an algorithmic purpose is also new to this work.

Independent of our work, Huang and Narayanan [2008] have considered AG codes constructed from Galois extensions, and observed how automorphisms of large order can be used for folding such codes. To our knowledge, the only instantiation of this approach that improves on folded RS codes is the one based on cyclotomic function fields from our work. As an alternate approach, they also

²List recovery is a generalization of list decoding where for each position a set of possible symbols is provided as input to the decoder, and the goal is to find all codewords that agree with some element of the input sets for at least a certain fraction of positions; see Remark 6.11.

³In binary concatenated codes, the message is first encoded by an “outer” code over a large alphabet Σ , and then each outer codeword symbol is encoded by an “inner” binary code $C_{\text{in}} : \Sigma \rightarrow \{0, 1\}^b$. Despite its simplicity, code concatenation remains the preeminent method for constructing good codes over small alphabets such as binary codes.

propose a decoding method that works with folding via automorphisms of small order. This involves computing several coefficients of the power series expansion of the message function at a low-degree place. Unfortunately, piecing together these coefficients into a function could lead to an exponential list size bound. The authors suggest a heuristic assumption under which they can show that for a *random* received word, the expected list size and running time are polynomially bounded.

2. Background on cyclotomic function fields

We assume familiarity with basic background on global fields and their extensions such as valuations and places, Galois extensions, decomposition of primes, ramification, Artin–Frobenius automorphism, etc. In this section, we will focus on background material concerning cyclotomic function fields. These are the function-field analog of the classic cyclotomic number fields from algebraic number theory. This theory was developed by Hayes [1974], building upon ideas due to Carlitz [1938]. The objective was to develop an explicit class field theory classifying all abelian extensions of the rational function field $\mathbb{F}_q(T)$, analogous to classic results for \mathbb{Q} and imaginary quadratic extensions of \mathbb{Q} . The common idea in these results is to allow a ring of “integers” in the ground field to act on part of its algebraic closure, and obtain abelian extensions by adjoining torsion points of this action. We will now describe these extensions of $\mathbb{F}_q(T)$.

Let T be an indeterminate over the finite field \mathbb{F}_q . Let $R_T = \mathbb{F}_q[T]$ denote the polynomial ring, and $F = \mathbb{F}_q(T)$ the field of rational functions. Let F^{ac} be a fixed algebraic closure of F . Let $\text{End}_{\mathbb{F}_q}(F^{\text{ac}})$ be the ring of \mathbb{F}_q -endomorphisms of F^{ac} , thought of as a \mathbb{F}_q -vector space. We consider two special elements of $\text{End}_{\mathbb{F}_q}(F^{\text{ac}})$:

- (i) the Frobenius automorphism τ defined by $\tau(z) = z^q$ for all $z \in F^{\text{ac}}$, and
- (ii) the map μ_T defined by $\mu_T(z) = Tz$ for all $z \in F^{\text{ac}}$.

The substitution $T \rightarrow \tau + \mu_T$ yields a ring homomorphism from R_T to $\text{End}_{\mathbb{F}_q}(F^{\text{ac}})$ given by

$$f(T) \mapsto f(\tau + \mu_T).$$

Using this, we can define the *Carlitz action* of R_T on F^{ac} as follows: for $M \in R_T$,

$$C_M(z) = M(\tau + \mu_T)(z) \quad \text{for all } z \in F^{\text{ac}}.$$

This action endows F^{ac} with the structure of an R_T -module, which is called the Carlitz module. For a nonzero polynomial $M \in R_T$, define the set

$$\Lambda_M = \{z \in F^{\text{ac}} \mid C_M(z) = 0\},$$

to consist of the M -torsion points of F^{ac} , that is, the elements annihilated by the Carlitz action of M (this is also the set of zeroes of the polynomial $C_M(Z) \in R_T[Z]$).

Proposition 2.2. *Let $M \in R_T$, $M \neq 0$, be a monic, irreducible polynomial of degree d .*

- (i) *Ramification at M : the place M is totally ramified in the extension $F(\Lambda_M)/F$. If $\lambda \in \Lambda_M$ is a root of $C_M(z)/z$ and \tilde{M} is the unique place of $F(\Lambda_M)$ lying above M , then λ is a \tilde{M} -prime element, that is, $v_{\tilde{M}}(\lambda) = 1$.*
- (ii) *Ramification at P_∞ : the infinite place P_∞ of F , that is, the pole of T , splits into $(q^d - 1)/(q - 1)$ places of degree 1 in $F(\Lambda_M)/F$, each with ramification index $(q - 1)$. Its decomposition group equals \mathbb{F}_q^* .*
- (iii) *Splitting at other places: if $P \in R_T$ is a monic, irreducible polynomial different from M , then P is unramified in $F(\Lambda_M)/F$, and splits into $(q^d - 1)/f$ primes of degree $f \deg P$ where f is the order of P modulo M (that is, the smallest positive integer e such that $P^e \equiv 1 \pmod{M}$).*

3. Reed–Solomon codes as cyclotomic function field codes

We now discuss how RS codes arise out of the simplest cyclotomic extension $F(\Lambda_T)/F$. This serves both as a warm-up for our later results, and as a method to illustrate that one can view the folding employed in [Guruswami and Rudra 2008] as arising naturally from the Artin automorphism at a certain prime in the extension $F(\Lambda_T)/F$.

We have $\Lambda_T = \{u \in F^{\text{ac}} \mid u^q + Tu = 0\}$. Pick a nonzero $\lambda \in \Lambda_T$. By Proposition 2.2, the only ramified places in $F(\Lambda_T)/F$ are T and the pole P_∞ of T . Both of these are totally ramified and have a unique place above them in $F(\Lambda_T)$. Denote by Q_∞ the place above P_∞ in $F(\Lambda_T)$.

We have $\lambda^{q-1} = -T$, so λ has a pole of order one at Q_∞ , and no poles elsewhere. The place $T + 1$ splits completely into $n = q - 1$ places of degree 1 in $F(\Lambda_T)$. The evaluation of λ at these places corresponds to the roots of $x^{q-1} = 1$, that is, to nonzero elements of \mathbb{F}_q . Thus the places above $T + 1$ can be described as $P_1, P_\gamma, \dots, P_{\gamma^{q-2}}$, where γ is a primitive element of \mathbb{F}_q and $\lambda(P_{\gamma^i}) = \gamma^i$ for $i = 0, 1, \dots, q - 2$.

For $k < q - 1$, define $\mathcal{M}_k = \{\sum_{i=0}^{k-1} \beta_i \lambda^i \mid \beta_i \in \mathbb{F}_q\}$. \mathcal{M}_k has q^k elements, each with at most $(k - 1)$ poles at Q_∞ and no poles elsewhere. Consider the \mathbb{F}_q -linear map $E_{\text{RS}} : \mathcal{M}_k \rightarrow \mathbb{F}_q^n$ defined as

$$E_{\text{RS}}(f) = (f(P_1), f(P_\gamma), \dots, f(P_{\gamma^{q-2}})).$$

Clearly this just defines an $[n, k]_q$ RS code, consisting of evaluations of polynomials of degree $< k$ at elements of \mathbb{F}_q^* .

Consider the place $T + \gamma$ of F . The condition $(T + \gamma)^f \equiv 1 \pmod{T}$ is satisfied if and only if $\gamma^f = 1$, which happens if and only if $(q - 1) \mid f$. Therefore, the place

$T + \gamma$ remains inert in $F(\Lambda_T)/F$. Let A denote the unique place above $T + \gamma$ in $F(\Lambda_T)$. The degree of A equals $q - 1$.

The Artin automorphism at A , σ_A , is given by $\sigma_A(\lambda) = C_{T+\gamma}(\lambda) = C_\gamma(\lambda) = \gamma\lambda$. Note that this implies $f(P_{\gamma^{i+1}}) = \sigma_A(f)(P_{\gamma^i})$ for $0 \leq i < q - 2$. By the property of the Artin automorphism, we have $\sigma_A(f) \equiv f^q \pmod{A}$ for all $f \in R_T[\lambda]$. Note that this is same as the condition $f(\gamma\lambda) \equiv f(\lambda)^q \pmod{(\lambda^{q-1} - \gamma)}$ treating f as a polynomial in λ . This corresponds to the algebraic relation between $f(X)$ and $f(\gamma X)$ in the ring $\mathbb{F}_q[X]$ that was used by Guruswami and Rudra [2008] in their decoding algorithm, specifically in the task of finding all $f(X)$ of degree less than k satisfying $Q(X, f(X), f(\gamma X)) = 0$ for a given $Q \in \mathbb{F}_q[X, Y, Z]$. In the cyclotomic language, this corresponds to finding all $f \in R_T[\lambda]$ with fewer than k poles at Q_∞ satisfying $Q(f, \sigma_A(f)) = 0$ for $Q \in R_T[\lambda](Y, Z)$. Since $\deg A = q - 1 \geq k$, f is determined by its residue at A , and we know $\sigma_A(f) \equiv f^q \pmod{A}$. Therefore, we can find all such f by finding the roots of the univariate polynomial $Q(Y, Y^q) \pmod{A}$ over the residue field \mathbb{O}_A/A .

