

Article

Repeated-Root Constacyclic Codes of Length np^s from Half-Cyclotomic Defining Sets

Alhanouf Ali Alhomaïdhi *  and Sami H. Saif * 

Department of Mathematics, College of Science, King Saud University, P.O. Box 2455, Riyadh 11451, Saudi Arabia

* Correspondence: aalhomaïdhi@ksu.edu.sa (A.A.A.); ssiaf1@ksu.edu.sa (S.H.S.)

Abstract

Repeated-root constacyclic codes form an important class of algebraic error-correcting codes, combining the cyclic symmetry of constacyclic codes with the additional multiplicity structure produced by repeated roots. These multiplicities are usually a source of technical difficulty in distance analysis; in this paper, we use them instead as controllable design parameters. Let $q = p^m$, let n be prime to p , and let $\lambda \in \mathbb{F}_q^*$. We study repeated-root λ -constacyclic codes of length np^s obtained by lifting the half-cyclotomic defining sets. Writing $\lambda = \lambda_0^{p^s}$ gives $x^{np^s} - \lambda = (x^n - \lambda_0)^{p^s}$; hence every code is determined by a multiplicity profile $e : \Gamma \rightarrow 0, 1, \dots, p^s$ on the set Γ of q -cyclotomic coset leaders in $\mathbb{Z}_{n,r}$, where $r = \text{ord}(\lambda_0)$. Equivalently, e determines a nested sequence of simple-root defining sets $T_{p^s-1} \subseteq \dots \subseteq T_0 \subseteq \mathbb{Z}_{n,r}$. The main result is a layerwise distance formula: if C_t is the simple-root code defined by T_t and $P_t = \prod_j (t_j + 1)$ for the p -adic expansion $t = \sum_j t_j p^j$, then $d_H(C(e)) = \min_{0 \leq t < p^s} P_t d_H(C_t)$. This reduces the repeated-root distance problem to certified distance information for the simple-root layers. Combining the formula with the half-cyclotomic bounds in the two-coset-size case gives explicit certified lower bounds for complete families of repeated-root codes. We also formulate a verified refined layer bound, prove a conditional half-rate family theorem with distance growth at least a constant multiple of $N / \log_q N$ for $N = np^s$, and give a lifting criterion for transferring certified simple-root parameters to the repeated-root setting. The construction is carried out for the length families $n = (q^p - 1)/(r\sigma)$, $n = \Phi_{p^b}(q)$, and $n = \Phi_{p_1 p_2}(q)$.

Keywords: constacyclic codes; repeated-root codes; cyclotomic cosets; Hamming distance; half-cyclotomic defining sets; optimal codes; best-known codes

MSC: 94B15; 94B05; 11T71; 12E20



Received: 16 June 2026

Revised: 6 July 2026

Accepted: 9 July 2026

Published: 10 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Constacyclic codes form a central class of algebraic linear codes over finite fields. They contain cyclic and negacyclic codes as the special cases $\lambda = 1$ and $\lambda = -1$, respectively. Let q be a prime power and let $\mathbb{F}_q$ be the finite field with q elements. A linear code $\mathcal{C} \subseteq \mathbb{F}_q^N$ is called λ -constacyclic, where $\lambda \in \mathbb{F}_q^*$, if $(c_0, c_1, \dots, c_{N-1}) \in \mathcal{C}$ implies $(\lambda c_{N-1}, c_0, \dots, c_{N-2}) \in \mathcal{C}$. Under the usual polynomial representation, λ -constacyclic codes of length N over $\mathbb{F}_q$ are precisely the ideals of $\mathbb{F}_q[x]/\langle x^N - \lambda \rangle$. Thus each such code is generated by a divisor of $x^N - \lambda$ when the ambient ring is principal. This ideal-theoretic description makes constacyclic codes a tractable and useful framework for studying generator polynomials, defining sets, duality, and distance bounds [1–3].

Constacyclic and related structured codes have been studied from several perspectives. They have been used to construct new linear codes [4], locally repairable codes [5], quantum MDS codes [6], and MDS or pseudo-cyclic families [7–9]. Algorithmic equivalence questions were considered in [10], while quasi-cyclic and quasi-twisted constructions have also produced many good linear codes [11]. Recent related works include five infinite families of binary cyclic codes with parameters $[n, k]$, where $(n - 6)/3 \leq k \leq 2(n + 6)/3$, together with binary duadic and self-dual codes having square-root-like distance lower bounds [12]; two classes of q -ary constacyclic codes with variable parameters $[(q^m - 1)/r, k, d]$ [13]; self-dual negacyclic codes with variable lengths and square-root-like lower bounds [14]; 2^s -ary generalizations of the Tang–Ding cyclic and duadic constructions [15]; two near-half-rate ternary cyclic families of length $3^m - 1$ with $d \geq 6(3^{(m-1)/2} - 1) + 1$ and $d \geq 3^{(m-1)/2} - 1$, together with binary cyclic codes $[2^m - 1, 2^{m-1}, d \geq 7 \cdot 2^{(m-3)/2} + 1]_2$ [16]; and sharper minimum-distance lower bounds for binary BCH codes and their ternary analogues [17]. Grassl’s database remains a standard reference for comparing the parameters of linear and quantum codes [18].

When $\gcd(N, q) = 1$, the polynomial $x^N - \lambda$ is square-free, and the corresponding codes are called simple-root constacyclic codes. In this case, defining sets give a direct description of generator polynomials, dimensions, duals, and BCH-type lower bounds. If λ has multiplicative order r , then the defining positions of a simple-root λ -constacyclic code of length n lie in $Z_{n,r} = \{1 + ir : 0 \leq i \leq n - 1\} \subseteq \mathbb{Z}/(nr)\mathbb{Z}$. The defining set is a union of q -cyclotomic cosets in $Z_{n,r}$, and the arithmetic of these cosets is essential in the construction and analysis of simple-root cyclic, negacyclic, and constacyclic codes [19,20].

A particularly important direction concerns cyclic and constacyclic codes whose dimensions are close to one half of the length and whose minimum distances exceed square-root-type lower ranges. This direction is related to the long-standing question of whether cyclic codes can form an asymptotically good family [21]. Tang and Ding constructed binary cyclic codes with parameters $[n, (n + 1)/2]$ and good minimum distances [22]. Later developments produced further cyclic and constacyclic families with strong lower bounds and good parameters [23–25]. The simple-root construction of [26], based on selecting approximately half of the q -cyclotomic cosets in suitable coset-size classes, is the main simple-root input for the present paper. We use this input only through rigorously established simple-root distance certificates; heuristic or experimentally suggested distances are not treated as certified bounds unless they are independently verified.

The repeated-root case is more delicate. If the characteristic p divides the length, then $x^N - \lambda$ has repeated irreducible factors, and the minimum distance is no longer controlled by a single ordinary simple-root defining set. The fundamental distance reduction principle for repeated-root cyclic codes was established in [27]. Repeated-root cyclic and constacyclic codes with optimal or best-known parameters were recently studied in [28,29]. These works show that repeated-root structures can produce strong parameters, provided that the multiplicities of the irreducible factors are handled carefully.

The precise novelty of this paper is to make these multiplicities part of the design, rather than treating them only as a complication of the repeated-root ambient ring. The aim of this paper is to combine the half-cyclotomic simple-root philosophy with repeated-root multiplicity design. Throughout the paper, $q = p^m$, $\gcd(n, p) = 1$, and $\lambda = \lambda_0^{p^s}$. Hence, $x^{np^s} - \lambda = (x^n - \lambda_0)^{p^s}$. Let Γ be the set of q -cyclotomic coset leaders in $Z_{n,r}$, where $r = \text{ord}(\lambda_0)$. A repeated-root λ -constacyclic code of length np^s is determined by a multiplicity profile $e : \Gamma \rightarrow \{0, 1, \dots, p^s\}$. Equivalently, such a profile determines a nested chain of simple-root defining sets $T_{p^s-1} \subseteq T_{p^s-2} \subseteq \dots \subseteq T_1 \subseteq T_0$, where T_t consists of those cyclotomic cosets whose multiplicity is greater than t . This layer viewpoint is the main structural device of the paper.

The key point is that we do not merely lift one simple-root constacyclic code uniformly to repeated-root length. A uniform lift may create trivial layers whose minimum distance is 1, thereby weakening the repeated-root distance. Instead, we allow nonuniform multiplicity profiles. Different cyclotomic cosets may survive through different repeated-root levels, and different levels may carry different half-cyclotomic defining sets. In this way, the dimension is controlled by the total size of the layers, while the distance is controlled by the weakest scaled layer distance.

The first contribution is a profile-theoretic description of repeated-root constacyclic codes of length np^s . In Proposition 1, we prove that multiplicity profiles are equivalent to nested chains of simple-root defining sets, and in Proposition 2, we obtain the dimension formula

$$\dim_{\mathbb{F}_q} \mathcal{C}(e) = np^s - \sum_{i \in \Gamma} e_i |C_i| = np^s - \sum_{t=0}^{p^s-1} |T_t(e)|. \quad (1)$$

Here C_i denotes the q -cyclotomic coset represented by $i \in \Gamma$, and $T_t(e)$ is the union of those cosets whose multiplicity is strictly larger than t . This formula is the basic dimension-control mechanism used throughout the paper.

The second contribution is the repeated-root distance reduction theorem. Theorem 4 shows that if $C_t(e)$ is the simple-root constacyclic code determined by the layer $T_t(e)$ and $t = \sum_j t_j p^j$ is the p -adic expansion of t , then

$$d_H(\mathcal{C}(e)) = \min_{0 \leq t < p^s} P_t d_H(C_t(e)), \quad P_t = \prod_j (t_j + 1). \quad (2)$$

The factor P_t is the Hamming weight of $(x^n - \lambda_0)^t$ in characteristic p , equivalently the Castagnoli–Massey–Schoeller–Seemann p -adic weight factor attached to the repeated-root layer t . Thus the Hamming distance of a repeated-root code of length np^s is computed from the distances of its simple-root layers of length n , weighted by explicit p -adic factors.

The third contribution is the construction of half-cyclotomic repeated-root constacyclic codes. In Definition 7 and Theorem 5, these codes are obtained by choosing the layers T_t as nested initial half-cyclotomic defining sets. Initial segments are used because they preserve the natural ordering of the half-cyclotomic coset classes and give nested layers with explicit dimension and distance certificates. In the two-coset-size case, where the coset sizes are $\ell < m$ and the t -th layer contains $a_{\ell,t}$ cosets of size ℓ and $a_{m,t}$ cosets of size m , with $a_{\ell,t}, a_{m,t} \geq 0$, the dimension becomes

$$\dim_{\mathbb{F}_q} \mathcal{C}_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} (\ell a_{\ell,t} + m a_{m,t}). \quad (3)$$

Hence the dimension can be designed explicitly by deciding how many cosets from each size class appear at each repeated-root level.

The fourth contribution is a layerwise lower-bound method for the minimum distance. Theorem 6 proves that if $\Delta(T_t)$ is any certified lower bound for the simple-root layer $C(T_t)$, then

$$d_H(\mathcal{C}(e)) \geq \min_{0 \leq t < p^s} P_t \Delta(T_t). \quad (4)$$

By a certified lower bound, we mean a rigorously proved lower bound, such as a BCH-type, Hartmann–Tzeng-type, Roos-type, half-cyclotomic, or exact computationally verified distance bound. Non-monotone numerical estimates or heuristic simulations are not used as certificates unless they are converted into rigorous lower bounds. In particular, the lower bounds from the half-cyclotomic simple-root construction of [26] can be inserted

layer by layer. The same framework also allows BCH bounds, exact computed distances, or table-certified simple-root distances to be used at different layers.

The fifth contribution is the realization of the construction for explicit arithmetic length families, including $n = (q^p - 1)/(r\sigma)$, $n = \Phi_{\rho^b}(q)$, and $n = \Phi_{\rho_1\rho_2}(q)$. Here ρ, ρ_1, ρ_2 denote primes controlling the cyclotomic factor involved, $b \geq 1$ is an exponent in the prime-power cyclotomic family, and σ is a divisor parameter chosen so that the length is integral and compatible with the constacyclic order r . For these families, the repeated-root parameters are determined by the arithmetic distribution of the q -cyclotomic cosets in $Z_{n,r}$, rather than by a formal length extension alone.

Finally, we derive a conditional half-rate consequence and lifting criteria for optimal, best-known, and record-breaking codes. Under suitable uniform layer-distance hypotheses, balanced half-cyclotomic chains give repeated-root families with dimension $\dim_{\mathbb{F}_q} \mathcal{C}_{\text{HC}}^{(n)} = np^s/2 + o(np^s)$ and distance bounded below by a logarithmic-type expression of order $N/\log_q N$, where $N = np^s$. This half-rate statement is conditional on the stated layer-distance hypotheses; these hypotheses are not asserted to hold automatically for every admissible profile. We also give criteria for lifting simple-root layers with known or best-known parameters to repeated-root constacyclic codes, allowing comparison with standard code tables and recent repeated-root constructions [18,28,29]. When the lifted parameters meet a proved upper bound, we call the resulting code optimal; otherwise, improvements over available tables are described as best-known or record-breaking.