4. Subfield construction from cyclic cyclotomic function fields

In this section, we will construct the function field construction that will be used for our AG codes, and establish the key algebraic facts concerning it. The approach will be to take the cyclotomic field $K = F(\Lambda_M)$, where M is an irreducible of degree $d > 1$, and get a code over \mathbb{F}_q . But the only places of degree 1 in $F(\Lambda_M)$ are the ones above the pole P_∞ of T . There are only $(q^d - 1)/(q - 1)$ such places above P_∞ , which is much smaller than the genus. So we descend to a subfield where many degree 1 places split completely. This is done by taking a subgroup H of $(\mathbb{F}_q[T]/(M))^*$ with many degree 1 polynomials and considering the fixed field $E = K^H$. For every irreducible $N \in R_T$ such that $\bar{N} = N \pmod{M} \in H$, the place N splits completely in the extension E/F (this follows from the fact that C_N is the Artin automorphism at the place N). This technique has also been used in works mentioned earlier [Quebbemann 1988; Niederreiter and Xing 1996; 1997], though our approach is more general and works with any irreducible M . The study of algorithms for cyclotomic codes and the role played by the Artin automorphism in their list decoding is also novel to our work.

4A. Table of parameters. Since there is an unavoidable surfeit of notation and parameters used in this section and Section 5, we summarize them for easy reference in the Appendix.

4B. Function field construction. Let \mathbb{F}_r be a subfield of \mathbb{F}_q . Let $M \in \mathbb{F}_r[T]$ be a monic polynomial that is irreducible over \mathbb{F}_q (note that we require $M(T)$ to have coefficients in the smaller field \mathbb{F}_r , but demand irreducibility in the ring $\mathbb{F}_q[T]$).

The following lemma follows from the general characterization of when binomials $T^m - \alpha$ are irreducible in $\mathbb{F}_q[T]$ [Lidl and Niederreiter 1986, Chapter 3].

Lemma 4.1. *Let $d \geq 1$ be an odd integer such that every prime factor of d divides $(r - 1)$ and $\gcd(d, (q - 1)/(r - 1)) = 1$. Let γ be a primitive element of \mathbb{F}_r . Then $T^d - \gamma \in \mathbb{F}_r[T]$ is irreducible in $\mathbb{F}_q[T]$.*

A simple choice for which the above conditions are met is $r = 2^a$, $q = r^2$, and $d = r - 1$ (we will need a more complicated choice for our list decoding result in [Theorem 7.1](#)). For the sake of generality as well as clarity of exposition, we will develop the theory without making specific choices for the parameters, a somewhat intricate task we will undertake in [Section 7](#).

For the rest of this section, fix $M(T) = T^d - \gamma$ as guaranteed by [Lemma 4.1](#). We continue with the notation $F = \mathbb{F}_q(T)$, $R_T = \mathbb{F}_q[T]$, and $K = F(\Lambda_M)$. Fix a generator $\lambda \in \Lambda_M$ of K/F so that $K = F(\lambda)$.

Let G be the Galois group of K/F , which is isomorphic to the cyclic multiplicative group $(\mathbb{F}_q[T]/(M))^*$. Let $H \subset G$ be the subgroup $\mathbb{F}_q^* \cdot (\mathbb{F}_r[T]/(M))^*$. The cardinality of H is $(r^d - 1) \cdot (q - 1)/(r - 1)$. Note that since G is cyclic, there is a unique subgroup H of this size. Indeed, if $\Gamma \in G$ is an arbitrary generator of G , then $H = \{1, \Gamma^b, \Gamma^{2b}, \dots, \Gamma^{q^d - 1 - b}\}$, where

$$b = \frac{|G|}{|H|} = \frac{q^d - 1}{r^d - 1} \cdot \frac{r - 1}{q - 1}. \quad (4-1)$$

Let $A \in R_T$ be an arbitrary polynomial such that $A \bmod M$ is a generator of $(\mathbb{F}_q[T]/(M))^*$. We can then take Γ so that $\Gamma(\lambda) = C_A(\lambda)$. We fix a choice of A in the sequel and assume that A is precomputed and known. In [Section 5C](#) we will pick such an A of appropriately large degree D . The effective version of Dirichlet's theorem for irreducible polynomials in arithmetic progressions guarantees the existence of such polynomials A for large enough degree [[Rosen 2002](#), Theorem 4.8].

Note that by [Proposition 2.1\(ii\)](#), the Galois action commutes with the Carlitz action and therefore $\Gamma^j(\lambda) = C_{A^j}(\lambda)$ for all $j \geq 1$. Thus knowing the polynomial A lets us compute the action of the automorphisms of H on any desired element of $K = F(\lambda)$.

Let $E \subset K$ be the subfield of K fixed by the subgroup H , that is,

$$E = \{x \in K \mid \sigma(x) = x \text{ for all } \sigma \in H\}.$$

The field E will be the one used to construct our codes. We first record some basic properties of the extension E/F , and how certain places decompose in this extension.

Then the fixed field $E = K^H$ equals the extension field $F(\mu)$. The minimal polynomial $h \in R_T[Z]$ of μ over F is given by

$$h(Z) = \prod_{j=0}^{b-1} (Z - \Gamma^j(\mu)).$$

Further, the polynomial $h(Z)$ can be computed in $q^{O(d)}$ time.

Proof (sketch). By definition, μ is fixed by each $\pi \in H$ and so $\mu \in E$. Therefore $F(\mu) \subseteq E$. To show $E = F(\mu)$, we will argue that $[F(\mu) : F] = b$, which in turn follows if we show that $h(Z)$ has coefficients in F and is irreducible over F . It is easy to see that the coefficients of h are fixed by Γ and hence by all of $\text{Gal}(K/F)$, and so must belong to F . Since λ and all its Galois conjugates $C_{A^i}(\lambda)$ are integral over F , each $\Gamma^j(\mu)$ is integral over F , and thus so is each coefficient of h . But since we already know they belong to F , the coefficients must in fact lie in R_T .

The irreducibility of h over R_T can be shown using Eisenstein's criterion with respect to M . Indeed, except the leading coefficient, every other coefficient of h is divisible by λ , and since $\lambda \in \tilde{M}$ (by [Proposition 2.2](#)), these coefficients belong to the ideal $F \cap \tilde{M} = M$. The constant term of h equals $\prod_{0 \leq i < q^d - 1} C_{A^i}(\lambda)$, which is also the constant term of

$$C_M(Z)/Z = \prod_{0 \leq i < q^d - 1} (Z - C_{A^i}(\lambda)).$$

The latter equals M by [Proposition 2.1\(i\)](#). Thus the constant term of h is not divisible by M^2 . By Eisenstein's criterion, h must be irreducible over F .

Finally, we address how the coefficients of $h(Z)$ can be computed efficiently. Note that for $j = 0, 1, \dots, b-1$,

$$\Gamma^j(\mu) = \prod_{\substack{0 \leq i < q^d - 1 \\ i \bmod b = j}} \Gamma^i(\lambda) = \prod_{\substack{0 \leq i < q^d - 1 \\ i \bmod b = j}} C_{A^i}(\lambda). \quad (4-4)$$

Using this, we can compute $\Gamma^j(\mu)$ for $0 \leq j \leq b-1$ as a formal polynomial in λ with coefficients from R_T . We can divide this polynomial by the monic polynomial $C_M(\lambda)/\lambda$ (formally, over the polynomial ring $R_T[\lambda]$) and represent $\Gamma^j(\mu)$ as a polynomial of degree less than $(q^d - 1)$ in λ . Using this representation, we can compute the polynomials

$$h^{(i)}(Z) = \prod_{j=0}^i (Z - \Gamma^j(\mu)) \quad \text{for } 1 \leq i \leq b-1$$

iteratively, as an element of $R_T[\lambda][Z]$, with all coefficients having degree less than $(q^d - 1)$ in λ . When $i = b-1$, we would have computed $h(Z)$ — we know at the end all the coefficients will have degree 0 in λ and belong to R_T . \square

Using $\prod_{j=0}^{b-1} \Gamma^j(\mu) = M$ from the above argument, and $v_{M'}(\Gamma^j(\mu)) = v_{M'}(\mu)$, we conclude that $v_{M'}(\mu) = 1$, that is, μ (as well as each of its Galois conjugates $\Gamma^j(\mu)$) is M' -prime. We record this fact below. It will be useful to establish that the integral closure of R_T in E equals $R_T[\mu]$ (Proposition 5.1), a fact we will use en route characterizing the message space in Theorem 5.2.

Lemma 4.4. *The element μ has a simple zero at M' , that is, $v_{M'}(\mu) = 1$.*

With the minimal polynomial $h(Z)$ of μ at our disposal, we turn to computing the evaluations of μ at the b places above $T - \beta$; call them $P_j^{(\beta)}$ for $j = 0, 1, \dots, b - 1$, for each $\beta \in \mathbb{F}_r$. (Recall that the place $T - \beta$ splits completely in E/F by Proposition 4.2(v).) The following lemma identifies the set of evaluations of μ at these places. This method is related to Kummer’s theorem on splitting of primes [Stichtenoth 1993, Section III.3].

Lemma 4.5. *Consider the polynomial $\bar{h}^{(\beta)}(Z) \in \mathbb{F}_q[Z]$ obtained by evaluating the coefficients of $h(Z)$, which are polynomials in T , at β . Then*

$$\bar{h}^{(\beta)}(Z) = \prod_{j=0}^{b-1} (Z - \mu(P_j^{(\beta)})).$$

In particular, the set of evaluations of μ at the places above $(T - \beta)$ equals the roots of $\bar{h}^{(\beta)}$ in \mathbb{F}_q , and can be computed in $b^{O(1)}$ time given $h \in R_T[Z]$.

Proof. We know $h(Z) = \prod_{j=0}^{b-1} (Z - \Gamma^j(\mu))$. Therefore

$$\bar{h}^{(\beta)}(Z) = \prod_{j=0}^{b-1} (Z - \Gamma^j(\mu)(P_0^{(\beta)})) = \prod_{j=0}^{b-1} (Z - \mu(\Gamma^{-j}(P_0^{(\beta)}))) = \prod_{j=0}^{b-1} (Z - \mu(P_j^{(\beta)})),$$

where the last step uses the fact that $\Gamma^{-j}(P_0^{(\beta)})$ for $j = 0, 1, \dots, b - 1$ is precisely the set of places above $T - \beta$. □

5. Code construction from cyclotomic function fields

We will now describe the AG codes based on the function field E . A tempting choice for the message space is perhaps $\{\sum_{i=0}^{b-1} a_i(T)\mu^i\} \subset R_T[\mu]$, where $a_i(T)$ are polynomials of some bounded degree. This is certainly a \mathbb{F}_q -linear space and messages in this space have no poles outside the places lying above P_∞ . However, the valuations of μ at these places are complicated — one needs the Newton polygon method to estimate them [Villa Salvador 2006, Section 12.4] — and since μ has both zeroes and poles among these places, it is hard to get good bounds on the total pole order of such messages at each of the places above P_∞ .

5A. Message space. Let M' be the unique totally ramified place M' in E lying above M ; $\deg M' = \deg M = d$. We will use as message space elements of $R_T[\mu]$ that have no more than a certain number ℓ of poles at the place M' and no poles elsewhere. These can equivalently be thought of (via a natural correspondence) as elements of E that have bounded (depending on ℓ) pole order at each place above P_∞ , and no poles elsewhere, and we can develop our codes and algorithms in this equivalent setting. Since the literature on AG codes typically focuses on one-point codes where the messages have poles at a unique place, we work with functions with poles restricted to M' .

Formally, for an integer $\ell \geq 1$, let $\mathcal{L}(\ell M')$ be the space of functions in E that have no poles outside M' and at most ℓ poles at M' . $\mathcal{L}(\ell M')$ is an \mathbb{F}_q -vector space, and by the Riemann–Roch theorem, $\dim \mathcal{L}(\ell M') \geq \ell d - g + 1$, where $g = d(b - 1)/2 + 1$ is the genus of E . We will assume that $\ell \geq b$, in which case $\dim \mathcal{L}(\ell M') = \ell d - g + 1$.

We will represent the code by a basis of $\mathcal{L}(\ell M')$ over \mathbb{F}_q . Of course, we first need to understand how to represent a single function in $\mathcal{L}(\ell M')$. [Theorem 5.2](#) below suggests a representation for elements of $\mathcal{L}(\ell M')$ that we can use. Its proof uses the following claim, which can be established using [Lemma 4.4](#) and an argument similar to the one used to prove that the integral closure of R_T in $K = F(\lambda)$ equals $R_T[\lambda]$ [[Rosen 2002](#), Proposition 12.9].

Proposition 5.1. *The integral closure of R_T in E equals*

$$R_T[\mu] = \left\{ \sum_{i=0}^{b-1} a_i \mu^i \mid a_i \in R_T \right\}.$$

Theorem 5.2. *A function f in E with poles only at M' has a unique representation of the form*

$$f = \frac{\sum_{i=0}^{b-1} a_i \mu^i}{M^e}, \tag{5-1}$$

where $e \geq 0$ is an integer, each $a_i \in R_T$, and not all the a_i 's are divisible by M (as polynomials in T).

Proof. If f has poles only at M' , there must be a smallest integer $e \geq 0$ such that $M^e f$ has no poles outside the places above P_∞ . This means that $M^e f$ must belong to the integral closure (ring of integers) of R_T in E , that is, the minimal polynomial of $M^e f$ over R_T is monic. By [Proposition 5.1](#), we have $M^e f \in R_T[\mu]$ and so we can write $f = M^{-e} \sum_{i=0}^{b-1} a_i \mu^i$ as claimed. The uniqueness of the representation follows since $\{1, \mu, \dots, \mu^{b-1}\}$ forms a basis of E over F . \square

5B. Succinctness of representation. In order to be able to efficiently compute with the representation (5-1) of functions in $\mathcal{L}(\ell M')$, we need the guarantee that the representation will be *succinct*, that is, of size polynomial in the code length.

We show that this will be the case by obtaining an upper bound on the degree of the coefficients $a_i \in R_T$ in Lemma 5.3 below. This is not as straightforward as one might hope, and we thank G. Anderson and D. Thakur for help with its proof. For the choice of parameters we will make (in Theorems 6.10 and 7.1), this upper bound will be polynomially bounded in the code length. Therefore, the assumed representation of the basis functions is of polynomial size.

Lemma 5.3. *Suppose $f \in \mathcal{L}(\ell M')$ is given by $f = M^{-e} \sum_{i=0}^{b-1} a_i \mu^i$ for $a_i \in R_T$ (not all divisible by M) and $e \geq 0$. Then the degree of each a_i is at most $\ell + q^d b$.*

Proof. Let $g = M^e f = \sum_{i=0}^{b-1} a_i \mu^i$. We know that g has at most eb poles at each place of E that lies above P_∞ (since f has no poles at these places). Using the fact that f has at most ℓ poles at M' , and the uniqueness of the representation $f = M^{-e} \sum_{i=0}^{b-1} a_i \mu^i$, it is easy to argue that $eb \leq \ell + b$. So, g has at most $\ell + b$ poles at each place of E lying above P_∞ .

Let $\sigma = \sigma_A$; we know that σ is a generator of $\text{Gal}(E/F)$. For $j = 0, 1, \dots, b-1$, we have $\sigma^j(g) = \sum_{i=0}^{b-1} a_i \sigma^j(\mu^i)$. Let $\mathbf{a} = (a_0, a_1, \dots, a_{b-1})^T$ be the (column) vector of coefficients, and let $\mathbf{g} = (g, \sigma(g), \dots, \sigma^{b-1}(g))^T$. Denoting by Φ the $b \times b$ matrix with $\Phi_{ji} = \sigma^j(\mu^i)$ for $0 \leq i, j \leq b-1$, we have the system of equations $\Phi \mathbf{a} = \mathbf{g}$.

We can thus determine the coefficients a_i by solving this linear system. By Cramér's rule, $a_i = \det \Phi_i / \det \Phi$, where Φ_i is obtained by replacing the i -th column of Φ by the column vector \mathbf{g} . The square of the denominator $\det \Phi$ is the discriminant of the field extension E/F , and belongs to R_T . Thus the degree of a_i is at most the pole order of $\det \Phi_i$ at an arbitrary place, say \tilde{P} , above P_∞ . By the definition (4-3) of μ , and the fact that λ and its conjugates have at most one pole at the places above P_∞ in $F(\Lambda_M)$, it follows that μ has at most $(q^d - 1)/b$ poles at \tilde{P} . The same holds for all its conjugates $\sigma^j(\mu)$. The function g and its conjugates $\sigma^j(g)$ have at most $\ell + b$ poles at \tilde{P} . All in all, this yields a crude upper bound of

$$\frac{q^d - 1}{b} \frac{(b-1)b}{2} + \ell + b \leq \ell + q^d b$$

for the pole order of $\det \Phi_i$ at \tilde{P} , and hence also the degree of the polynomial $a_i \in R_T$. \square

5C. Rational places for encoding and their ordering. So far, the polynomial $A \in R_T$ was any monic irreducible polynomial that was a primitive element modulo M , so that its Artin automorphism σ_A generates $\text{Gal}(E/F)$. We will now pick A to have degree D satisfying

$$D > \frac{\ell d}{b} \quad \text{and} \quad D > 3d, \tag{5-2}$$

where the latter condition (in fact even $D > 2d + o(d)$ suffices) ensures that there are at least $q^D/2Dq^d$ irreducible polynomials of degree D with any desired residue modulo M . This follows from the effective version of Dirichlet's theorem for polynomials; see for instance [Rosen 2002, Theorem 4.8].

For D satisfying the above conditions, an irreducible polynomial A of degree D that is primitive modulo M can be found by a Las Vegas algorithm in $(Dq^d)^{O(1)}$ time by picking a random polynomial and checking that it works, or deterministically by brute force in $q^{O(d+D)}$ time. Both of these bounds are within the decoding time claimed in Theorem 6.10, and will be polynomial in the block length for our parameter choices in Theorem 7.1. By Proposition 2.1, A remains inert in E/F , and let us denote by A' the unique place of E that lies over A . The degree of A' equals Db .