The paper is organized as follows. Section 2 fixes notation and recalls the simple-root input. Section 3 develops multiplicity profiles, associated simple-root layers, and the half-cyclotomic repeated-root construction. Section 4 proves the repeated-root distance reduction theorem. Section 5 gives the layerwise lower bounds. Section 6 treats half-rate families and explicit arithmetic length families. Section 7 discusses lifting optimal and best-known simple-root layers. Section 8 gives numerical illustrations and comparisons, including the discussion of scope and limitations in Section 8.1. Section 9 concludes the paper and outlines future directions.

2. Preliminaries

This section fixes the notation used throughout the paper and records the simple-root input that will later be lifted to the repeated-root setting. The basic point is that $x^{np^s} - \lambda$ is obtained from $x^n - \lambda_0$ by raising each irreducible factor to the p^s -th power. Thus the usual cyclotomic-coset description of simple-root constacyclic codes remains available, but each coset now carries a multiplicity. Throughout the paper, a lower bound is called certified only when it follows from a proved distance theorem or from an exact verified computation.

Throughout, $q = p^m$ is a prime power with characteristic p . The integer n satisfies $\gcd(n, p) = 1$, and s is a positive integer (the repeated-root exponent). Let $r \mid q - 1$ and let $\lambda_0 \in \mathbb{F}_q^*$ have multiplicative order r . Set $\lambda = \lambda_0^{p^s}$. Because the Frobenius map $a \mapsto a^{p^s}$ is an automorphism of $\mathbb{F}_q$, this relation is always valid and determines the constacyclic shift attached to λ . Hence

$$x^{np^s} - \lambda = (x^n - \lambda_0)^{p^s}. \quad (5)$$

Thus $x^n - \lambda_0$ is the simple-root base polynomial, while p^s records the repeated-root multiplicity.

All cyclotomic cosets below are q -cyclotomic cosets modulo nr . In explicit length families we denote auxiliary primes by ρ, ρ_1, ρ_2 to avoid confusion with the characteristic p . Definitions 1–4 fix the simple-root notation used throughout, while Assumption 1 specifies the two-coset-size regime used in the main bounds.

Definition 1. Let $Z_{n,r} = 1 + ir : 0 \leq i \leq n-1 \subseteq \mathbb{Z}/(nr)\mathbb{Z}$. For $i \in Z_{n,r}$, the q -cyclotomic coset of i modulo nr is $C_i = C_i^{(q,nr)} = iq^j \bmod nr : j \geq 0$. The smallest positive representative of C_i is its coset leader. Write $\Gamma = \Gamma(q, n, r)$ for the set of coset leaders that lie in $Z_{n,r}$.

Let β be a primitive nr -th root of unity with $\beta^n = \lambda_0$. For $i \in \Gamma$ define $M_i(x) = \prod_{j \in C_i} (x - \beta^j) \in \mathbb{F}_q[x]$, so that

$$x^n - \lambda_0 = \prod_{i \in \Gamma} M_i(x) \quad (6)$$

is the factorization into distinct irreducibles. Raising each factor to the p^s -th power gives the repeated-root factorization of $x^{np^s} - \lambda$.

Definition 2. Let $T \subseteq Z_{n,r}$ be a union of whole cyclotomic cosets. The simple-root λ_0 -constacyclic code with defining set T is

$$C(T) = \left\langle \prod_{C_i \subseteq T} M_i(x) \right\rangle \subseteq \mathbb{F}_q[x] / \langle x^n - \lambda_0 \rangle. \quad (7)$$

If $T = \emptyset$, then $C(T) = \mathbb{F}_q^n$ and we set $d_H(C(T)) = 1$.

Definition 3. For a positive integer l , let $\Gamma_l = \{i \in \Gamma : |C_i| = l\}$, $N_l = |\Gamma_l|$. When $N_l > 0$, list the elements of Γ_l increasingly: $\Gamma_l = \delta_1^{(l)} < \delta_2^{(l)} < \dots < \delta_{N_l}^{(l)}$.

Initial segments are used because they preserve nestedness of the later repeated-root layers and allow the available half-cyclotomic distance bounds to be applied directly.

Definition 4. For each l with $N_l > 0$, define

$$H_l^- = \bigcup_{j=1}^{\lfloor N_l/2 \rfloor} C_{\delta_j^{(l)}}; H_l^+ = \bigcup_{j=1}^{\lceil N_l/2 \rceil} C_{\delta_j^{(l)}}. \quad (8)$$

The lower and upper half-cyclotomic defining sets are $H^- = \bigcup_l H_l^-$; $H^+ = \bigcup_l H_l^+$. More generally, a defining set of the form $\bigcup_l \bigcup_{j=1}^{a_l} C_{\delta_j^{(l)}}$ with $0 \leq a_l \leq N_l$ is called an initial half-cyclotomic defining set.

The strongest distance estimates are available when exactly two distinct coset sizes occur.

Assumption 1. The triple (q, n, r) satisfies the two-coset-size hypothesis if there exist positive integers $l < m$ such that $|C_i| : i \in Z_{n,r} = l, m$. Under this hypothesis, the number of cosets of the larger size is $N_m = \frac{n - \gcd(n, (q^l - 1)/r)}{m}$, provided the arithmetic conditions that guarantee exactly two sizes hold.

The following simple-root input from [26] is used later only under the stated two-coset-size hypothesis and with the half-cyclotomic ordering by coset leaders.

Theorem 1 ([26]). Assume that the q -cyclotomic cosets in $Z_{n,r}$ have exactly two sizes $\ell < m$. Let N_ℓ and N_m denote the numbers of cosets of sizes ℓ and m , respectively. Let T^+ be the defining set obtained by taking the first $\lceil N_\ell/2 \rceil$ cosets of size ℓ and the first $\lceil N_m/2 \rceil$ cosets of size m , ordered increasingly by their coset leaders. Then $d_H(C(T^+)) \geq \left\lfloor \frac{qN_m}{2(q-1)} \right\rfloor$.

For later layers, we also use the following standard BCH-type certificate. It summarizes the proof idea behind the consecutive-run lower bounds used throughout the paper.

Theorem 2. Let $T \subseteq \mathbb{Z}_{n,r}$ be a union of q -cyclotomic cosets, and let $C(T)$ be the corresponding simple-root λ_0 -constacyclic code. Let $L(T)$ be the largest integer $L \geq 0$ such that $\{1, 1+r, \dots, 1+(L-1)r\} \subseteq T$. Then $d_H(C(T)) \geq L(T) + 1$. In particular, for the upper half-cyclotomic defining set T^+ , one has $d_H(C(T^+)) \geq L(T^+) + 1$. Consequently, whenever this consecutive-run value improves the lower bound in Theorem 1, it may be used as the stronger certified layer bound. All such bounds are used only as certified lower bounds; equality is not assumed unless it is verified separately.

Proof. This is the BCH bound for constacyclic codes. If the defining set contains the consecutive positions $1, 1+r, \dots, 1+(L-1)r$, then the generator polynomial has the corresponding L consecutive roots $\beta, \beta^{1+r}, \dots, \beta^{1+(L-1)r}$. The BCH bound gives $d_H(C(T)) \geq L + 1$. Taking $L = L(T)$ proves the result. $\square$

3. Multiplicity Profiles and Half-Cyclotomic Construction

The central object of the paper is a multiplicity profile. In the simple-root setting a cyclotomic coset is either present or absent from the defining set. In the repeated-root setting it may occur with any multiplicity between 0 and p^s . The profile is therefore best viewed as a nested sequence of simple-root layers, and these layers are the objects to which the half-cyclotomic distance bounds will be applied.

Definition 5. A multiplicity profile is a map $e : \Gamma \rightarrow \{0, 1, \dots, p^s\}$, $i \mapsto e_i$. Here e_i is the exponent assigned to the irreducible factor $M_i(x)$ corresponding to the coset representative $i \in \Gamma$. The corresponding repeated-root λ -constacyclic code of length np^s is

$$\mathcal{C}(e) = \left\langle \prod_{i \in \Gamma} M_i(x)^{e_i} \right\rangle \subseteq \mathbb{F}_q[x] / \langle x^{np^s} - \lambda \rangle. \quad (9)$$

By (5) and (6), one has $x^{np^s} - \lambda = \prod_{i \in \Gamma} M_i(x)^{p^s}$. Since $\mathbb{F}_q[x] / \langle x^{np^s} - \lambda \rangle$ is a principal ideal ring, every repeated-root λ -constacyclic code of length np^s is generated by a monic divisor of $x^{np^s} - \lambda$. Hence each such code is uniquely of the form $\mathcal{C}(e)$ for a multiplicity profile e .

Definition 6. For a multiplicity profile e and $0 \leq t < p^s$, define $T_t(e) = \bigcup_{e_i > t} C_i \subseteq \mathbb{Z}_{n,r}$. The integer t is the layer index, and $T_t(e)$ consists exactly of the cosets whose multiplicity is strictly larger than t . The t -th associated simple-root layer is the simple-root code $C_t(e) = C(T_t(e)) \subseteq \mathbb{F}_q[x] / \langle x^n - \lambda_0 \rangle$.

Definition 6 is the link between multiplicities and the nested defining sets reconstructed in Proposition 1.

Thus a coset C_i appears in layer t precisely when its multiplicity exceeds t ; larger multiplicity means persistence through more layers.

Proposition 1. Let $e : \Gamma \rightarrow \{0, 1, \dots, p^s\}$ be a multiplicity profile. Then the sets $T_t(e)$ form a nested chain $T_{p^s-1}(e) \subseteq T_{p^s-2}(e) \subseteq \dots \subseteq T_1(e) \subseteq T_0(e)$. Moreover, e is uniquely determined by this chain via

$$e_i = \#\{t : 0 \leq t < p^s, C_i \subseteq T_t(e)\}. \quad (10)$$

Conversely, every nested chain

$$T_{p^s-1} \subseteq T_{p^s-2} \subseteq \dots \subseteq T_1 \subseteq T_0 \quad (11)$$

of unions of q -cyclotomic cosets in $\mathbb{Z}_{n,r}$ arises uniquely in this way. Hence multiplicity profiles are in bijection with such nested chains.

Proof. If $C_i \subseteq T_{t+1}(e)$, then $e_i > t + 1$, and hence $e_i > t$. Thus $C_i \subseteq T_t(e)$. Since the sets $T_t(e)$ are unions of whole cyclotomic cosets, this gives $T_{t+1}(e) \subseteq T_t(e)$ for every $0 \leq t < p^s - 1$, proving the nesting.

For fixed $i \in \Gamma$, one has $C_i \subseteq T_t(e) \iff e_i > t$. Therefore the admissible values of t are exactly $0, 1, \dots, e_i - 1$ if $e_i > 0$, and there are none if $e_i = 0$. Hence their number is e_i , which proves (10). Thus the chain determines the profile uniquely.

Conversely, assume (11) is a nested chain of unions of q -cyclotomic cosets. For each $i \in \Gamma$, define e_i by the right-hand side of (10). Then $0 \leq e_i \leq p^s$, so e is a multiplicity profile. We show that its associated layers are the prescribed sets. Fix $i \in \Gamma$ and put $A_i = \{t : 0 \leq t < p^s, C_i \subseteq T_t\}$. By nesting, if $t \in A_i$ and $0 \leq u \leq t$, then $u \in A_i$. Hence A_i is an initial interval of $\{0, 1, \dots, p^s - 1\}$, namely $A_i = \{0, 1, \dots, e_i - 1\}$, with the usual interpretation that this set is empty when $e_i = 0$. Consequently, $C_i \subseteq T_t \iff e_i > t$. Since each T_t is a union of whole cyclotomic cosets, it follows that

$$T_t = \bigcup_{\substack{i \in \Gamma \\ C_i \subseteq T_t}} C_i = \bigcup_{\substack{i \in \Gamma \\ e_i > t}} C_i = T_t(e)$$

for every $0 \leq t < p^s$. Thus the chain is recovered from the profile e . The two constructions are inverse to each other, and the desired bijection follows. $\square$

Proposition 2. For the repeated-root code $\mathcal{C}(e)$,

$$\dim_{\mathbb{F}_q} \mathcal{C}(e) = np^s - \sum_{i \in \Gamma} e_i |C_i|. \quad (12)$$

Equivalently, in terms of the associated layers,

$$\dim_{\mathbb{F}_q} \mathcal{C}(e) = np^s - \sum_{t=0}^{p^s-1} |T_t(e)|. \quad (13)$$

Proof. Let $g_e(x) = \prod_i M_i(x)^{e_i}$ and write $x^{np^s} - \lambda = g_e h_e$. Multiplication by g_e induces an $\mathbb{F}_q$ -linear isomorphism $\mathbb{F}_q[x] / \langle h_e \rangle \cong \mathcal{C}(e)$, so $\dim \mathcal{C}(e) = \deg h_e = np^s - \deg g_e$. Since $\deg M_i = |C_i|$, we have $\deg g_e = \sum_i e_i |C_i|$, giving (12).

For (13), note $|T_t(e)| = \sum_{e_i > t} |C_i|$. Summing over t yields $\sum_t |T_t(e)| = \sum_i e_i |C_i|$, which inserted into (12) gives (13). $\square$

The layer form (13) is central for the construction: the dimension is controlled by the total size of the nested defining sets. In the next section, the minimum distance will be governed by the weakest scaled layer distance, providing the essential mechanism of the repeated-root half-cyclotomic construction.

4. The Repeated-Root Distance Reduction Theorem

This section proves the distance mechanism behind the construction. Although the repeated-root code has length np^s , its Hamming distance is controlled by simple-root codes of length n . Passing from the repeated-root code to the t -th simple-root layer introduces the p -adic multiplier P_t . This is the bridge that allows simple-root bounds to be inserted into a repeated-root distance formula.