For each $\beta \in \mathbb{F}_r$, fix an arbitrary place $P_0^{(\beta)}$ lying above $T - \beta$ in E . For $j = 0, 1, \dots, b-1$, define

$$P_j^{(\beta)} = \sigma_A^{-j}(P_0^{(\beta)}) . \quad (5-3)$$

Since $\text{Gal}(E/F)$ acts transitively on the set of primes above a prime, and σ_A generates $\text{Gal}(E/F)$, these constitute all the places above $T - \beta$. Lemma 4.5 already tells us the set of evaluations of μ at these places, but not which evaluation corresponds to which point. We have $\mu(\sigma_A^{-j}(P_0^{(\beta)})) = \sigma_A^j(\mu)(P_0^{(\beta)})$; hence, to compute the evaluations of μ at all these b places according to the ordering (5-3), it suffices to know

- (i) the value at $\mu(P_0^{(\beta)})$, which we can find by simply picking one of the roots from Lemma 4.5 arbitrarily, and
- (ii) a representation of $\sigma_A(\mu)$ as an element of $R_T[\mu]$ (since $\sigma_A(\mu)$ is integral over R_T , it belongs to $R_T[\mu]$ by virtue of Proposition 5.1). Note that $T(P_0^{(\beta)}) = \beta$, so once we know $\mu(P_0^{(\beta)})$, we can evaluate any element of $R_T[\mu]$ at $P_0^{(\beta)}$.

We now show that $\sigma_A(\mu) \in R_T[\mu]$ can be computed efficiently.

- Lemma 5.4.** (i) *The values of $\sigma_A^j(\mu)$ for $0 \leq j \leq b-1$ as elements of $R_T[\mu]$ can be computed in $q^{O(d)}$ time.*
- (ii) *The values $\mu(P_j^{(\beta)})$ for $\beta \in \mathbb{F}_r$ and $j = 0, 1, \dots, b-1$ can be computed in $q^{O(d)}$ time. Knowing these values, we can compute any function in the message space $\mathcal{L}(\ell M')$ represented in the form (5-1) at the places $P_j^{(\beta)}$ in $\text{poly}(\ell, q^d)$ time.*

Proof. Part (ii) follows from (i) and the discussion above. To prove (i), note that once we compute $\sigma_A(\mu)$, we can recursively compute $\sigma_A^j(\mu)$ for $j \geq 2$, using the relation $h(\mu) = 0$ to replace μ^b and higher powers of μ in terms of $1, \mu, \dots, \mu^{b-1}$.

By definition (4-3), we have $\mu = \prod_{0 \leq i < (q^d-1)/b} C_{A^{ib \bmod M}}(\lambda)$. Thus one can compute an expression

$$\mu = \sum_{i=0}^{q^d-2} e_i \lambda^i \in R_T[\lambda]$$

with coefficients $e_i \in R_T$ in $q^{O(d)}$ time. By successive multiplication in the ring $R_T[\lambda]$ (using the relation $C_M(\lambda) = 0$ to express λ^{q^d-1} and higher powers in terms of $1, \lambda, \dots, \lambda^{q^d-2}$), we can compute, for $l = 0, 1, \dots, b-1$, expressions

$$\mu^l = \sum_{i=0}^{q^d-2} e_{il} \lambda^i$$

with $e_{il} \in R_T$ in $q^{O(d)}$ time.

We have

$$\sigma_A(\mu) = \sum_{i=0}^{q^d-2} e_i \sigma_A(\lambda)^i = \sum_{i=0}^{q^d-2} e_i C_{A \bmod M}(\lambda)^i.$$

So one can likewise compute an expression $\sigma_A(\mu) = \sum_{i=0}^{q^d-2} f_i \lambda^i$ with $f_i \in R_T$ in $q^{O(d)}$ time. The task now is to rewrite this expression for $\sigma_A(\mu)$ as an element of $R_T[\mu]$, of the form $\sum_{l=0}^{b-1} a_l \mu^l$, for unknowns $a_l \in R_T$ that are to be determined. We will argue that this can be accomplished by solving a linear system.

Indeed, using the expressions $\mu^l = \sum_{i=0}^{q^d-2} e_{il} \lambda^i$, the coefficients a_l satisfy the following system of linear equations over R_T :

$$\sum_{l=0}^{b-1} e_{il} a_l = f_i \quad \text{for } i = 0, 1, \dots, q^d - 2. \quad (5-4)$$

Since the representation $\sigma_A(\mu) = \sum_{l=0}^{b-1} a_l \mu^l$ is unique, the system has a unique solution. By Cramér's rule, the degree of each a_l is at most $q^{O(d)}$. Therefore, we can express the system (5-4) as a linear system of size $q^{O(d)}$ over \mathbb{F}_q in unknowns the coefficients of all the polynomials $a_l \in R_T$. By solving this system in $q^{O(d)}$ time, we can compute the representation of $\sigma_A(\mu)$ as an element of $R_T[\mu]$. \square

5D. The basic cyclotomic algebraic-geometric code. The basic AG code \mathcal{C}^0 based on subfield E of the cyclotomic function field $F(\Lambda_M)$ is defined as

$$\mathcal{C}^0 = \left\{ (f(P_j^{(\beta)}))_{\beta \in F_r, 0 \leq j < b} \mid f \in \mathcal{L}(\ell M') \right\}, \quad (5-5)$$

where the ordering of the places $P_j^{(\beta)}$ above $T - \beta$ is as in (5-3). We record the standard parameters of the above AG code, which follows from Riemann–Roch, the genus of E from Proposition 4.2, and that a nonzero $f \in \mathcal{L}(\ell M')$ can have at most $\ell \deg M' = \ell d$ zeroes.

Lemma 5.5. *Suppose $\ell \geq b$. Then \mathcal{C}^0 is an \mathbb{F}_q -linear code of block length $n = rb$, dimension $k = \ell d - d(b-1)/2$, and distance at least $n - \ell d$.*

Lemma 5.4(ii) implies the following.

Lemma 5.6 (Efficient encoding). *Given a basis for the message space $\mathcal{L}(\ell M')$ represented in the form (5-1), the generator matrix of the cyclotomic code \mathcal{C}^0 can be computed in $\text{poly}(\ell, q^d, q^D)$ time.*

5E. The folded cyclotomic code. Let $m \geq 1$ be an integer. For convenience, we assume $m|b$ (though this is not really necessary). Analogously to the construction of folded RS codes [Guruswami and Rudra 2008], the folded cyclotomic code \mathcal{C} is obtained from \mathcal{C}^0 by bundling together successive m -tuples of symbols into a single symbol to give a code of length $N = n/m$ over \mathbb{F}_q^m . Formally,

$$\mathcal{C} = \left\{ \left(f(P_{mi}^{(\beta)}), f(P_{mi+1}^{(\beta)}), \dots, f(P_{mi+m-1}^{(\beta)}) \right)_{\beta \in \mathbb{F}_r, 0 \leq i < b/m} \mid f \in \mathcal{L}(\ell M') \right\}. \quad (5-6)$$

We will index the N positions of codewords in \mathcal{C} by pairs (β, ι) for $\beta \in \mathbb{F}_r$ and $\iota \in \{0, 1, \dots, (b/m) - 1\}$.

The generator matrix of unfolded code \mathcal{C}^0 , which can be computed given a basis for $\mathcal{L}(\ell M')$ according to Lemma 5.6, obviously suffices for encoding. Later we will argue that the *same representation* also suffices for polynomial time list decoding.

5F. Folding and Artin–Frobenius automorphism. The unique place A' that lies above A has degree $D' \stackrel{\text{def}}{=} Db$. The residue field at A' , denoted by $K_{A'}$, is isomorphic to $\mathbb{F}_{q^{D'}}$. By our choice $Db > \ell d$, this immediately implies that a message in $\mathcal{L}(\ell M')$ is uniquely determined by its evaluation at A' .

Lemma 5.7. *The map $\text{ev}_{A'} : \mathcal{L}(\ell M') \rightarrow K_{A'}$ given by $\text{ev}_{A'}(f) = f(A')$ is one-one.*

The key algebraic property of our folding is the following.

Lemma 5.8. *For every $f \in \mathcal{L}(\ell M')$:*

- (i) *For every $\beta \in \mathbb{F}_r$ and $0 \leq j < b-1$, $\sigma_A(f)(P_j^{(\beta)}) = f(P_{j+1}^{(\beta)})$.*
- (ii) *$\sigma_A(f)(A') = f(A')^{q^D}$.*

Proof. Part (i) follows since we ordered the places above $T - \beta$ such that $P_{j+1}^{(\beta)} = \sigma_A^{-1}(P_j^{(\beta)})$.

Part (ii) follows from the property of the Artin automorphism at A , since the norm of the place A equals $q^{\deg A} = q^D$. (A nice discussion of the Artin–Frobenius automorphism, albeit in the setting of number fields, appears in [Marcus 1977, Chapter 4].) □

6. List decoding algorithm

We now turn to list decoding the folded cyclotomic code \mathcal{C} defined in (5-6). The underlying approach is similar to that of the algorithm for list decoding folded RS codes [Guruswami and Rudra 2008] and AG generalizations of Parvaresh–Vardy codes [2005; Guruswami and Patthak 2008]. We will therefore not repeat the entire rationale and motivation behind the algorithm development. But our technical presentation and analysis is self-contained. In fact, our presentation here does offer some simplifications over previous descriptions of AG list decoding algorithms from [Guruswami and Sudan 1999; 2001; Guruswami and Patthak 2008]. A principal strength of the new description is that it *avoids the use of zero-increasing bases* at each code place $P_j^{(\beta)}$. This simplifies the algorithm as well as the representation of the code needed for decoding.