Throughout this section, the Hamming weight of a residue class in $\mathbb{F}_q[x] / \langle x^{np^s} - \lambda \rangle$ means the Hamming weight of its unique representative of degree $< np^s$.

For $0 \leq t < p^s$, write its p -adic expansion as $t = \sum_{j=0}^{s-1} t_j p^j$, where $0 \leq t_j \leq p-1$, and define $P_t = \prod_{j=0}^{s-1} (t_j + 1)$.

Since $\deg(x^n - \lambda_0)^t = nt < np^s$, no reduction occurs in the quotient. By Lucas' theorem, the number of nonzero binomial coefficients in the expansion of $(x^n - \lambda_0)^t$ is exactly P_t . Thus $\text{wt}_H((x^n - \lambda_0)^t) = P_t$.

We shall use the following form of the repeated-root distance formula of Castagnoli–Massey–Schoeller–Seemann.

Theorem 3 ([27]). Let $\pi = x^n - \lambda_0$, and let $\mathcal{C}$ be an ideal of $\mathbb{F}_q[x]/\langle \pi^{p^s} \rangle$. For $0 \leq t < p^s$, define

$$\text{Tor}_t(\mathcal{C}) = \{\bar{a} \in \mathbb{F}_q[x]/\langle \pi \rangle : \text{there exists a lift } a \in \mathbb{F}_q[x] \text{ of } \bar{a} \text{ such that } \pi^t a \in \mathcal{C}\}.$$

Then

$$d_H(\mathcal{C}) = \min_{0 \leq t < p^s} \{\text{wt}_H(\pi^t), d_H(\text{Tor}_t(\mathcal{C}))\}.$$

Theorem 4. Let $\mathcal{C}(e)$ be the repeated-root λ -constacyclic code defined by the multiplicity profile $e : \Gamma \rightarrow \{0, 1, \dots, p^s\}$, namely $\mathcal{C}(e) = \langle \prod_{i \in \Gamma} M_i(x)^{e_i} \rangle \subseteq \mathbb{F}_q[x]/\langle x^{np^s} - \lambda \rangle$. For $0 \leq t < p^s$, set $T_t(e) = \bigcup_{e_i > t} C_i$ and let $C_t(e) = \mathcal{C}(T_t(e)) \subseteq \mathbb{F}_q[x]/\langle x^n - \lambda_0 \rangle$ be the associated simple-root layer. If $t = \sum_{j=0}^{s-1} t_j p^j$, with $0 \leq t_j \leq p-1$, and $P_t = \prod_{j=0}^{s-1} (t_j + 1)$, then $d_H(\mathcal{C}(e)) = \min_{0 \leq t < p^s} \{P_t, d_H(C_t(e))\}$.

Proof. Put $\pi = x^n - \lambda_0$. Then $x^{np^s} - \lambda = \pi^{p^s}$. Let $R_s = \mathbb{F}_q[x]/\langle \pi^{p^s} \rangle$, $R_0 = \mathbb{F}_q[x]/\langle \pi \rangle$. Thus R_s is the repeated-root ambient ring and R_0 is the corresponding simple-root ambient ring.

By Theorem 3, it is enough to identify the torsion layers of $\mathcal{C}(e)$ and compute the weights $\text{wt}_H(\pi^t)$.

Since $\pi = x^n - \lambda_0 = \prod_{i \in \Gamma} M_i(x)$, the Chinese remainder theorem gives $R_s \cong \prod_{i \in \Gamma} \mathbb{F}_q[x]/\langle M_i(x)^{p^s} \rangle$, $R_0 \cong \prod_{i \in \Gamma} \mathbb{F}_q[x]/\langle M_i(x) \rangle$. Under this decomposition, the ideal $\mathcal{C}(e) = \langle \prod_{i \in \Gamma} M_i(x)^{e_i} \rangle$ corresponds componentwise to $\prod_{i \in \Gamma} \langle M_i(x)^{e_i} \rangle \subseteq \prod_{i \in \Gamma} \mathbb{F}_q[x]/\langle M_i(x)^{p^s} \rangle$. Moreover, in the i -th component, π^t is a unit multiple of $M_i(x)^t$, since the factors $M_j(x)$, $j \neq i$, are units modulo $M_i(x)^{p^s}$.

Let $\bar{a} \in R_0$, and write its i -th component as $\bar{a}_i \in \mathbb{F}_q[x]/\langle M_i(x) \rangle$. The condition $\bar{a} \in \text{Tor}_t(\mathcal{C}(e))$ is equivalent, component by component, to the existence of a lift a_i such that $M_i(x)^t a_i \in \langle M_i(x)^{e_i} \rangle$ in $\mathbb{F}_q[x]/\langle M_i(x)^{p^s} \rangle$. If $e_i \leq t$, there is no restriction on $\bar{a}_i$. If $e_i > t$, then a_i must be divisible by $M_i(x)$, and hence $\bar{a}_i = 0$ modulo $M_i(x)$. Therefore, the torsion layer has zeros precisely in the components with $e_i > t$. Hence $\text{Tor}_t(\mathcal{C}(e)) = \mathcal{C}(T_t(e)) = C_t(e)$.

It remains to compute $\text{wt}_H(\pi^t)$. Since $0 \leq t < p^s$, the polynomial $\pi^t = (x^n - \lambda_0)^t$ has degree $nt < np^s$, so its coefficient support is not changed by reduction modulo $x^{np^s} - \lambda$. Expanding gives $\pi^t = \sum_{k=0}^t \binom{t}{k} (-\lambda_0)^{t-k} x^{nk}$. Since $\lambda_0 \neq 0$, the nonzero coefficients are exactly those for which $\binom{t}{k} \not\equiv 0 \pmod{p}$. By Lucas' theorem, if $t = \sum_{j=0}^{s-1} t_j p^j$ and $k = \sum_{j=0}^{s-1} k_j p^j$, then $\binom{t}{k} \not\equiv 0 \pmod{p} \iff k_j \leq t_j$ for all j . Thus the number of nonzero binomial coefficients is $\prod_{j=0}^{s-1} (t_j + 1) = P_t$. The monomials x^{nk} , $0 \leq k \leq t < p^s$, have distinct exponents between 0 and $np^s - 1$, so their supports do not overlap. Therefore, $\text{wt}_H(\pi^t) = P_t$. Substituting $\text{Tor}_t(\mathcal{C}(e)) = C_t(e)$ and $\text{wt}_H(\pi^t) = P_t$ into Theorem 3 gives $d_H(\mathcal{C}(e)) = \min_{0 \leq t < p^s} \{P_t, d_H(C_t(e))\}$. This completes the proof. $\square$

The next auxiliary function records the smallest possible p -adic factor P_t among all layers $t \geq \ell$. It is useful when a profile is optimized over a range of repeated-root levels rather than by checking every value of t separately.

Proposition 3. For each $\ell = 0, 1, \dots, p^s$ define

$$\mathcal{M}_s(\ell) = \begin{cases} (p+1 - \lfloor p^{\log_p(p^s-\ell)} \rfloor) p^{s-1-\lfloor \log_p(p^s-\ell) \rfloor}, & 0 \leq \ell < p^s, \\ +\infty, & \ell = p^s. \end{cases} \quad (14)$$

Then $\min_{\ell \leq t < p^s} P_t = \mathcal{M}_s(\ell)$.

Proof. If $\ell = p^s$, the interval $[\ell, p^s)$ is empty, and the convention $\mathcal{M}_s(p^s) = +\infty$ applies. Assume $0 \leq \ell < p^s$ and put $h = p^s - \ell$. Write $h = cp^a + r_h$, $1 \leq c \leq p-1$, $0 \leq r_h < p^a$, where $a = \lfloor \log_p h \rfloor$ and $c = \lfloor h/p^a \rfloor$. Equivalently, $c = \lfloor p^{\log_p h} \rfloor$. Since $cp^a \leq h < (c+1)p^a$, we have $p^s - (c+1)p^a < \ell \leq p^s - cp^a$. Set $t_0 = p^s - cp^a$. Then $\ell \leq t_0 < p^s$, and the base- p digits of t_0 are 0 below position a , $p-c$ in position a , and $p-1$ above position a . Hence $P_{t_0} = (p-c+1)p^{s-a-1}$. Thus $\min_{\ell \leq t < p^s} P_t \leq (p-c+1)p^{s-a-1}$.

Conversely, let $\ell \leq t < p^s$ and write $u = p^s - t$. Then $1 \leq u \leq h$. If $u = bp^a$ with $1 \leq b \leq c$, then the base- p digits of $p^s - u$ give $P_t = (p-b+1)p^{s-a-1} \geq (p-c+1)p^{s-a-1}$. If $u = bp^a + r_u$ with $0 \leq b \leq c$ and $0 < r_u < p^a$, then the digits above position a of $p^s - u$ are all $p-1$, the digit in position a contributes at least $p-b$, and the lower a digits represent the positive integer $p^a - r_u$; therefore their contribution to the product P_t is at least 2. Hence $P_t \geq 2(p-b)p^{s-a-1} \geq (p-c+1)p^{s-a-1}$, where the last inequality follows from $b \leq c \leq p-1$. Therefore $P_t \geq (p-c+1)p^{s-a-1}$ for every $t \in [\ell, p^s)$. Combining the two inequalities gives $\min_{\ell \leq t < p^s} P_t = (p-c+1)p^{s-a-1}$. Substituting $h = p^s - \ell$, $a = \lfloor \log_p h \rfloor$, and $c = \lfloor p^{\log_p h} \rfloor$ gives exactly $\mathcal{M}_s(\ell)$. $\square$

The role of Proposition 3 is auxiliary: it is useful when one wants to optimize a profile over a range of levels rather than inspect every p -adic value P_t separately. The same expression applies here because the p -adic weight of $(x^n - \lambda_0)^t$ is independent of the particular constacyclic shift.

Half-Cyclotomic Repeated-Root Construction

We now impose the half-cyclotomic ordering on the associated layers. Each layer is an initial union of cyclotomic cosets within each size class, and the resulting nested chain determines the repeated-root multiplicities. Initial segments are used because they preserve nestedness of the repeated-root layers and allow the available half-cyclotomic distance certificates to be applied layer by layer.

Definition 7. A nested chain $T_{p^s-1} \subseteq T_{p^s-2} \subseteq \dots \subseteq T_1 \subseteq T_0$ of unions of q -cyclotomic cosets in $\mathbb{Z}_{n,r}$ is called half-cyclotomic if, for each $0 \leq t < p^s$ and each coset-size class ℓ , there exists an integer $a_{\ell,t} \in \{0, 1, \dots, N_\ell\}$ such that $T_t = \bigcup_\ell \bigcup_{j=1}^{a_{\ell,t}} C_{\delta_j^{(\ell)}}$. Here $a_{\ell,t}$ is the number of cosets of size ℓ selected in the t -th layer. Equivalently, T_t is obtained by taking an initial segment of the ordered cosets in every size class. The nesting condition translates to $a_{\ell,t+1} \leq a_{\ell,t}$ for every ℓ and $0 \leq t < p^s - 1$.

Construction 1. Let $(T_t)_{0 \leq t < p^s}$ be a nested half-cyclotomic chain with parameters $(a_{\ell,t})$. Define a multiplicity profile by $e_i = \#\{t : 0 \leq t < p^s, C_i \subseteq T_t\}$, $i \in \Gamma$. The associated half-cyclotomic repeated-root code is

$$C_{\text{HC}} = \mathcal{C}(e) = \left\langle \prod_{i \in \Gamma} M_i(x)^{e_i} \right\rangle \subseteq \mathbb{F}_q[x] / \langle x^{np^s} - \lambda \rangle. \quad (15)$$

This construction admits two equivalent interpretations. From the simple-root side, one chooses a nested tower of half-cyclotomic defining sets and records how long each

coset remains present. From the repeated-root side, one starts with the multiplicities e_i and recovers the layers by the threshold rule $e_i > t$. This equivalence is the mechanism that allows the simple-root half-cyclotomic distance bounds to be applied level by level in the repeated-root distance formula.

Theorem 5. Let C_{HC} be the code associated with the nested half-cyclotomic chain $(T_t)_{0 \leq t < p^s}$, and write $T_t = \bigcup_{\ell} \bigcup_{j=1}^{a_{\ell,t}} C_{\delta_j^{(\ell)}}$. Then $\dim_{\mathbb{F}_q} C_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} \sum_{\ell} \ell a_{\ell,t}$. In particular, under the two-coset-size hypothesis with coset sizes $\ell < m$, this becomes $\dim_{\mathbb{F}_q} C_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} (\ell a_{\ell,t} + m a_{m,t})$.

Proof. By Construction 1, the code C_{HC} is the repeated-root code $\mathcal{C}(e)$ attached to the multiplicity profile $e_i = \#\{t : 0 \leq t < p^s, C_i \subseteq T_t\}$. By the layer reconstruction result (Proposition 1), the associated simple-root layers of $\mathcal{C}(e)$ are precisely the prescribed sets T_t . Hence the dimension formula for multiplicity profiles (Proposition 2) gives $\dim_{\mathbb{F}_q} C_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} |T_t|$. It remains to compute $|T_t|$ from the half-cyclotomic parameters. For fixed t , the set T_t is a disjoint union of whole q -cyclotomic cosets. In the coset-size class ℓ , it contains exactly $C_{\delta_1^{(\ell)}}, \dots, C_{\delta_{a_{\ell,t}}^{(\ell)}}$, each of cardinality ℓ . Therefore the contribution of this size class to $|T_t|$ is $\ell a_{\ell,t}$. Summing over all coset-size classes yields $|T_t| = \sum_{\ell} \ell a_{\ell,t}$. Substituting this into the dimension formula gives the first stated expression. If only two coset sizes $\ell < m$ occur, the inner sum reduces to the two terms $\ell a_{\ell,t} + m a_{m,t}$, which proves the special case. $\square$

5. Main Lower Bounds

The previous section reduced the repeated-root distance problem to the distances of the simple-root layers, weighted by the factors P_t . We now turn this identity into usable lower bounds by inserting certified simple-root information. Thus, the results in this section should be read as layerwise consequences of Theorem 4: whenever a layer distance is known exactly, the corresponding term is exact; otherwise, only a certified lower bound is used.

In particular, the half-cyclotomic bound from Theorem 1 is applied only to the lower or upper half-cyclotomic layers for which its hypotheses hold. For arbitrary initial layers, we use a verified BCH run, an exact computed distance, or another certified simple-root bound.

Let $T \subseteq Z_{n,r}$ be a union of q -cyclotomic cosets. A number $\Delta(T)$ is called a certified lower bound for $C(T)$ if $d_H(C(T)) \geq \Delta(T)$. We use the conventions $\Delta(\emptyset) = 1$ and $\Delta(Z_{n,r}) = +\infty$. Thus the estimates in this section are lower bounds unless exact layer distances are known.

Theorem 6. Let $\mathcal{C}(e)$ be the repeated-root code determined by a multiplicity profile e , and let $T_t(e) = \bigcup_{\substack{i \in \Gamma \\ e_i > t}} C_i$ for $0 \leq t < p^s$ be its associated simple-root layers. Suppose that, for every t , a certified lower bound $\Delta(T_t(e))$ is known. Then $d_H(\mathcal{C}(e)) \geq \min_{0 \leq t < p^s} P_t \Delta(T_t(e))$.

Proof. By Theorem 4, $d_H(\mathcal{C}(e)) = \min_{0 \leq t < p^s} P_t d_H(C(T_t(e)))$. Since $d_H(C(T_t(e))) \geq \Delta(T_t(e))$ for every t and $P_t > 0$, the claim follows by taking the minimum over $0 \leq t < p^s$. $\square$

Theorem 7. Let C_{HC} be the repeated-root code attached to a nested half-cyclotomic chain $(T_t)_{0 \leq t < p^s}$. For each layer choose a certified bound D_t as follows:

$$D_t = \begin{cases} \left\lfloor \frac{qN_m}{2(q-1)} \right\rfloor, & \text{if } T_t = H^+ \text{ and the upper-half hypotheses hold,} \\ \left\lfloor \frac{qN_m}{2(q-1)} \right\rfloor, & \text{if } T_t = H^- \text{ and the lower-half hypotheses hold,} \\ L(T_t) + 1, & \text{if a consecutive run } \{1, 1+r, \dots, 1+(L(T_t)-1)r\} \subseteq T_t \text{ is verified,} \\ 1, & \text{if } T_t = \emptyset, \\ +\infty, & \text{if } T_t = Z_{n,r}. \end{cases} \quad (16)$$

Here the upper-half and lower-half hypotheses mean that the corresponding half-cyclotomic defining set contains a consecutive run of length at least $\left\lfloor \frac{qN_m}{2(q-1)} \right\rfloor - 1$. Then $d_H(C_{\text{HC}}) \geq \min_{0 \leq t < p^s} P_t D_t$. More generally, the same conclusion holds for any choice of certified bounds $D_t \leq d_H(C(T_t))$.

Proof. By Construction 1 and Proposition 1, the simple-root layers of C_{HC} are precisely the prescribed sets T_t . Hence Theorem 4 gives $d_H(C_{\text{HC}}) = \min_{0 \leq t < p^s} P_t d_H(C(T_t))$. It remains to check that each displayed value of D_t is a certified lower bound for $d_H(C(T_t))$.

The consecutive-run case follows directly from Theorem 2. If $T_t = H^+$ and the upper-half hypotheses hold, then the corresponding consecutive run has length at least $\left\lfloor \frac{qN_m}{2(q-1)} \right\rfloor - 1$, and hence $d_H(C(H^+)) \geq \left\lfloor \frac{qN_m}{2(q-1)} \right\rfloor$. The same argument applies to $T_t = H^-$ under the lower-half hypotheses. If $T_t = \emptyset$, then $C(T_t) = \mathbb{F}_q^n$, so $d_H(C(T_t)) = 1$. If $T_t = Z_{n,r}$, then $C(T_t) = \{0\}$, and we use the convention $d_H(\{0\}) = +\infty$. Thus every displayed value satisfies $D_t \leq d_H(C(T_t))$. Substitution into the repeated-root distance formula gives $d_H(C_{\text{HC}}) \geq \min_{0 \leq t < p^s} P_t D_t$. The final assertion follows in the same way for any family of certified bounds $D_t \leq d_H(C(T_t))$. $\square$

Theorem 8. Let C_{HC} be attached to a nested half-cyclotomic chain $(T_t)_{0 \leq t < p^s}$. Suppose that, for each nonempty proper layer T_t , an integer $L_t \geq 0$ has been verified such that $\{1, 1+r, \dots, 1+(L_t-1)r\} \subseteq T_t$, with the convention that this set is empty when $L_t = 0$. Define

$$E_t = \begin{cases} L_t + 1, & \text{if } \emptyset \neq T_t \neq Z_{n,r}, \\ 1, & \text{if } T_t = \emptyset, \\ +\infty, & \text{if } T_t = Z_{n,r}. \end{cases} \quad (17)$$

Then $d_H(C_{\text{HC}}) \geq \min_{0 \leq t < p^s} P_t E_t$. If, in addition, a layer T_t has another certified simple-root lower bound $B_t \leq d_H(C(T_t))$, then E_t may be replaced by $\max\{E_t, B_t\}$ for that layer. In particular, if a layer is an admissible half-cyclotomic layer for which the half-cyclotomic bound is already proved, then one may replace E_t by the larger of E_t and that bound.

Proof. By Construction 1 and Proposition 1, the simple-root layers of C_{HC} are exactly the prescribed defining sets T_t . Therefore, Theorem 4 gives $d_H(C_{\text{HC}}) = \min_{0 \leq t < p^s} P_t d_H(C(T_t))$. For nonempty proper layers, Theorem 2 gives $d_H(C(T_t)) \geq L_t + 1$. If $T_t = \emptyset$, then $C(T_t) = \mathbb{F}_q^n$, so $d_H(C(T_t)) = 1 = E_t$. If $T_t = Z_{n,r}$, then $C(T_t) = \{0\}$, and with the convention $d_H(\{0\}) = +\infty$ we again have $E_t = d_H(C(T_t))$. Thus $E_t \leq d_H(C(T_t))$ for every $0 \leq t < p^s$, and the desired inequality follows by substitution into the distance reduction formula.

Finally, if $B_t \leq d_H(C(T_t))$ is another certified bound, then $\max\{E_t, B_t\} \leq d_H(C(T_t))$. Replacing E_t by this maximum in the same argument is therefore valid. $\square$

Corollary 1. Let $\mathcal{C}_{\text{HC}}$ be attached to the nested half-cyclotomic chain $(T_t)_{0 \leq t < p^s}$. Suppose that the chosen certified layer bound is exact for every layer: $\Delta_{\text{HC}}(T_t) = d_{\text{H}}(C(T_t))$ for all $0 \leq t < p^s$. Then $d_{\text{H}}(\mathcal{C}_{\text{HC}}) = \min_{0 \leq t < p^s} P_t \Delta_{\text{HC}}(T_t)$.

Proof. By Theorem 4, $d_{\text{H}}(\mathcal{C}_{\text{HC}}) = \min_{0 \leq t < p^s} P_t d_{\text{H}}(C(T_t))$. Since $\Delta_{\text{HC}}(T_t) = d_{\text{H}}(C(T_t))$ for every t , substitution gives $d_{\text{H}}(\mathcal{C}_{\text{HC}}) = \min_{0 \leq t < p^s} P_t \Delta_{\text{HC}}(T_t)$. $\square$

Corollary 1 will be used when exact layer distances are known from computation or from a sharp simple-root bound. Consequently, the layerwise lower bound is sharp whenever the certified bounds are exact on the layers that determine the minimum in Theorem 4; otherwise, the result should be interpreted only as a certified lower bound.

6. Half-Rate and Explicit Length Families

The previous section established layerwise certified lower bounds for fixed repeated-root codes. We now pass to families of lengths. The dimension is controlled by the total size of the nested layers, while the distance is controlled by the weakest scaled simple-root layer distance. Thus, half-rate families arise by balancing the total layer size around one half of the repeated-root length and avoiding full-space layers of distance 1. We then record explicit arithmetic length families for which the coset distribution, dimension, and certified lower bounds can be computed directly.

Definition 8. Let $(T_t)_{0 \leq t < p^s}$ be a nested half-cyclotomic chain. We call the chain balanced if $\sum_{t=0}^{p^s-1} |T_t| = \frac{np^s}{2} + o(np^s)$ along the family under consideration. Equivalently, for the associated half-cyclotomic repeated-root code $\mathcal{C}_{\text{HC}}$, $\dim_{\mathbb{F}_q} \mathcal{C}_{\text{HC}} = \frac{np^s}{2} + o(np^s)$.

Definition 8 is the asymptotic condition used in the following theorem.

The next theorem is conditional: its conclusion holds only for families whose simple-root layers have certified distance lower bounds of the stated order. The hypotheses are not automatic and must be verified for the particular layer chains under consideration.

Theorem 9. Let q and s be fixed, and let $n \rightarrow \infty$ through a family with $\gcd(n, p) = 1$. For each n , let $(T_t^{(n)})_{0 \leq t < p^s}$ be a nested half-cyclotomic chain, and let $\mathcal{C}_{\text{HC}}^{(n)} \subseteq \mathbb{F}_q[x] / \langle x^{np^s} - \lambda \rangle$ be the associated repeated-root code. Assume that the chains are balanced: $\sum_{t=0}^{p^s-1} |T_t^{(n)}| = \frac{np^s}{2} + o(np^s)$. Assume also that the layer distances are certified uniformly in the following sense: there exist constants $c_t > 0$, independent of n , such that $d_{\text{H}}(C(T_t^{(n)})) \geq c_t \frac{n}{\log_q n}$ for every $0 \leq t < p^s$. These layer-distance hypotheses are assumptions of the theorem; they are not asserted to hold automatically for all half-cyclotomic chains. Then $\dim_{\mathbb{F}_q} \mathcal{C}_{\text{HC}}^{(n)} = \frac{np^s}{2} + o(np^s)$, and $d_{\text{H}}(\mathcal{C}_{\text{HC}}^{(n)}) \geq C \frac{n}{\log_q n}$, where $C = \min_{0 \leq t < p^s} P_t c_t > 0$. Equivalently, if $N = np^s$, then $d_{\text{H}}(\mathcal{C}_{\text{HC}}^{(n)}) \geq c' \frac{N}{\log_q N}$ for some constant $c' > 0$ depending only on q , s , and the constants c_t .

Proof. Let $N = np^s$. By Theorem 5, the dimension of $\mathcal{C}_{\text{HC}}^{(n)}$ is $\dim_{\mathbb{F}_q} \mathcal{C}_{\text{HC}}^{(n)} = np^s - \sum_{t=0}^{p^s-1} |T_t^{(n)}|$. Using the balancedness assumption gives

$$\begin{aligned} \dim_{\mathbb{F}_q} \mathcal{C}_{\text{HC}}^{(n)} &= np^s - \left(\frac{np^s}{2} + o(np^s) \right) \\ &= \frac{np^s}{2} + o(np^s). \end{aligned}$$

Thus the rate tends to $1/2$.

For the distance, Theorem 4 gives $d_H(\mathcal{C}_{\text{HC}}^{(n)}) = \min_{0 \leq t < p^s} P_t d_H(C(T_t^{(n)}))$. By hypothesis, $d_H(C(T_t^{(n)})) \geq c_t \frac{n}{\log_q n}$ for every t . Therefore $P_t d_H(C(T_t^{(n)})) \geq P_t c_t \frac{n}{\log_q n}$ for each t . Taking the minimum over $0 \leq t < p^s$, we obtain $d_H(\mathcal{C}_{\text{HC}}^{(n)}) \geq (\min_{0 \leq t < p^s} P_t c_t) \frac{n}{\log_q n}$. Set $C = \min_{0 \leq t < p^s} P_t c_t$. Since each $P_t > 0$ and each $c_t > 0$, we have $C > 0$. Hence $d_H(\mathcal{C}_{\text{HC}}^{(n)}) \geq C \frac{n}{\log_q n}$. Finally, $N = np^s$, so $\log_q N = \log_q n + \log_q(p^s) = \log_q n + s \log_q p$. Since q and s are fixed, $s \log_q p$ is constant. For all large n , $\log_q n \leq \log_q N$. Hence

$$\frac{n}{\log_q n} = \frac{N/p^s}{\log_q n} \geq \frac{N/p^s}{\log_q N} = \frac{1}{p^s} \frac{N}{\log_q N}.$$

Therefore $d_H(\mathcal{C}_{\text{HC}}^{(n)}) \geq \frac{C}{p^s} \frac{N}{\log_q N}$. Taking $c' = \frac{C}{p^s} > 0$ gives the claimed form. $\square$

Remark 1. The layer condition in Theorem 9 is essential. If some layer T_t is empty, then $C(T_t) = \mathbb{F}_q^n$ has distance 1, and the repeated-root distance formula includes the term P_t . In that case one cannot deduce growth of order $N / \log_q N$. Thus the construction must keep all repeated-root layers protected by simple-root defining sets with growing distance. Accordingly, Theorem 9 should be read as a conditional transfer principle: once the required layer-distance certificates are available, the repeated-root family inherits the stated half-rate and logarithmic-distance behavior.