The list decoding problem for \mathcal{C} up to e errors corresponds to solving the following function reconstruction problem. Recall that the length of the code is $N = n/m = rb/m$, and the codeword positions are indexed by $\mathbb{F}_r \times \{0, 1, \dots, (b/m) - 1\}$.

Input: Collection \mathcal{T} of N tuples $(y_{mi}^{(\beta)}, y_{mi+1}^{(\beta)}, \dots, y_{mi+m-1}^{(\beta)}) \in \mathbb{F}_q^m$ for $\beta \in \mathbb{F}_r$ and $0 \leq i < b/m$.

Output: A list of all $f \in \mathcal{L}(\ell M')$ whose encoding according to \mathcal{C} agrees with the (β, i) -th tuple for at least $N - e$ codeword positions.

6A. Algorithm description. We describe the algorithm at a high level below and later justify how the individual steps can be implemented efficiently, and under what condition the decoding will succeed. We stress that regardless of complexity considerations, even the *combinatorial* list-decodability property “proved” by the algorithm is nontrivial.

Algorithm List-Decode(\mathcal{C})

Parameters:

- An integer parameter s , $2 \leq s \leq m$, for s -variate interpolation;
- an integer parameter $w \geq 1$ that governs the zero order (multiplicity) guaranteed by interpolation; and
- an integer parameter $\Delta \geq 1$ that is the total degree of the interpolated s -variate polynomial.

Step 1 (Interpolation): Find a nonzero polynomial $Q(Z_1, Z_2, \dots, Z_s)$ of total degree at most Δ with coefficients in $\mathcal{L}(\ell M')$ such that for each $\beta \in \mathbb{F}_r$, $0 \leq i < b/m$, and $j' \in \{0, 1, \dots, m - s\}$, the shifted polynomial

$$Q(Z_1 + y_{mi+j'}^{(\beta)}, Z_2 + y_{mi+j'+1}^{(\beta)}, \dots, Z_s + y_{mi+j'+s-1}^{(\beta)}) \quad (6-1)$$

has the property that the coefficient of the monomial $Z_1^{n_1} Z_2^{n_2} \dots Z_s^{n_s}$ vanishes at $P_{mi+j'}^{(\beta)}$ whenever its total degree $n_1 + n_2 + \dots + n_s < w$.

Step 2 (Root-finding): Find a list of all $f \in \mathcal{L}(\ell M')$ satisfying

$$Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f)) = 0.$$

Output those whose encoding according to \mathcal{C} agrees with at least $N - e$ of the m -tuples in \mathcal{T} .

6B. Analysis of error correction radius.

Lemma 6.1. *If $k(\Delta + 1)^s \geq N(m - s + 1)(w + s - 1)^s$ (where, as we recall, $k = \ell d - d(b - 1)/2$ is the dimension of $\mathcal{L}(\ell M')$), then a nonzero polynomial Q with the stated properties exists. If we know the evaluations of the functions in a basis $\{\phi_1, \phi_2, \dots, \phi_k\}$ of $\mathcal{L}(\ell M')$ at the places $P_j^{(\beta)}$, then such a Q can be found by solving a homogeneous system of linear equations over \mathbb{F}_q with at most $Nm(w + s)^s$ equations and unknowns.*

Proof. The proof is standard and follows by counting degrees of freedom versus number of constraints. One can express the desired polynomial as

$$\sum_{n_1, n_2, \dots, n_s} q_{(n_1, \dots, n_s)} Z_1^{n_1} \cdots Z_s^{n_s},$$

with unknowns $q_{(n_1, \dots, n_s)} \in \mathbb{F}_q$. The number of coefficients is $k \binom{\Delta + s}{s} > k(\Delta + 1)^s / s!$. One can express for each place $P_{mi+j'}^{(\beta)}$ the required condition at that place by $\binom{w+s-1}{s}$ linear conditions (this quantity is the number of monomials of total degree less than w), for a total of

$$N(m - s + 1) \binom{w + s - 1}{s} < N(m - s + 1) \frac{(w + s - 1)^s}{s!}$$

constraints. When the number of unknowns exceeds the number of constraints, a nonzero solution must exist. A solution can also be found efficiently once the linear system is set up, which can clearly be done if we know the evaluations of ϕ_i 's at the code places (that is, a *generator matrix* of the code). \square

Lemma 6.2. *Let Q be the polynomial found in [Step 1](#). If the encoding of some f as per \mathcal{C} agrees with $(y_{mi}^{(\beta)}, y_{mi+1}^{(\beta)}, \dots, y_{mi+m-1}^{(\beta)})$ for some position (β, ι) , then $Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f))$ has at least w zeroes at each of the $(m - s + 1)$ places $P_{mi+j'}^{(\beta)}$ for $j' = 0, 1, \dots, m - s$.*

Proof. The proof differs slightly from earlier proofs of similar statements (for example, [[Guruswami and Patthak 2008](#), Lemma 6.6]) in that it avoids the use of zero-increasing bases and is thus simpler. We will prove the claim for $j' = 0$, and the same proof works for any $j' \leq m - s$. Note that agreement on the m -tuple at position (b, ι) implies that

$$f(P_{mi}^{(\beta)}) = y_{mi}^{(\beta)}, f(P_{mi+1}^{(\beta)}) = y_{mi+1}^{(\beta)}, \dots, f(P_{mi+s-1}^{(\beta)}) = y_{mi+s-1}^{(\beta)}.$$

By Lemma 5.8(i), this implies

$$f(P_{m_i}^{(\beta)}) = y_{m_i}^{(\beta)}, \sigma_A(f)(P_{m_i}^{(\beta)}) = y_{m_i+1}^{(\beta)}, \dots, \sigma_{A^{s-1}}(f)(P_{m_i}^{(\beta)}) = y_{m_i+s-1}^{(\beta)}.$$

Denote by Q^* the shifted polynomial (6-1) for the triple $(\beta, \iota, 0)$. We have

$$\begin{aligned} Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f)) &= Q^*(f - y_{m_i}^{(\beta)}, \sigma_A(f) - y_{m_i+1}^{(\beta)}, \dots, \sigma_{A^{s-1}}(f) - y_{m_i+s-1}^{(\beta)}) \\ &= \sum_{\substack{n_1, n_2, \dots, n_s \\ w \leq n_1 + \dots + n_s \leq \Delta}} q_{(n_1, \dots, n_s)}^* (f - f(P_{m_i}^{(\beta)}))^{n_1} (\sigma_A(f) - \sigma_A(f)(P_{m_i}^{(\beta)}))^{n_2} \\ &\quad \dots (\sigma_{A^{s-1}}(f) - \sigma_{A^{s-1}}(f)(P_{m_i}^{(\beta)}))^{n_s}. \end{aligned}$$

for some coefficients $q_{(n_1, \dots, n_s)}^* \in \mathbb{F}_q$. Since each term of the function in the last expression has valuation at least w at $P_{m_i}^{(\beta)}$, so does $Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f))$. \square

Lemma 6.3. *If the encoding of $f \in \mathcal{L}(\ell M')$ has at least $N - e$ agreements with the input tuples \mathcal{T} , and $(N - e)(m - s + 1)w > d\ell(\Delta + 1)$, then*

$$Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f)) = 0.$$

Proof. Since f has no poles outside M' , neither do $\sigma_{A^i}(f)$ for $1 \leq i < s$. Moreover, $v_{M'}(\sigma_A(f)) = v_{\sigma_A^{-1}(M')}(f) = v_{M'}(f)$ (since M' is the unique place above M and is thus fixed by every Galois automorphism). Since $f \in \mathcal{L}(\ell M')$, this implies $\sigma_{A^i}(f) \in \mathcal{L}(\ell M')$ for every i . Since each coefficient of Q also belongs to $\mathcal{L}(\ell M')$, we conclude that $Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f)) \in \mathcal{L}((\ell + \ell\Delta)M')$. On the other hand, by Lemma 6.2, $Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f))$ has at least $(N - e)(m - s + 1)w$ zeroes. If $(N - e)(m - s + 1)w > \ell(\Delta + 1)d$, then $Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f))$ has more zeroes than poles and must thus equal 0. \square

Putting together the above lemmas, we can conclude the following about the list decoding radius guaranteed by the algorithm. Note that we have not yet discussed how Step 2 may be implemented, or why it implies a reasonable bound on the output list size. We will do this in Section 6C.