This observation will be used implicitly when constructing the explicit families below.

6.1. Explicit Length Families

We now give explicit length families with completely computable repeated-root parameters. The essential arithmetic feature is that the defining positions $Z_{n,r} = \{1 + jr : 0 \leq j \leq n - 1\} \subseteq \mathbb{Z}/(nr)\mathbb{Z}$ split into exactly two q -cyclotomic coset sizes. Once this is known, the dimension follows by counting selected cosets in each layer, and the distance bound follows from the layerwise repeated-root lower bound.

Throughout the remainder of this section $q = p^m$, where p is the characteristic of $\mathbb{F}_q$. The repeated-root exponent is $s \geq 1$, and $\gcd(n, p) = 1$. We write $\lambda = \lambda_0^{p^s}$, $r = \text{ord}(\lambda_0)$, $r \mid q - 1$. Auxiliary primes appearing in the arithmetic length formulas are denoted by ρ, ρ_1, ρ_2 . The parameter $b \geq 1$ is used in the prime-power cyclotomic family $n = \Phi_{\rho^b}(q)$, and σ denotes the divisor parameter in the family $nr = (q^\rho - 1)/\sigma$. Thus the repeated-root multiplier is p^s , while the simple-root length n may involve the auxiliary primes.

6.2. Arithmetic Counting of Cyclotomic Cosets

We begin with two elementary counting lemmas. They are included to make the coset-size computations self-contained.

Lemma 1. Let q be a prime power, let n be a positive integer with $\gcd(n, q) = 1$, and let $r \mid q - 1$. Put $Z_{n,r} = \{1 + jr : 0 \leq j \leq n - 1\} \subseteq \mathbb{Z}/(nr)\mathbb{Z}$. For $i \in Z_{n,r}$, let $C_i = \{iq^j \bmod nr : j \geq 0\}$, and put $M = \text{ord}_{nr}(q)$. For $u \geq 1$, define

$$F(u) = \{i \in Z_{n,r} : iq^u \equiv i \pmod{nr}\}. \quad (18)$$

Then the following assertions hold.

(i) If $d = \gcd(nr, i)$, then

$$|C_i| = \text{ord}_{nr/d}(q). \quad (19)$$

(ii) Put $A_u = (q^u - 1)/r$ and $d_u = \gcd(n, A_u)$. Then

$$|F(u)| = \begin{cases} d_u, & \text{if } \gcd(r, n/d_u) = 1, \\ 0, & \text{if } \gcd(r, n/d_u) > 1. \end{cases} \quad (20)$$

(iii) If $u \mid M$, then $F(u)$ is exactly the union of all q -cyclotomic cosets in $Z_{n,r}$ whose sizes divide u . Consequently, if $f(u) = |F(u)|$, then

$$f(u) = \sum_{a|u} aN_a, \quad (21)$$

where N_a denotes the number of q -cyclotomic cosets in $Z_{n,r}$ of size a .

(iv) For every divisor $a \mid M$, the number of q -cyclotomic cosets in $Z_{n,r}$ of size exactly a is

$$N_a = \frac{1}{a} \sum_{u|a} \mu\left(\frac{a}{u}\right) f(u), \quad (22)$$

where μ is the Möbius function.

(v) In particular, suppose that $\ell < m$ are divisors of M and that the values $f(u)$, computed by (ii), satisfy

$$N_\ell = \frac{1}{\ell} \sum_{u|\ell} \mu\left(\frac{\ell}{u}\right) f(u) > 0, \quad N_m = \frac{1}{m} \sum_{u|m} \mu\left(\frac{m}{u}\right) f(u) > 0, \quad (23)$$

and $N_a = 0$ for every divisor $a \mid M$ with $a \notin \{\ell, m\}$. Then the cyclotomic cosets in $Z_{n,r}$ have exactly two sizes, namely ℓ and m . In this case $f(\ell) = \ell N_\ell$, $f(m) = n$, and

$$N_\ell = \frac{f(\ell)}{\ell}, \quad N_m = \frac{n - f(\ell)}{m}. \quad (24)$$

Proof. For (i), by definition $|C_i|$ is the least positive integer u such that $iq^u \equiv i \pmod{nr}$. This congruence is equivalent to $nr \mid i(q^u - 1)$. If $d = \gcd(nr, i)$, then this is equivalent to $nr/d \mid q^u - 1$. Hence the least such u is $\text{ord}_{nr/d}(q)$.

We prove (ii). Every element of $Z_{n,r}$ has the form $i = 1 + jr$, where $0 \leq j \leq n - 1$. The condition $iq^u \equiv i \pmod{nr}$ is equivalent to $nr \mid i(q^u - 1)$. Since $q^u - 1 = rA_u$, this becomes $n \mid A_u(1 + rj)$. Thus $|F(u)|$ is the number of solutions $j \pmod{n}$ of

$$A_u(1 + rj) \equiv 0 \pmod{n}.$$

Let $d_u = \gcd(n, A_u)$, and write $n = d_u n'$ and $A_u = d_u A'_u$, with $\gcd(A'_u, n') = 1$. The congruence is equivalent to $A'_u(1 + rj) \equiv 0 \pmod{n'}$, and hence to $1 + rj \equiv 0 \pmod{n'}$. This linear congruence is solvable if and only if $\gcd(r, n') = 1$, that is, if and only if $\gcd(r, n/d_u) = 1$. If it is solvable, it has exactly one solution modulo n' , and therefore exactly d_u solutions modulo $n = d_u n'$. If it is not solvable, it has no solution. This proves (ii).

For (iii), let $u \mid M$. An element $i \in Z_{n,r}$ belongs to $F(u)$ if and only if $iq^u \equiv i \pmod{nr}$. This holds if and only if the period of i under multiplication by q , namely $|C_i|$, divides u . Hence $F(u)$ is precisely the union of all cosets whose sizes divide u . Therefore, $f(u) = |F(u)| = \sum_{a|u} aN_a$.

Assertion (iv) follows from (iii) by Möbius inversion applied to the arithmetic functions $f(u)$ and $g(a) = aN_a$ on the divisor lattice of M . Thus $aN_a = \sum_{u|a} \mu(a/u) f(u)$, which gives the stated formula for N_a .

Finally, (v) follows immediately from (iv). If the only nonzero values among the N_a , $a \mid M$, occur at $a = \ell$ and $a = m$, then the only possible coset sizes are ℓ and m . Since

$N_\ell, N_m > 0$, both sizes actually occur. In that case $F(\ell)$ is the union of all cosets of size ℓ , so $f(\ell) = \ell N_\ell$. Also $F(m) = Z_{n,r}$, because all coset sizes divide m , and hence $f(m) = n$. Subtracting the $f(\ell) = \ell N_\ell$ elements lying in cosets of size ℓ from the total number n gives $mN_m = n - f(\ell)$. Therefore $N_\ell = f(\ell)/\ell$ and $N_m = (n - f(\ell))/m$. $\square$

Lemma 1 is the counting step used in Proposition 4.

Proposition 4. Assume that the q -cyclotomic cosets in $Z_{n,r}$ have exactly two sizes $\ell < m$. Let N_ℓ and N_m be the numbers of cosets of sizes ℓ and m . Let $(T_t)_{0 \leq t < p^s}$ be a nested half-cyclotomic chain of the form

$$T_t = \left(\bigcup_{j=1}^{a_{\ell,t}} C_{\delta_j^{(\ell)}} \right) \cup \left(\bigcup_{j=1}^{a_{m,t}} C_{\delta_j^{(m)}} \right). \quad (25)$$

Let C_{HC} be the associated repeated-root code. Then $\dim_{\mathbb{F}_q} C_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} (\ell a_{\ell,t} + m a_{m,t})$. Moreover, if $d_H(C(T_t)) \geq D_t$ ($0 \leq t < p^s$) are certified simple-root layer bounds, then $d_H(C_{\text{HC}}) \geq \min_{0 \leq t < p^s} P_t D_t$.

Proof. For fixed t , the set T_t contains $a_{\ell,t}$ cosets of size ℓ and $a_{m,t}$ cosets of size m . Since distinct cyclotomic cosets are disjoint, $|T_t| = \ell a_{\ell,t} + m a_{m,t}$. By Theorem 5, $\dim_{\mathbb{F}_q} C_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} |T_t|$. Substituting the expression for $|T_t|$ gives the stated dimension formula.

For the distance, Theorem 6 gives $d_H(C_{\text{HC}}) \geq \min_{0 \leq t < p^s} P_t D_t$ whenever $D_t \leq d_H(C(T_t))$ for every t . This is exactly the hypothesis. $\square$

We next investigate when a constacyclic code can be transformed into a cyclic code by a coordinatewise monomial map, allowing cyclic-code arguments to be applied after an explicit equivalence.

Lemma 2. Let $N \geq 1$, and let $\lambda \in \mathbb{F}_q^*$. If there exists $\eta \in \mathbb{F}_q^*$ such that $\eta^N = \lambda$, then every λ -constacyclic code of length N over $\mathbb{F}_q$ is monomially equivalent to a cyclic code of length N . In the present setting, where $N = np^s$ and $\lambda = \lambda_0^{p^s}$, this condition is equivalent to the existence of $\eta \in \mathbb{F}_q^*$ such that $\eta^n = \lambda_0$.

Proof. Assume $\eta^N = \lambda$. Define $\Phi : \mathbb{F}_q^N \rightarrow \mathbb{F}_q^N$ by $\Phi(c_0, c_1, \dots, c_{N-1}) = (c_0, \eta c_1, \eta^2 c_2, \dots, \eta^{N-1} c_{N-1})$. This is a monomial transformation, hence an $\mathbb{F}_q$ -linear isometry.

Let $\tau_\lambda(c_0, c_1, \dots, c_{N-1}) = (\lambda c_{N-1}, c_0, c_1, \dots, c_{N-2})$ be the λ -constacyclic shift, and let $\tau_1(d_0, d_1, \dots, d_{N-1}) = (d_{N-1}, d_0, d_1, \dots, d_{N-2})$ be the cyclic shift. For $c = (c_0, \dots, c_{N-1})$, one computes

$$\begin{aligned} \Phi(\tau_\lambda(c)) &= (\lambda c_{N-1}, \eta c_0, \eta^2 c_1, \dots, \eta^{N-1} c_{N-2}), \\ \eta \tau_1(\Phi(c)) &= (\eta^N c_{N-1}, \eta c_0, \eta^2 c_1, \dots, \eta^{N-1} c_{N-2}) \\ &= (\lambda c_{N-1}, \eta c_0, \eta^2 c_1, \dots, \eta^{N-1} c_{N-2}). \end{aligned}$$

Thus $\Phi \circ \tau_\lambda = \eta \tau_1 \circ \Phi$. If $\mathcal{C}$ is λ -constacyclic, then $\tau_\lambda(\mathcal{C}) \subseteq \mathcal{C}$. Applying Φ , we get $\eta \tau_1(\Phi(\mathcal{C})) \subseteq \Phi(\mathcal{C})$. Since $\eta \neq 0$ and $\Phi(\mathcal{C})$ is linear, this implies $\tau_1(\Phi(\mathcal{C})) \subseteq \Phi(\mathcal{C})$. Hence $\Phi(\mathcal{C})$ is cyclic. Therefore $\mathcal{C}$ is monomially equivalent to a cyclic code.

Now let $N = np^s$ and $\lambda = \lambda_0^{p^s}$. If $\eta^n = \lambda_0$, then $\eta^N = (\eta^n)^{p^s} = \lambda_0^{p^s} = \lambda$. Conversely, if $\eta^N = \lambda$, then $(\eta^n)^{p^s} = \lambda_0^{p^s}$. The Frobenius map $x \mapsto x^{p^s}$ is injective on $\mathbb{F}_q$, so $\eta^n = \lambda_0$. This proves the equivalence. $\square$

The following families are now verified through Lemma 1. Thus the two-coset-size condition and the numbers N_ℓ, N_m are not imposed as independent assumptions; they are obtained by computing the fixed-position numbers $f(u) = |F(u)|$ and then applying the Möbius formula in Lemma 1.

6.3. Family I: Lengths of the Form $nr = (q^\rho - 1)/\sigma$

This subsection treats the first arithmetic length family, where $nr = (q^\rho - 1)/\sigma$, and records the resulting cyclotomic-coset structure needed to construct the corresponding repeated-root half-cyclotomic constacyclic codes.

Theorem 10. Let ρ be an auxiliary prime and let σ be a positive integer such that $nr = \frac{q^\rho - 1}{\sigma}$. Assume $\text{ord}_{nr}(q) = \rho$. Put $A_1 = \frac{q-1}{r}$, $d_1 = \gcd(n, A_1)$. Assume also that $\gcd(r, n/d_1) = 1$. Then the q -cyclotomic cosets in $Z_{n,r}$ have exactly the two sizes 1 and ρ , with $N_1 = d_1 = \gcd\left(n, \frac{q-1}{r}\right)$, $N_\rho = \frac{n-N_1}{\rho}$. Consequently, if $(T_t)_{0 \leq t < p^s}$ is a nested half-cyclotomic chain

$$T_t = \left(\bigcup_{j=1}^{a_{1,t}} C_{\delta_j^{(1)}} \right) \cup \left(\bigcup_{j=1}^{a_{\rho,t}} C_{\delta_j^{(\rho)}} \right), \quad (26)$$

and if C_{HC} is the associated repeated-root code, then $\dim_{\mathbb{F}_q} C_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} (a_{1,t} + \rho a_{\rho,t})$. For any certified layer bounds D_t , $d_H(C_{\text{HC}}) \geq \min_{0 \leq t < p^s} P_t D_t$. If the layer is the upper half-cyclotomic set, or the lower half-cyclotomic set under the hypotheses of Theorem 1, one may take $D_t = \left\lfloor \frac{qN_\rho}{2(q-1)} \right\rfloor$. A verified consecutive-run bound may be used instead when it is stronger. If $T_t = \emptyset$, take $D_t = 1$, and if $T_t = Z_{n,r}$, take $D_t = +\infty$. If λ has an np^s -th root in $\mathbb{F}_q$, then the codes are monomially equivalent to cyclic codes.

Proof. Since $\text{ord}_{nr}(q) = \rho$, we have $q^\rho \equiv 1 \pmod{nr}$, and ρ is the least positive exponent with this property. Let $i \in Z_{n,r}$, and put $d = \gcd(nr, i)$. By Lemma 1, $|C_i| = \text{ord}_{nr/d}(q)$. Because nr/d divides nr , the order of q modulo nr/d divides the order of q modulo nr , namely ρ . Since ρ is prime, $|C_i| \in \{1, \rho\}$. Thus only the sizes 1 and ρ can occur.

It remains to count the size-1 cosets. A coset has size 1 precisely when its element is fixed by multiplication by q . Hence the number of elements lying in size-1 cosets is $|F(1)|$. By Lemma 1, $|F(1)| = d_1$ because $d_1 = \gcd(n, (q-1)/r)$ and $\gcd(r, n/d_1) = 1$. Since a size-1 coset contains exactly one element, the number of such cosets is $N_1 = d_1 = \gcd\left(n, \frac{q-1}{r}\right)$. The remaining $n - N_1$ elements of $Z_{n,r}$ lie in cosets of size ρ . Therefore $N_\rho = \frac{n-N_1}{\rho}$. Now the repeated-root parameter formulas follow from Proposition 4 with $(\ell, m) = (1, \rho)$. Indeed, each layer has size $|T_t| = a_{1,t} + \rho a_{\rho,t}$, so $\dim_{\mathbb{F}_q} C_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} (a_{1,t} + \rho a_{\rho,t})$. The distance estimate follows from the same proposition. The stated choices of D_t are certified layer bounds in the corresponding cases: the displayed half-cyclotomic bound for admissible nontrivial layers, the value 1 for the full-space layer $C(\emptyset) = \mathbb{F}_q^n$, and $+\infty$ for the zero layer $C(Z_{n,r}) = \{0\}$. Finally, the monomial-equivalence statement follows from Lemma 2. $\square$

6.4. Family II: Lengths of the Form $n = \Phi_{\rho^b}(q)$

This subsection treats the second arithmetic length family, where $n = \Phi_{\rho^b}(q)$, and uses the identity $\Phi_{\rho^b}(q) = \frac{q^{\rho^b} - 1}{q^{\rho^{b-1}} - 1}$ to describe the associated cyclotomic-coset distribution for the repeated-root construction.

Theorem 11. Assume $b \geq 1$, $\rho \nmid q-1$, $n = \Phi_{\rho^b}(q)$. Assume further that $\text{ord}_{nr}(q) = \rho^b$, and that for every proper divisor u of ρ^b , $|F(u)| = 1$. Then the q -cyclotomic cosets in $Z_{n,r}$ have exactly the two sizes 1 and ρ^b , with $N_1 = 1$, $N_{\rho^b} = \frac{n-1}{\rho^b}$. Consequently, for every nested half-cyclotomic chain

$$T_t = \left(\bigcup_{j=1}^{a_{1,t}} C_{\delta_j^{(1)}} \right) \cup \left(\bigcup_{j=1}^{a_{\rho^b,t}} C_{\delta_j^{(\rho^b)}} \right), \quad (27)$$

the associated repeated-root code satisfies $\dim_{\mathbb{F}_q} \mathcal{C}_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} (a_{1,t} + \rho^b a_{\rho^b,t})$, and, for any certified layer bounds D_t , $d_{\text{H}}(\mathcal{C}_{\text{HC}}) \geq \min_{0 \leq t < p^s} P_t D_t$. If the layer is the upper half-cyclotomic set, or the lower half-cyclotomic set under the hypotheses of Theorem 1, one may take $D_t = \left\lfloor \frac{qN_{\rho^b}}{2(q-1)} \right\rfloor$. A verified consecutive-run bound may be used instead when it is stronger. If $T_t = \emptyset$, take $D_t = 1$, and if $T_t = Z_{n,r}$, take $D_t = +\infty$. If λ has an np^s -th root in $\mathbb{F}_q$, then the codes are monomially equivalent to cyclic codes.

Proof. Since $\text{ord}_{nr}(q) = \rho^b$, every coset size divides ρ^b . Therefore, the possible coset sizes are powers of ρ : $1, \rho, \rho^2, \dots, \rho^b$. By hypothesis, for every proper divisor u of ρ^b , the fixed set $F(u)$ has cardinality 1. In particular, $|F(1)| = 1$. Thus exactly one element of $Z_{n,r}$ lies in a coset of size 1, and therefore $N_1 = 1$. Now suppose that there were a coset of intermediate size ρ^j with $1 \leq j < b$. Then all elements of this coset would be fixed by multiplication by q^{ρ^j} , so $F(\rho^j)$ would contain the one element from the size-1 coset together with all ρ^j elements of this intermediate coset. Hence $|F(\rho^j)| \geq 1 + \rho^j > 1$, contradicting the hypothesis $|F(\rho^j)| = 1$. Therefore no intermediate sizes occur. Since $\text{ord}_{nr}(q) = \rho^b$, at least one coset has size ρ^b . Hence the only coset sizes are 1 and ρ^b .

The single size-1 coset accounts for one element of $Z_{n,r}$. The remaining $n - 1$ elements are partitioned into cosets of size ρ^b , so $N_{\rho^b} = \frac{n-1}{\rho^b}$. The repeated-root dimension and distance formulas now follow from Proposition 4 with $(\ell, m) = (1, \rho^b)$. The validity of the displayed choices of D_t follows as in the previous theorem. The monomial-equivalence statement follows from Lemma 2. $\square$

Theorem 12. Assume $b \geq 1$, $\rho \mid \frac{q-1}{r}$, $n = \Phi_{\rho^b}(q)$. Assume further that $\text{ord}_{nr}(q) = \rho^b$, and that for every proper divisor u of ρ^b , $|F(u)| = \rho$. Then the q -cyclotomic cosets in $Z_{n,r}$ have exactly the two sizes 1 and ρ^b , with $N_1 = \rho$, $N_{\rho^b} = \frac{n-\rho}{\rho^b}$. Consequently, for every nested half-cyclotomic chain

$$T_t = \left(\bigcup_{j=1}^{a_{1,t}} \mathcal{C}_{\delta_j^{(1)}} \right) \cup \left(\bigcup_{j=1}^{a_{\rho^b,t}} \mathcal{C}_{\delta_j^{(\rho^b)}} \right), \quad (28)$$

the associated repeated-root code satisfies $\dim_{\mathbb{F}_q} \mathcal{C}_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} (a_{1,t} + \rho^b a_{\rho^b,t})$, and, for any certified layer bounds D_t , $d_{\text{H}}(\mathcal{C}_{\text{HC}}) \geq \min_{0 \leq t < p^s} P_t D_t$. If the layer is the upper half-cyclotomic set, or the lower half-cyclotomic set under the hypotheses of Theorem 1, one may take $D_t = \left\lfloor \frac{qN_{\rho^b}}{2(q-1)} \right\rfloor$. A verified consecutive-run bound may be used instead when it is stronger. If $T_t = \emptyset$, take $D_t = 1$, and if $T_t = Z_{n,r}$, take $D_t = +\infty$.

Proof. Since $\text{ord}_{nr}(q) = \rho^b$, every coset size divides ρ^b . Thus every coset size is one of $1, \rho, \rho^2, \dots, \rho^b$. The hypothesis $|F(1)| = \rho$ says that exactly ρ elements are fixed by multiplication by q . These elements are precisely the elements lying in cosets of size 1. Hence $N_1 = \rho$. For every proper divisor u of ρ^b , we have $|F(u)| = \rho$. Thus no new elements become fixed by q^u beyond those already fixed by q . If an intermediate coset of size ρ^j , $1 \leq j < b$, existed, then its ρ^j elements would be contained in $F(\rho^j)$, and therefore $|F(\rho^j)| \geq \rho + \rho^j > \rho$, contradicting the hypothesis. Hence no intermediate sizes occur. The only possible sizes are therefore 1 and ρ^b .

The ρ size-1 cosets account for ρ elements of $Z_{n,r}$. The remaining $n - \rho$ elements are partitioned into cosets of size ρ^b , so $N_{\rho^b} = \frac{n-\rho}{\rho^b}$. The repeated-root parameter formulas follow from Proposition 4 with $(\ell, m) = (1, \rho^b)$. This gives both the stated dimension and the lower bound $d_{\text{H}}(\mathcal{C}_{\text{HC}}) \geq \min_{0 \leq t < p^s} P_t D_t$. The choices $D_t = 1$ and $D_t = +\infty$ for empty and full defining sets come from $C(\emptyset) = \mathbb{F}_q^n$ and $C(Z_{n,r}) = \{0\}$, respectively. $\square$

Theorem 13. Assume $b \geq 1$, $\rho \mid q-1$, $\rho \nmid \frac{q-1}{r}$, $n = \Phi_{\rho^b}(q)$, and assume either $b \geq 2$ or ρ is odd. Assume further that $\text{ord}_{nr}(q) = \rho^b$, that $|F(1)| = 0$, and that for every proper divisor u of ρ^b with $u > 1$, $|F(u)| = \rho$. Then the q -cyclotomic cosets in $Z_{n,r}$ have exactly the two sizes ρ and ρ^b , with $N_\rho = 1$, $N_{\rho^b} = \frac{n-\rho}{\rho^b}$. Consequently, for every nested half-cyclotomic chain

$$T_t = \left(\bigcup_{j=1}^{a_{\rho,t}} C_{\delta_j^{(\rho)}} \right) \cup \left(\bigcup_{j=1}^{a_{\rho^b,t}} C_{\delta_j^{(\rho^b)}} \right), \quad (29)$$

the associated repeated-root code satisfies

$$\dim_{\mathbb{F}_q} \mathcal{C}_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} (\rho a_{\rho,t} + \rho^b a_{\rho^b,t}), \quad (30)$$

and, for any certified layer bounds D_t , $d_{\text{H}}(\mathcal{C}_{\text{HC}}) \geq \min_{0 \leq t < p^s} P_t D_t$. If the layer is the upper half-cyclotomic set, or the lower half-cyclotomic set under the hypotheses of Theorem 1, one may take $D_t = \left\lfloor \frac{qN_{\rho^b}}{2(q-1)} \right\rfloor$. A verified consecutive-run bound may be used instead when it is stronger. If $T_t = \emptyset$, take $D_t = 1$, and if $T_t = Z_{n,r}$, take $D_t = +\infty$. Monomial equivalence to cyclic codes holds exactly when λ has an np^s -th root in $\mathbb{F}_q$.

Proof. Again, $\text{ord}_{nr}(q) = \rho^b$ implies that all coset sizes divide ρ^b . The hypothesis $|F(1)| = 0$ says that no element is fixed by multiplication by q , so there are no cosets of size 1.

For every proper divisor $u > 1$ of ρ^b , the fixed set has cardinality ρ . In particular, for $u = \rho$, the set $F(\rho)$ contains exactly ρ elements. Since there are no size-1 cosets, these ρ elements must form a single coset of size ρ . Hence $N_\rho = 1$. If an intermediate coset of size ρ^j , $2 \leq j < b$, existed, then its ρ^j elements would be contained in $F(\rho^j)$, together with the ρ elements of the size- ρ coset. This would give $|F(\rho^j)| \geq \rho + \rho^j > \rho$, contradicting the hypothesis $|F(\rho^j)| = \rho$. Hence no intermediate sizes occur. Therefore the only coset sizes are ρ and ρ^b .

The unique size- ρ coset accounts for ρ elements. The remaining $n - \rho$ elements are partitioned into cosets of size ρ^b , so $N_{\rho^b} = \frac{n-\rho}{\rho^b}$. The repeated-root formulas now follow from Proposition 4 with $(\ell, m) = (\rho, \rho^b)$. The empty and full defining-set conventions give $D_t = 1$ and $D_t = +\infty$. Finally, the monomial-equivalence statement is precisely Lemma 2. $\square$

6.5. Family III: Lengths of the Form $n = \Phi_{\rho_1 \rho_2}(q)$

Theorems 11–13 give the three prime-power cyclotomic length families used before the product-family case.