Theorem 6.4. *For every $s, 2 \leq s \leq m$, and any $\zeta > 0$, for the choice $w = \lceil s/\zeta \rceil$ and a suitable choice of the parameter Δ , the algorithm List-Decode(\mathcal{C}) successfully list decodes up to e errors whenever*

$$e < (N - 1) - (1 + \zeta) \left(\frac{k}{m - s + 1} \right)^{1-1/s} N^{1/s} \left(1 + \frac{d(b-1)}{2k} \right). \tag{6-2}$$

Proof. Picking $w = \lceil s/\zeta \rceil$ and

$$\Delta + 1 = \left\lceil \left(\frac{N(m-s+1)}{k} \right)^{1/s} (w + s - 1) \right\rceil,$$

the requirement of Lemma 6.1 is met. By Lemma 5.5, the dimension k satisfies $\ell d = k + d(b - 1)/2$. A straightforward computation reveals that for this choice,

the bound (6-2) implies the decoding condition $(N - e)(m - s + 1)w > \ell d(\Delta + 1)$, under which Lemma 6.3 guarantees successful decoding. \square

Remark 6.5. The error correction radius above is nontrivial only when $s \geq 2$. We will see later how to pick parameters so that the error fraction approaches $1 - R^{1-1/s}$. For AG codes, even $s = 1$ led to a nontrivial guarantee of about $1 - \sqrt{R}$ in [Guruswami and Sudan 1999], and for folded RS codes the error fraction with s -variate interpolation was $1 - R^{s/(s+1)}$. The weaker bound we get is due to restricting the pole order of coefficients of Q to at most ℓ , the number of poles allowed for messages. This is similar to the algorithm in [Guruswami and Patthak 2008, Section 5]. Since we let grow s anyway, this does not hurt us. It also avoids some difficult technical complications that would arise otherwise (discussed in, for example, [Guruswami and Patthak 2008]), and allows us to implement the interpolation step just using the natural generator matrix of the code.

6C. Root-finding using the Artin automorphism. So far we have not discussed how Step 2 of decoding can be performed, and why in particular it implies a reasonably small upper bound on the number of solutions $f \in \mathcal{L}(\ell M')$ that it may find in the worst case. We address this now. This is where the properties of the Artin automorphism σ_A will play a crucial role. Recall that $K_{A'} = \mathbb{C}_{A'}/A'$ denotes the residue field at the place A' of E lying above A , and that we picked A so that $D = \deg A$ obeyed $Db > \ell d$.

Lemma 6.6. *Suppose $f \in \mathbb{C}_{A'}$ satisfies*

$$Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f)) = 0$$

for some $Q \in \mathbb{C}_{A'}[Z_1, Z_2, \dots, Z_s]$. Let $\bar{Q} \in K_{A'}[Z_1, Z_2, \dots, Z_s]$ be the polynomial obtained by reducing the coefficients of Q modulo A' . Then $f(A') \in K_{A'}$ obeys

$$\bar{Q}(f(A'), f(A')^{q^D}, f(A')^{q^{2D}}, \dots, f(A')^{q^{D(s-1)}}) = 0. \quad (6-3)$$

Proof. If $Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f)) = 0$, then surely

$$\bar{Q}(f(A'), \sigma_A(f)(A'), \dots, \sigma_{A^{s-1}}(f)(A')) = 0.$$

The claim (6-3) now follows immediately from Lemma 5.8(ii). \square

Lemma 6.7. *If $Q(Z_1, \dots, Z_s)$ is a nonzero polynomial of total degree at most $\Delta < q^D$ all of whose coefficients belong to $\mathcal{L}(\ell M')$, then the polynomial $\Phi \in K_{A'}[Y]$ defined as*

$$\Phi(Y) \stackrel{\text{def}}{=} \bar{Q}(Y, Y^{q^D}, \dots, Y^{q^{D(s-1)}})$$

is a nonzero polynomial of degree at most $\Delta \cdot q^{D(s-1)}$.

Proof. If $\psi \in \mathcal{L}(\ell M')$ is nonzero, then $\psi(A') \neq 0$. (Otherwise, the degree of the zero divisor of ψ will be at least $\deg A' = bD > \ell d$, and thus exceed the degree of the pole divisor of ψ .) It follows that if $Q \neq 0$, then $\bar{Q}(Z_1, \dots, Z_s)$ obtained by reducing coefficients of Q modulo A' is also nonzero.⁵ Since the degree of \bar{Q} in each Z_i is at most $\Delta < q^D$, it is easy to see that $\Phi(Y) = \bar{Q}(Y, Y^{q^D}, \dots, Y^{q^{D(s-1)}})$ is also nonzero. The degree of Φ is at most $q^{D(s-1)}$ times the total degree of \bar{Q} , which is at most Δ . \square

By the above two lemmas, we see that one can compute the set of residues $f(A')$ of all f satisfying $Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f)) = 0$ by computing the roots in $K_{A'}$ of $\Phi(Y)$. Since $\text{ev}_{A'}$ is injective on $\mathcal{L}(\ell M')$ (Lemma 5.7), this also lets us recover the message $f \in \mathcal{L}(\ell M')$.

Lemma 6.8. *Given a nonzero polynomial $Q(Z_1, \dots, Z_s)$ with coefficients from $\mathcal{L}(\ell M')$ and degree $\Delta < q^D$, the set of functions*

$$\mathcal{S} = \{f \in \mathcal{L}(\ell M') \mid Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f)) = 0\}$$

has cardinality at most q^{Ds} .

Moreover, knowing the evaluations of a basis $\mathcal{B} = \{\phi_1, \phi_2, \dots, \phi_k\}$ of $\mathcal{L}(\ell M')$ at the place A' , one can compute the coefficients expressing each $f \in \mathcal{S}$ in the basis \mathcal{B} in $q^{O(Ds)}$ time.

Proof. As argued above, any desired $f \in \mathcal{L}(\ell M')$ has the property that $\Phi(f(A')) = 0$, so the evaluations of functions in \mathcal{S} take at most $\text{degree}(\Phi) \leq \Delta q^{D(s-1)} \leq q^{Ds}$ values. Since $\text{ev}_{A'}$ is injective on \mathcal{S} , this implies $|\mathcal{S}| \leq q^{Ds}$. The second part follows since we can compute the roots of Φ in $K_{A'}$ in time $\text{poly}(q^{Ds}, \log |K_{A'}|) \leq q^{O(Ds)}$. Knowing $f(A')$, we can recover f (in terms of the basis \mathcal{B}) by solving a linear system if we know the evaluations of the functions in the basis \mathcal{B} at A' . The next section discusses a convenient representation for computations in $K_{A'}$. \square

6C.1. Representation of the residue field $K_{A'}$. The following gives a convenient representation for elements of $K_{A'}$ which can be used in computations involving this field.

Lemma 6.9. *The elements $\{1, \mu(A), \dots, \mu(A)^{b-1}\}$ form a basis for $K_{A'}$ over the field $R_T/(A) \simeq \mathbb{F}_{q^D}$. In other words, elements of $K_{A'}$ can be expressed in a unique way as*

$$\sum_{i=0}^{b-1} b_i(T) \mu(A)^i,$$

where each $b_i \in R_T$ has degree less than D .

⁵This is simplicity we gain by restricting the coefficients of Q to also belong to $\mathcal{L}(\ell M')$.

Proof. Since A is inert in E/F , the minimal polynomial $h(Z)$ of μ over F has the property that $\bar{h}(Z)$, obtained by reducing the coefficients of h modulo A , is irreducible over the residue field $R_T/(A)$. Thus $\mu(A)$ generates $K_{A'}$ over $R_T/(A)$, and in fact the minimal polynomial of $\mu(A)$ with respect to K_A equals $\bar{h}(Z)$. Note that the coefficients of \bar{h} , which belong to $R_T/(A)$, have a natural representation as a polynomial in R_T of degree less than $\deg A = D$. \square

We note that given the representation of the basis $\mathcal{B} = \{\phi_1, \phi_2, \dots, \phi_k\}$ in the form guaranteed by [Theorem 5.2](#), one can trivially compute the evaluations of $\phi_i(A')$ in the above form. There is no need to explicitly compute $\mu(A) \in \mathbb{C}_A/A$. Therefore, the decoding algorithm requires no additional preprocessed information beyond a basis for the message space $\mathcal{L}(\ell M')$ —the rest can all be computed efficiently from the basis alone.

6D. Wrap-up. We are now ready to state our final decoding claim.

Theorem 6.10. *For any s , $2 \leq s \leq m$, and $\zeta > 0$, the folded cyclotomic code $\mathcal{C} \subseteq (\mathbb{F}_q^m)^N$ defined in (5-6) can be list decoded in time $(Nm)^{O(1)}(s/\zeta)^{O(s)} + q^{O(Ds)}$ from a fraction ρ of errors*

$$\rho = 1 - (1 + \zeta) \left(\frac{R_0 m}{m - s + 1} \right)^{1-1/s} \left(1 + \frac{d}{2R_0 r} \right), \quad (6-4)$$

where $R_0 = k/n$ is the rate of the code. The size of the output list is at most q^{Ds} . The decoding algorithm assumes polynomial amount of preprocessed information consisting of basis functions $\{\phi_1, \dots, \phi_k\}$ for the message space $\mathcal{L}(\ell M')$ represented in the form (5-1). (This is the same representation used for encoding, and it is succinct by [Lemma 5.3](#).)