Theorem 14. Let $\rho_1 < \rho_2$ be distinct auxiliary primes, and let $n = \Phi_{\rho_1 \rho_2}(q)$. Assume $\text{ord}_{nr}(q) = \rho_1 \rho_2$. Assume further that $|F(1)| = 1$, $|F(\rho_1)| = 1$, $|F(\rho_2)| = 1$. Then the q -cyclotomic cosets in $Z_{n,r}$ have exactly the two sizes 1 and $\rho_1 \rho_2$, with $N_1 = 1$, $N_{\rho_1 \rho_2} = \frac{n-1}{\rho_1 \rho_2}$. Consequently, for every nested half-cyclotomic chain

$$T_t = \left(\bigcup_{j=1}^{a_{1,t}} C_{\delta_j^{(1)}} \right) \cup \left(\bigcup_{j=1}^{a_{\rho_1 \rho_2,t}} C_{\delta_j^{(\rho_1 \rho_2)}} \right), \quad (31)$$

the associated repeated-root code satisfies

$$\dim_{\mathbb{F}_q} \mathcal{C}_{\text{HC}} = np^s - \sum_{t=0}^{p^s-1} (a_{1,t} + \rho_1 \rho_2 a_{\rho_1 \rho_2,t}), \quad (32)$$

and, for any certified layer bounds D_t , $d_H(C_{HC}) \geq \min_{0 \leq t < p^s} P_t D_t$. If the layer is the upper half-cyclotomic set, or the lower half-cyclotomic set under the hypotheses of Theorem 1, one may take $D_t = \left\lfloor \frac{q^{N_{\rho_1 \rho_2}}}{2(q-1)} \right\rfloor$. A verified consecutive-run bound may be used instead when it is stronger. If $T_t = \emptyset$, take $D_t = 1$, and if $T_t = Z_{n,r}$, take $D_t = +\infty$. If λ has an np^s -th root in $\mathbb{F}_q$, then the codes are monomially equivalent to cyclic codes.

Proof. Since $\text{ord}_{nr}(q) = \rho_1 \rho_2$, every coset size divides $\rho_1 \rho_2$. Hence the possible sizes are $1, \rho_1, \rho_2, \rho_1 \rho_2$. The hypothesis $|F(1)| = 1$ gives exactly one element fixed by q , hence one coset of size 1: $N_1 = 1$. If a coset of size ρ_1 existed, then all of its ρ_1 elements would lie in $F(\rho_1)$, together with the element from the size-1 coset. Thus $|F(\rho_1)| \geq 1 + \rho_1 > 1$, contradicting the hypothesis $|F(\rho_1)| = 1$. Hence no coset of size ρ_1 exists. The same argument using $|F(\rho_2)| = 1$ excludes cosets of size ρ_2 . Therefore the only coset sizes are 1 and $\rho_1 \rho_2$.

The single size-1 coset accounts for one element of $Z_{n,r}$. The remaining $n - 1$ elements are partitioned into cosets of size $\rho_1 \rho_2$, so $N_{\rho_1 \rho_2} = \frac{n-1}{\rho_1 \rho_2}$. The repeated-root parameter formulas follow from Proposition 4 with $(\ell, m) = (1, \rho_1 \rho_2)$. The monomial-equivalence statement follows from Lemma 2. $\square$

7. Lifting Optimal and Best-Known Simple-Root Codes

We now discuss how simple-root constacyclic codes can be lifted to repeated-root length np^s . There are two natural procedures. The first is the uniform lift, where one raises the whole simple-root generator to a single exponent. This construction is simple, but it has a serious distance limitation: after the lifting exponent is passed, the associated layers become the full space $\mathbb{F}_q^n$, whose minimum distance is 1. The second is the nonuniform lift, where different irreducible factors are assigned different multiplicities. This is the more flexible construction, and it is the one naturally compatible with the layer method developed above.

Definition 9. Let $D = \langle g(x) \rangle \subseteq \mathbb{F}_q[x] / \langle x^n - \lambda_0 \rangle$ be a simple-root λ_0 -constacyclic code, where $g(x) = \prod_{i \in A} M_i(x)$ for some union of q -cyclotomic coset leaders $A \subseteq \Gamma$.

For $1 \leq e \leq p^s$, the uniform repeated-root lift of D with multiplicity e is

$$\text{RR}_e(D) = \langle g(x)^e \rangle \subseteq \mathbb{F}_q[x] / \langle x^{np^s} - \lambda \rangle. \quad (33)$$

Equivalently, this is the repeated-root code with multiplicity profile

$$e_i = \begin{cases} e, & i \in A, \\ 0, & i \notin A. \end{cases} \quad (34)$$

More generally, a nonuniform repeated-root lift of D is obtained by assigning integers $1 \leq e_i \leq p^s$ ($i \in A$) and setting $e_i = 0$ for $i \notin A$. The resulting code is

$$\mathcal{C}(e) = \left\langle \prod_{i \in A} M_i(x)^{e_i} \right\rangle \subseteq \mathbb{F}_q[x] / \langle x^{np^s} - \lambda \rangle. \quad (35)$$

Definition 9 separates the uniform construction from the nonuniform profiles treated in Theorem 15.

Proposition 5. Let D be a simple-root $[n, k, d]$ λ_0 -constacyclic code over $\mathbb{F}_q$, with generator polynomial $g(x)$. Let

$$\mathcal{C} = \text{RR}_e(D) = \langle g(x)^e \rangle \subseteq \mathbb{F}_q[x] / \langle x^{np^s} - \lambda \rangle, \quad 1 \leq e \leq p^s. \quad (36)$$

Then $\dim_{\mathbb{F}_q} \mathcal{C} = np^s - e(n - k)$. Moreover, the associated simple-root layers are

$$C_t = \begin{cases} D, & 0 \leq t < e, \\ \mathbb{F}_q^n, & e \leq t < p^s. \end{cases} \quad (37)$$

Consequently, $d_H(\mathcal{C}) = \min\{\min_{0 \leq t < e} P_t d, \min_{e \leq t < p^s} P_t\}$, where the second minimum is interpreted as $+\infty$ if $e = p^s$. Since $P_0 = 1$, this can also be written as

$$d_H(\text{RR}_e(D)) = \begin{cases} \min\{d, \mathcal{M}_s(e)\}, & 1 \leq e < p^s, \\ d, & e = p^s, \end{cases} \quad (38)$$

where $\mathcal{M}_s(e) = \min_{e \leq t < p^s} P_t$.

Proof. Since D is an $[n, k, d]$ simple-root constacyclic code, its generator polynomial $g(x)$ has degree $\deg g = n - k$. The uniform repeated-root lift is generated by $g(x)^e$. Therefore $\deg g(x)^e = e \deg g = e(n - k)$. Because $g(x)^e \mid x^{np^s} - \lambda$, the ideal $\langle g(x)^e \rangle \subseteq \mathbb{F}_q[x] / \langle x^{np^s} - \lambda \rangle$ has dimension $np^s - \deg g(x)^e = np^s - e(n - k)$. This proves the dimension formula.

We now identify the layers. Write $g(x) = \prod_{i \in A} M_i(x)$. The multiplicity profile of $\text{RR}_e(D)$ is

$$e_i = \begin{cases} e, & i \in A, \\ 0, & i \notin A. \end{cases}$$

By the definition of the associated layers, $T_t(e) = \bigcup_{\substack{i \in \Gamma \\ e_i \geq t}} C_i$. If $0 \leq t < e$, then $e_i = e > t$ exactly for $i \in A$. Hence $T_t(e) = \bigcup_{i \in A} C_i$, and therefore $C_t = C(T_t(e)) = D$. If $e \leq t < p^s$, then no exponent e_i is larger than t . Hence $T_t(e) = \emptyset$, and so $C_t = C(\emptyset) = \mathbb{F}_q^n$. Thus the layers are exactly as stated.

Applying the repeated-root distance reduction theorem gives $d_H(\mathcal{C}) = \min_{0 \leq t < p^s} P_t d_H(C_t)$. For $0 \leq t < e$, we have $C_t = D$, so $d_H(C_t) = d$. For $e \leq t < p^s$, we have $C_t = \mathbb{F}_q^n$, so $d_H(C_t) = 1$. Therefore $d_H(\mathcal{C}) = \min\{\min_{0 \leq t < e} P_t d, \min_{e \leq t < p^s} P_t\}$. If $e = p^s$, the second range is empty and its minimum is taken to be $+\infty$. Since $0 < e$, the first range contains $t = 0$, and $P_0 = 1$. As all $P_t \geq 1$, we have $\min_{0 \leq t < e} P_t d = d$. Hence, for $e < p^s$, $d_H(\text{RR}_e(D)) = \min\{d, \mathcal{M}_s(e)\}$, where $\mathcal{M}_s(e) = \min_{e \leq t < p^s} P_t$. For $e = p^s$, only the first part remains, and the distance is d . $\square$

Remark 2. The preceding proposition shows that a uniform lift with $e < p^s$ always has $d_H(\text{RR}_e(D)) \leq \mathcal{M}_s(e)$. Thus even if D is optimal or best-known as a simple-root code, the uniform lift may have small repeated-root distance because the layers $t \geq e$ are the full space $\mathbb{F}_q^n$. This is the main reason to use nonuniform multiplicity profiles: they allow one to keep nontrivial defining sets in higher layers and thereby avoid unprotected full-space layers.

Proposition 5 and Remark 2 motivate the nonuniform lifting theorem below.

Theorem 15. Let

$$\mathcal{C}(e) = \left\langle \prod_{i \in \Gamma} M_i(x)^{e_i} \right\rangle \subseteq \mathbb{F}_q[x] / \langle x^{np^s} - \lambda \rangle \quad (39)$$

be a repeated-root code with multiplicity profile $e : \Gamma \rightarrow \{0, 1, \dots, p^s\}$. For $0 \leq t < p^s$, let $C_t(e) = C(T_t(e))$ be its associated simple-root layer, and put $d_t = d_H(C_t(e))$. Then $\dim_{\mathbb{F}_q} \mathcal{C}(e) = np^s - \sum_{i \in \Gamma} e_i |C_i|$, and $d_H(\mathcal{C}(e)) = \min_{0 \leq t < p^s} P_t d_t$. Thus, if $\mathcal{C}(e) \neq \{0\}$, its parameters are $[np^s, np^s - \sum_{i \in \Gamma} e_i |C_i|, \min_{0 \leq t < p^s} P_t d_t]_q$.

Proof. The dimension formula is exactly Proposition 2. Indeed, the generator of $\mathcal{C}(e)$ is $g_e(x) = \prod_{i \in \Gamma} M_i(x)^{e_i}$. Since $\deg M_i = |C_i|$, we have $\deg g_e = \sum_{i \in \Gamma} e_i \deg M_i = \sum_{i \in \Gamma} e_i |C_i|$. Therefore $\dim_{\mathbb{F}_q} \mathcal{C}(e) = np^s - \deg g_e = np^s - \sum_{i \in \Gamma} e_i |C_i|$. For the distance, Theorem 4 gives the exact identity $d_H(\mathcal{C}(e)) = \min_{0 \leq t < p^s} P_t d_H(C_t(e))$. By definition $d_t = d_H(C_t(e))$. Hence $d_H(\mathcal{C}(e)) = \min_{0 \leq t < p^s} P_t d_t$. Combining the dimension and distance formulas gives the stated parameters. $\square$

Corollary 2. Assume that the simple-root layers $C_t(e)$ have known exact parameters $C_t(e) = [n, k_t, d_t]_q$ for $0 \leq t < p^s$. Then the repeated-root lift $\mathcal{C}(e)$ has exact dimension and distance $\dim_{\mathbb{F}_q} \mathcal{C}(e) = np^s - \sum_{t=0}^{p^s-1} |T_t(e)|$ and $d_H(\mathcal{C}(e)) = \min_{0 \leq t < p^s} P_t d_t$. Equivalently, $\dim_{\mathbb{F}_q} \mathcal{C}(e) = np^s - \sum_{i \in \Gamma} e_i |C_i|$.

Proof. The layer form of the dimension formula gives $\dim_{\mathbb{F}_q} \mathcal{C}(e) = np^s - \sum_{t=0}^{p^s-1} |T_t(e)|$. The profile form gives the equivalent expression $\dim_{\mathbb{F}_q} \mathcal{C}(e) = np^s - \sum_{i \in \Gamma} e_i |C_i|$. Finally, the distance formula follows from Theorem 15. $\square$

Corollary 2 gives the form used when all layer parameters are known exactly.

The following definition separates proved optimality from database-based comparisons.

Definition 10. Let $U_q(N, K)$ be a proved upper bound for the minimum distance of a q -ary linear code of length N and dimension K . A code $[N, K, D]_q$ is called distance-optimal relative to U_q if $D = U_q(N, K)$. It is called almost distance-optimal relative to U_q if $D = U_q(N, K) - 1$.

If $B_q(N, K)$ denotes the largest minimum distance appearing in a fixed table or database of known q -ary linear codes, then a code $[N, K, D]_q$ is called best-known relative to that table if $D = B_q(N, K)$, and record-breaking relative to that table if $D > B_q(N, K)$. The latter two terms refer only to comparison with the chosen table or database and do not imply optimality unless a proved upper bound is also met.

Definition 10 supplies the benchmark language used in Theorem 16.