Proof. We first note that bound on fraction of errors follows from [Theorem 6.4](#), and the fact that $k = R_0 n = R_0 N m = R_0 b r$. By [Lemma 6.1](#) and its proof, in [Step 1](#) of the algorithm we can find a nonzero polynomial Q (of degree less than q^D) such that for any $f \in \mathcal{L}(\ell M')$ that needs to be output by the list decoder, we must have $Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f)) = 0$. We can evaluate the basis functions ϕ_i at $P_j^{(\beta)}$ in $(\ell q^d)^{O(1)}$ time by [Lemma 5.4](#), and with this information, the running time of this interpolation step can be bounded by $(Nm)^{O(1)}(w + s)^{O(s)} = (Nm)^{O(1)}(s/\zeta)^{O(s)}$ (since $w = O(s/\zeta)$). We can also efficiently compute the evaluations of ϕ_i at A' in the representation suggested by [Lemma 6.9](#). Therefore, by [Lemma 6.8](#), we can then find a list of the at most q^{Ds} functions f satisfying $Q(f, \sigma_A(f), \dots, \sigma_{A^{s-1}}(f)) = 0$ in $q^{O(Ds)}$ time. \square

Remark 6.11 (List recovery). A similar claim holds for the more general *list recovery* problem, where for each position we are given as input a set of up to l elements of \mathbb{F}_q^m , and the goal is to find all codewords which agree with some element of the input sets for at least a fraction $(1 - \rho)$ of positions. In this case, $1 - \rho$ only needs

to be only a factor $l^{1/s}$ larger than the bound (6-4). By picking $s \gg l$, the effect of l can be made negligible. This feature is very useful in concatenation schemes; see Section 7A and [Guruswami and Rudra 2008] for further details.

7. Long codes with optimal rate for list decoding

We now describe the parameter choices which lead to capacity-achieving list-decodable codes, that is, codes of rate R_0 that can correct a fraction $1 - R_0 - \varepsilon$ of errors (for any desired $0 < R_0 < 1$), and whose alphabet size is polylogarithmic in the block length; the formal statement appears in Theorem 7.1 below. (Recall that for folded RS codes, the alphabet size is a large polynomial in the block length.) Using concatenation and expander-based ideas, [Guruswami and Rudra 2008] also presents capacity-achieving codes over a fixed alphabet size (that depends on the distance ε to capacity alone). The advantage of our codes is that they inherit strong list recovery properties similar to the folded RS codes (Remark 6.11). This is very useful in concatenation schemes, and indeed our codes can be used as outer codes for an explicit family of binary concatenated codes list-decodable up to the Zyablov radius, *with no brute-force search* for the inner code (see Section 7A below).

We now describe our main result on how to obtain the desired codes from the construction \mathcal{C} and Theorem 6.10. The underlying parameter choices to achieve this require a fair bit of care.

Theorem 7.1 (Main theorem). *For every R_0 , $0 < R_0 < 1$, and every constant $\varepsilon > 0$, the following holds for infinitely many integers \mathbf{q} which are powers of two. There is a code of rate at least R_0 over an alphabet of size \mathbf{q} with block length*

$$N \geq 2^{\mathbf{q}^{\Omega(\varepsilon^2 / \log(1/R_0))}}$$

that can be list decoded up to a fraction $1 - R_0 - \varepsilon$ of errors in time bounded by $(N \log(1/R_0) / \varepsilon^2)^{O(1/(R_0\varepsilon)^2)}$.

Proof. Suppose R_0 , $0 < R_0 < 1$, and $\varepsilon > 0$ are given. Let $c = 2 \lfloor \frac{10}{R_0\varepsilon} \rfloor + 1$, and let $\phi(c)$ denote the Euler's totient function of c .

Let $u \geq 1$ be an arbitrary integer; we will get a family of codes by varying u . The code we construct will be a folded cyclotomic code \mathcal{C} defined in (5-6). Let $x = \phi(c)u$. Note that $2^x \equiv 1 \pmod{c}$. We first pick q, r, d as follows: $r = 2^x$, $q = r^2$, and $d = (2^x - 1)/c$. For this choice, $d|r - 1$ and $(q - 1)/(r - 1) = r + 1$ is coprime to d , as required in Lemma 4.1. So we can take $M(T) = T^d - \gamma \in \mathbb{F}_r[T]$ for γ primitive in \mathbb{F}_r as the irreducible polynomial over \mathbb{F}_q .

For the above choice, $d/r < 1/c \leq \varepsilon R_0/20$, so that $d/2R_0r < \varepsilon/10$. By picking

$$s = \Theta(\varepsilon^{-1} \log(1/R_0)), \quad m = \Theta(s/\varepsilon), \quad \text{and} \quad \zeta = \varepsilon/20,$$

we can ensure that the decoding radius ρ guaranteed in (6-4) by Theorem 6.10 is at least $1 - (1 + \varepsilon)R_0$.

The degree b of the extension E/F , introduced in (4-1), is given by

$$b = (r^d + 1)/(r + 1).$$

The length of the unfolded cyclotomic code \mathcal{C}^0 (defined in (5-5)) equals $n = rb > r^d/2$. We need to ensure that the rate of \mathcal{C}^0 , which is equal to the rate of the folded cyclotomic code \mathcal{C} , is at least R_0 . To this end, we will pick

$$\ell = \left\lceil \frac{b}{2} + \frac{R_0 r b}{d} \right\rceil. \quad (7-1)$$

It is easily checked that for our choice of parameters, $\ell \geq b$. By Lemma 5.5, the rate of \mathcal{C}^0 equals $d(\ell - (b - 1)/2)/rb$, which is at least R_0 for the above choice of ℓ .

We next pick the value of D , the degree of the irreducible A , which is the key quantity governing the list size and decoding complexity. To satisfy the condition (5-2), we need $D > \max\{\ell d/b, 3d\}$. For the ℓ chosen above, this condition is surely met if $D > 3r$. We can thus pick

$$D = \Theta(r) = \Theta(dc) = \Theta(d/(R_0\varepsilon)) .$$

The running time of the list decoding algorithm is dominated by the $q^{O(Ds)}$ term, and for the above choice of parameters can be bounded by $q^{O(d/(R_0\varepsilon)^2)}$. The block length of the code N satisfies

$$N = \frac{n}{m} > \frac{r^d}{2m} = \frac{q^{d/2}}{2m} = \Omega\left(\frac{\varepsilon^2 q^{d/2}}{\log(1/R_0)}\right).$$

As a function of N , the decoding complexity is therefore bounded by

$$(N \log(1/R_0)/\varepsilon^2)^{O(1/(R_0\varepsilon)^2)}.$$

The alphabet size of the folded cyclotomic code is $\mathbf{q} = q^m$, and we can bound the block length N from below as a function of \mathbf{q} as:

$$\begin{aligned} N &\geq \frac{q^{d/2}}{2m} \geq \frac{q^{\Omega(r/c)}}{2m} \geq \frac{q^{\Omega(\varepsilon R_0 \sqrt{q})}}{2m} \\ &\geq 2^{\sqrt{q}} \quad (\text{for large enough } q \text{ compared to } 1/R_0, 1/\varepsilon) \\ &= 2^{\mathbf{q}^{1/(2m)}} \geq 2^{\mathbf{q}^{\Omega(\varepsilon^2/\log(1/R_0))}} . \end{aligned}$$

This establishes the claimed lower bound on block length, and completes the proof of the theorem. \square

7A. Concatenated codes list-decodable up to Zyablov radius. Using the strong list recovery property of folded RS codes, a polynomial time construction of binary codes list-decodable up to the Zyablov radius was given in [Guruswami and Rudra 2008, Theorem 5.3]. The construction used folded RS codes as outer codes in a concatenation scheme, and involved an undesirable brute-force search to find a binary inner code that achieves list decoding capacity. The time to construct the code grew faster than $N^{\Omega(1/\varepsilon)}$, where ε is the distance of the decoding radius to the Zyablov radius. This result as well as our result below hold not only for binary codes but also codes over any fixed alphabet; for sake of clarity, we state results only for binary codes.

As the folded cyclotomic codes from Theorem 7.1 are much longer than the alphabet size, by using them as outer codes, it is possible to achieve a similar result without having to search for an inner code, by using as inner codes *all possible binary linear codes* of a certain rate!

Theorem 7.2. *Let $0 < R_0, r < 1$ and $\varepsilon > 0$. Let \mathcal{C} be a folded cyclotomic code guaranteed by Theorem 7.1 with rate at least R_0 and a large enough block length N . Let \mathcal{C}^* be a binary code obtained by concatenating \mathcal{C} with all possible binary linear maps of rate r (each one used a roughly equal number of times). Then \mathcal{C}^* is a binary linear code of rate at least $R_0 \cdot r$ that can be list decoded from a fraction $(1 - R_0)H^{-1}(1 - r) - \varepsilon$ of errors in $N^{(1/\varepsilon)^{O(1)}}$ time.*

We briefly discuss the idea behind proving the above claim. As the alphabet size of folded cyclotomic codes is polylogarithmic in N , each outer codeword symbol can be expressed using $O_\varepsilon(\log \log N)$ bits. Hence the total number of such inner codes S will be at most $2^{O_\varepsilon((\log \log N)^2)} \ll N$ for large enough N . The N outer codeword positions will be partitioned into S (roughly) equal parts in an arbitrary way, and each inner code used to encode all the outer codeword symbols in one of the parts. Most of the inner codes achieve list decoding capacity — if their rate is r , they can list decode $H^{-1}(1 - r) - \varepsilon$ fraction of errors with constant sized lists (of size $2^{O(1/\varepsilon)}$). This suffices for analyzing the standard algorithm for decoding concatenated codes (namely, list decode the inner codes to produce a small set of candidate symbols for each position, and then list recover the outer code based on these sets). Arguing as in [Guruswami and Rudra 2008, Theorem 5.3], we can thus prove Theorem 7.2.

Appendix: List of parameters

Since the construction of the cyclotomic function field and the associated error-correcting code used a large number of parameters, we summarize them below for easy reference.

We begin by recalling the parameters concerning the function field construction.

q	size of the ground finite field
r	size of the subfield $\mathbb{F}_r \subset \mathbb{F}_q$
F	the field $\mathbb{F}_q(T)$ of rational functions
R_T	the ring of polynomials $\mathbb{F}_q[T]$
P_∞	the place of F that is the unique pole of T
M	polynomial $T^d - \gamma \in \mathbb{F}_r[T]$, irreducible over \mathbb{F}_q
d	degree of the irreducible polynomial M
C_M	the Carlitz action corresponding to M
Λ_M	the M -torsion points in F^{ac} under the action C_M
K	the cyclotomic function field $F(\Lambda_M)$
λ	nonzero element of Λ_M that generates K over F ; $K = F(\lambda)$
G	the Galois group of K/F , naturally isomorphic to $(R_T/(M))^*$
H	the subgroup $\mathbb{F}_q^* \cdot \mathbb{F}_r[T]$ of G
E	the fixed field K^H of H
μ	primitive element for E/F ; $E = F(\mu)$
b	the degree $[E : F]$ of the extension E/F
g	the genus of E/F , equals $d(b-1)/2 + 1$

The construction of the code \mathcal{C}^0 from `lrefeqbasic-cycl` and its folded version \mathcal{C} from `lrefeqcode-def` used further parameters, listed here:

M'	the unique place of E lying above M
ℓ	maximum pole order at M' of message functions; $\ell \geq b$
$\mathcal{L}(\ell M')$	\mathbb{F}_q -linear space of messages of the codes
n	block length of \mathcal{C}^0 , $n = br$
k	dimension of the \mathbb{F}_q -linear code \mathcal{C} , $k = \ell d - g + 1$
m	folding parameter
N	block length of folded code \mathcal{C} , $N = n/m$
$P_j^{(\beta)}$	the rational places lying above $T - \beta$ in E , for $\beta \in \mathbb{F}_r$ and $0 \leq j < b$
A	an irreducible polynomial (place of F) that remains inert in E/F
D	the degree of the polynomial A ; satisfies $Db > \ell d$
σ_A	the Artin automorphism of the extension E/F at A
A'	the unique place of E lying above A

Acknowledgments

Much of this work was carried out when I was a member in the School of Mathematics, Institute for Advanced Study, during 2007-08. I thank the IAS for providing an inspiring work environment and its wonderful hospitality.

- [Marcus 1977] D. A. Marcus, *Number fields*, Springer, New York, 1977. [MR 56 #15601](#) [Zbl 0383.12001](#)
- [Niederreiter and Xing 1996] H. Niederreiter and C. Xing, “Explicit global function fields over the binary field with many rational places”, *Acta Arith.* **75**:4 (1996), 383–396. [MR 97d:11177](#) [Zbl 0877.11065](#)
- [Niederreiter and Xing 1997] H. Niederreiter and C. Xing, “Cyclotomic function fields, Hilbert class fields, and global function fields with many rational places”, *Acta Arith.* **79**:1 (1997), 59–76. [MR 97m:11141](#) [Zbl 0891.11057](#)
- [Parvaresh and Vardy 2005] F. Parvaresh and A. Vardy, “Correcting errors beyond the Guruswami–Sudan radius in polynomial time”, pp. 285–294 in *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science* (Pittsburgh, PA, 2005), IEEE, 2005.
- [Quebbemann 1988] H.-G. Quebbemann, “Cyclotomic Goppa codes”, *IEEE Trans. Inform. Theory* **34**:5, part 2 (1988), 1317–1320. [MR 90b:11135](#) [Zbl 0665.94014](#)
- [Rosen 2002] M. Rosen, *Number theory in function fields*, Graduate Texts in Mathematics **210**, Springer, New York, 2002. [MR 2003d:11171](#) [Zbl 1043.11079](#)
- [Shen 1993] B.-Z. Shen, “A Justesen construction of binary concatenated codes that asymptotically meet the Zyablov bound for low rate”, *IEEE Trans. Inform. Theory* **39**:1 (1993), 239–242. [MR 93m:94046](#) [Zbl 0766.94022](#)
- [Stichtenoth 1993] H. Stichtenoth, *Algebraic function fields and codes*, Springer, Berlin, 1993. [MR 94k:14016](#) [Zbl 0816.14011](#)
- [Stichtenoth 2006] H. Stichtenoth, “Transitive and self-dual codes attaining the Tsfasman–Vlăduț–Zink bound”, *IEEE Trans. Inform. Theory* **52**:5 (2006), 2218–2224. [MR 2006m:94126](#)
- [Sudan 1997] M. Sudan, “Decoding of Reed Solomon codes beyond the error-correction bound”, *J. Complexity* **13**:1 (1997), 180–193. [MR 98f:94024](#) [Zbl 0872.68026](#)
- [Villa Salvador 2006] G. D. Villa Salvador, *Topics in the theory of algebraic function fields*, Birkhäuser, Boston, 2006. [MR 2007i:11002](#) [Zbl 1154.11001](#)

Communicated by Hendrik W. Lenstra

Received 2009-06-29

Revised 2010-01-05

Accepted 2010-02-17

guruswami@cmu.edu

Computer Science Department, Carnegie Mellon University,
Pittsburgh, PA 15213, United States
<http://www.cs.cmu.edu/~venkatg>

Algebra & Number Theory

www.jant.org

EDITORS

MANAGING EDITOR

Bjorn Poonen
Massachusetts Institute of Technology
Cambridge, USA

EDITORIAL BOARD CHAIR

David Eisenbud
University of California
Berkeley, USA

BOARD OF EDITORS

Georgia Benkart	University of Wisconsin, Madison, USA	Susan Montgomery	University of Southern California, USA
Dave Benson	University of Aberdeen, Scotland	Shigefumi Mori	RIMS, Kyoto University, Japan
Richard E. Borcherds	University of California, Berkeley, USA	Andrei Okounkov	Princeton University, USA
John H. Coates	University of Cambridge, UK	Raman Parimala	Emory University, USA
J-L. Colliot-Thélène	CNRS, Université Paris-Sud, France	Victor Reiner	University of Minnesota, USA
Brian D. Conrad	University of Michigan, USA	Karl Rubin	University of California, Irvine, USA
Hélène Esnault	Universität Duisburg-Essen, Germany	Peter Sarnak	Princeton University, USA
Hubert Flenner	Ruhr-Universität, Germany	Michael Singer	North Carolina State University, USA
Edward Frenkel	University of California, Berkeley, USA	Ronald Solomon	Ohio State University, USA
Andrew Granville	Université de Montréal, Canada	Vasudevan Srinivas	Tata Inst. of Fund. Research, India
Joseph Gubeladze	San Francisco State University, USA	J. Toby Stafford	University of Michigan, USA
Ehud Hrushovski	Hebrew University, Israel	Bernd Sturmfels	University of California, Berkeley, USA
Craig Huneke	University of Kansas, USA	Richard Taylor	Harvard University, USA
Mikhail Kapranov	Yale University, USA	Ravi Vakil	Stanford University, USA
Yujiro Kawamata	University of Tokyo, Japan	Michel van den Bergh	Hasselt University, Belgium
János Kollár	Princeton University, USA	Marie-France Vignéras	Université Paris VII, France
Hendrik W. Lenstra	Universiteit Leiden, The Netherlands	Kei-Ichi Watanabe	Nihon University, Japan
Yuri Manin	Northwestern University, USA	Andrei Zelevinsky	Northeastern University, USA
Barry Mazur	Harvard University, USA	Efim Zelmanov	University of California, San Diego, USA

PRODUCTION

ant@mathscipub.org

Paulo Ney de Souza, Production Manager

Silvio Levy, Senior Production Editor


See inside back cover or www.jant.org for submission instructions.

The subscription price for 2010 is US \$140/year for the electronic version, and \$200/year (+\$30 shipping outside the US) for print and electronic. Subscriptions, requests for back issues from the last three years and changes of subscribers address should be sent to Mathematical Sciences Publishers, Department of Mathematics, University of California, Berkeley, CA 94720-3840, USA.

Algebra & Number Theory (ISSN 1937-0652) at Mathematical Sciences Publishers, Department of Mathematics, University of California, Berkeley, CA 94720-3840 is published continuously online. Periodical rate postage paid at Berkeley, CA 94704, and additional mailing offices.

ANT peer-review and production is managed by EditFLOW™ from Mathematical Sciences Publishers.

PUBLISHED BY

 **mathematical sciences publishers**
<http://www.mathscipub.org>

A NON-PROFIT CORPORATION

Typeset in L^AT_EX

Copyright ©2010 by Mathematical Sciences Publishers