Theorem 16. Let $\mathcal{C}(e)$ be a repeated-root lift whose associated simple-root layers have known exact distances $d_t = d_H(C_t(e))$, $0 \leq t < p^s$. Put $N = np^s$, $K = np^s - \sum_{i \in \Gamma} e_i |C_i|$, and $D = \min_{0 \leq t < p^s} P_t d_t$. Then $\mathcal{C}(e)$ has exact parameters $[N, K, D]_q$. Consequently:

- If $D = U_q(N, K)$ for a proved upper bound $U_q(N, K)$, then $\mathcal{C}(e)$ is distance-optimal relative to U_q .
- If $D = U_q(N, K) - 1$, then $\mathcal{C}(e)$ is almost distance-optimal relative to U_q .
- If $D = B_q(N, K)$ for a chosen table of known q -ary linear codes, then $\mathcal{C}(e)$ is best-known relative to that table.
- If $D > B_q(N, K)$, then $\mathcal{C}(e)$ is record-breaking relative to that table.

Proof. By Theorem 15, the dimension of $\mathcal{C}(e)$ is $K = np^s - \sum_{i \in \Gamma} e_i |C_i|$, and its exact minimum distance is $D = \min_{0 \leq t < p^s} P_t d_t$. Thus $\mathcal{C}(e)$ has parameters $[N, K, D]_q$ with $N = np^s$.

If $D = U_q(N, K)$, then no q -ary linear code of length N and dimension K can have minimum distance larger than D , by the definition of the upper bound U_q . Hence $\mathcal{C}(e)$ is distance-optimal relative to that bound. If $D = U_q(N, K) - 1$, then its distance is exactly one below this proved upper bound, so it is almost distance-optimal relative to U_q . If

$D = B_q(N, K)$ for a chosen table of known codes, then $\mathcal{C}(e)$ matches the best distance recorded in that table and is best-known relative to it. If $D > B_q(N, K)$, then $\mathcal{C}(e)$ improves the distance recorded in that table, and is therefore record-breaking relative to that table. $\square$

Problem 1. Fix a target dimension K . Among all half-cyclotomic multiplicity profiles $e : \Gamma \rightarrow \{0, 1, \dots, p^s\}$ satisfying $np^s - \sum_{i \in \Gamma} e_i |C_i| \geq K$, determine or estimate $\max_e \min_{0 \leq t < p^s} P_t d_H(C_t(e))$. If exact layer distances are not known, replace $d_H(C_t(e))$ by a certified layer lower bound $\Delta_{HC}(T_t(e))$, and study $\max_e \min_{0 \leq t < p^s} P_t \Delta_{HC}(T_t(e))$.

We now give a precise solution to the half-cyclotomic multiplicity optimization problem. The main point is that the problem is finite and hence has an optimum. In general, however, there is no universal greedy rule: the factors P_t are not monotone in t , and the nesting condition couples the layers. The correct general formulation is a finite integer optimization, which can be solved exactly by dynamic programming.

Let $B = np^s - K$. By the layer form of the dimension formula, $\dim_{\mathbb{F}_q} \mathcal{C}(e) = np^s - \sum_{t=0}^{p^s-1} |T_t(e)|$, the constraint $np^s - \sum_{i \in \Gamma} e_i |C_i| \geq K$ is equivalent to $\sum_{t=0}^{p^s-1} |T_t(e)| \leq B$. Thus the optimization problem may be written directly in terms of nested half-cyclotomic chains $T_{p^s-1} \subseteq T_{p^s-2} \subseteq \dots \subseteq T_1 \subseteq T_0$ satisfying $\sum_{t=0}^{p^s-1} |T_t| \leq B$. For each layer T , write $\delta(T)$ for either the exact distance $d_H(C(T))$ or for a certified lower bound such as $\Delta_{HC}(T)$. If this certified bound is not monotone with respect to inclusion, we replace it, whenever monotonicity is needed, by the certified monotone envelope defined in Lemma 3. The objective is then $\max \min_{0 \leq t < p^s} P_t \delta(T_t)$, where the maximum is taken over all feasible nested half-cyclotomic chains.

Lemma 3. Let $T \subseteq T'$ be unions of q -cyclotomic cosets in $\mathbb{Z}_{n,r}$. Then $C(T') \subseteq C(T)$, and hence $d_H(C(T)) \leq d_H(C(T'))$. Thus the true minimum distance is monotone with respect to enlargement of the defining set. A lower-bound function $\delta(T)$ is not automatically monotone; when monotonicity is required, we use either a certified monotone bound or its certified monotone envelope

$$\widehat{\delta}(T) = \max\{\delta(U) : U \subseteq T, U \text{ is a union of } q\text{-cyclotomic cosets}\}.$$

If every $\delta(U)$ is a certified lower bound for $d_H(C(U))$, then $\widehat{\delta}(T)$ is a certified lower bound for $d_H(C(T))$ and satisfies $T \subseteq T' \Rightarrow \widehat{\delta}(T) \leq \widehat{\delta}(T')$.

Proof. The simple-root code $C(T)$ is generated by $g_T(x) = \prod_{C_i \subseteq T} M_i(x)$. If $T \subseteq T'$, then every factor of $g_T(x)$ is also a factor of $g_{T'}(x)$. Therefore $g_T(x) \mid g_{T'}(x)$. In the quotient ring $\mathbb{F}_q[x]/\langle x^n - \lambda_0 \rangle$, divisibility of generators reverses inclusion of the corresponding principal ideals: $\langle g_{T'}(x) \rangle \subseteq \langle g_T(x) \rangle$. Thus $C(T') \subseteq C(T)$. Since the nonzero codewords of $C(T')$ form a subset of the nonzero codewords of $C(T)$, we have

$$d_H(C(T)) = \min_{0 \neq c \in C(T)} \text{wt}_H(c) \leq \min_{0 \neq c \in C(T')} \text{wt}_H(c) = d_H(C(T')).$$

Now assume that each $\delta(U)$ is a certified lower bound for $d_H(C(U))$. If $U \subseteq T$, then the first part gives $d_H(C(U)) \leq d_H(C(T))$. Hence $\delta(U) \leq d_H(C(U)) \leq d_H(C(T))$, and taking the maximum over all such U shows that $\widehat{\delta}(T) \leq d_H(C(T))$. Finally, if $T \subseteq T'$, then the set of admissible $U \subseteq T$ is contained in the set of admissible $U \subseteq T'$, so $\widehat{\delta}(T) \leq \widehat{\delta}(T')$. $\square$

Lemma 3 is the monotonicity input for the maximal-chain statement below.

Proposition 6. Assume $0 \leq K \leq np^s$. Then the optimization problem $\max_e \min_{0 \leq t < p^s} P_t d_H(C_t(e))$ over all feasible half-cyclotomic multiplicity profiles has a solution. The same conclusion holds if $d_H(C_t(e))$ is replaced by any certified layer bound $\Delta_{HC}(T_t(e))$.

Proof. The set Γ of coset leaders is finite. A multiplicity profile is a map $e : \Gamma \rightarrow \{0, 1, \dots, p^s\}$. Hence the total number of multiplicity profiles is at most $(p^s + 1)^{|\Gamma|}$. The half-cyclotomic profiles form a subset of this finite set, and the profiles satisfying $np^s - \sum_{i \in \Gamma} e_i |C_i| \geq K$ form a finite feasible set. This feasible set is nonempty when $0 \leq K \leq np^s$, since the zero profile $e_i = 0$ has dimension np^s . The objective function is a real-valued function on this finite nonempty set. Therefore it attains its maximum. The same argument applies when exact layer distances are replaced by certified lower bounds. $\square$

Proposition 7. Assume that the layer quantity $\delta(T)$ used in the objective is monotone with respect to inclusion. Then there exists an optimal feasible nested half-cyclotomic chain $T_{p^s-1} \subseteq \dots \subseteq T_0$ which is maximal in the following sense: no layer can be strictly enlarged while preserving the half-cyclotomic initial-segment condition, the nesting condition, and the budget condition $\sum_{t=0}^{p^s-1} |T_t| \leq B$.

Proof. By Proposition 6, optimal feasible chains exist. Among all optimal feasible chains, choose one, say $\mathcal{T} = (T_t)_{0 \leq t < p^s}$, for which the total layer size $\sum_{t=0}^{p^s-1} |T_t|$ is maximal.

We claim that $\mathcal{T}$ is maximal. Suppose it is not. Then there exists an index j and a strictly larger half-cyclotomic initial set $T'_j \supsetneq T_j$ such that replacing T_j by T'_j preserves the nesting condition and the budget condition. Let $\mathcal{T}'$ be the resulting chain. Then $\sum_{t=0}^{p^s-1} |T'_t| > \sum_{t=0}^{p^s-1} |T_t|$. For all $t \neq j$, the layer is unchanged. For $t = j$, monotonicity gives $\delta(T_j) \leq \delta(T'_j)$. Since $P_j > 0$, we also have $P_j \delta(T_j) \leq P_j \delta(T'_j)$. Therefore $\min_{0 \leq t < p^s} P_t \delta(T_t) \leq \min_{0 \leq t < p^s} P_t \delta(T'_t)$. Thus $\mathcal{T}'$ is also optimal. But $\mathcal{T}'$ has strictly larger total layer size than $\mathcal{T}$, contradicting the maximal choice of $\mathcal{T}$. Hence $\mathcal{T}$ is maximal. $\square$

Remark 3. The preceding proposition is not a greedy theorem. It only says that an optimal chain may be chosen so that it cannot be enlarged without violating the constraints. A universal greedy algorithm would require additional structure, such as concavity or submodularity of the layer bound, and such properties are not available in general. In particular, the factors P_t are not monotone in t .

Remark 3 clarifies the limitation of Proposition 7.

We now formulate an exact dynamic programme. Let $\mathcal{A}$ be the finite set parametrising all allowed half-cyclotomic layers. For $a \in \mathcal{A}$, let $T(a)$ be the corresponding layer, and define $w(a) = |T(a)|$, $\delta(a) = \delta(T(a))$. Define a partial order on $\mathcal{A}$ by $a \preceq b \iff T(a) \subseteq T(b)$. Thus a nested chain corresponds to a sequence $a_0, a_1, \dots, a_{p^s-1}$ satisfying $a_t \preceq a_{t-1}$ for $1 \leq t < p^s$.

Theorem 17. Let R be a real number. For each level t , define $\mathcal{A}_t(R) = \{a \in \mathcal{A} : P_t \delta(a) \geq R\}$. For $a \in \mathcal{A}_0(R)$, set $F_0(a) = w(a)$. For $1 \leq t < p^s$ and $a \in \mathcal{A}_t(R)$, define recursively

$$F_t(a) = w(a) + \min_{\substack{b \in \mathcal{A}_{t-1}(R) \\ a \preceq b}} F_{t-1}(b), \quad (40)$$

with the convention that the minimum is $+\infty$ if no such b exists.

Then there exists a feasible nested half-cyclotomic chain satisfying $\sum_{t=0}^{p^s-1} |T_t| \leq B$ and $\min_{0 \leq t < p^s} P_t \delta(T_t) \geq R$ if and only if $\min_{a \in \mathcal{A}_{p^s-1}(R)} F_{p^s-1}(a) \leq B$. Consequently, the optimal value of Problem 1 is the largest attainable value of R .

Proof. A nested chain is equivalent to a sequence $a_0, a_1, \dots, a_{p^s-1}$ with $a_t \preceq a_{t-1}$ for $1 \leq t < p^s$. The condition $\min_{0 \leq t < p^s} P_t \delta(T(a_t)) \geq R$ is equivalent to $a_t \in \mathcal{A}_t(R)$ for $0 \leq t < p^s$. The total layer size is $\sum_{t=0}^{p^s-1} w(a_t)$. We prove by induction on t that $F_t(a)$ is the minimum total weight of a valid partial chain $a_0, a_1, \dots, a_t$ ending at $a_t = a$, satisfying $a_u \in \mathcal{A}_u(R)$ for $0 \leq u \leq t$, and $a_u \preceq a_{u-1}$ for $1 \leq u \leq t$. For $t = 0$, such a partial chain consists only of $a_0 = a$, so its weight is $w(a)$. Therefore $F_0(a) = w(a)$. Assume the assertion holds for $t - 1$. A partial chain ending at $a_t = a$ must have previous layer $a_{t-1} = b$, where $b \in \mathcal{A}_{t-1}(R)$ and $a \preceq b$. The minimum total weight of the first t layers ending at such a b is $F_{t-1}(b)$ by the induction hypothesis. Adding the last layer contributes $w(a)$. Therefore the minimum total weight of a partial chain ending at a is exactly the recursive value defining $F_t(a)$.

At the final level $p^s - 1$, a complete chain satisfying the threshold condition has minimum possible total weight $\min_{a \in \mathcal{A}_{p^s-1}(R)} F_{p^s-1}(a)$. Therefore such a chain satisfies the budget condition if and only if this minimum is at most B . This proves the criterion. Since $\mathcal{A}$ and the set of levels are finite, the set of attainable objective values is finite. Hence the largest attainable R exists and equals the optimal value. $\square$

Theorem 17 converts the profile search into a finite dynamic programme. Let $L = p^s$ and $A = |\mathcal{A}|$. For a fixed threshold R , after the partial order on $\mathcal{A}$ is known, the recursion in Theorem 17 can be evaluated in $O(LA^2)$ time and $O(A)$ working memory. The partial order can be precomputed in $O(A^2)$ time. Since the optimal threshold belongs to the finite set $\{P_t \delta(a) : 0 \leq t < L, a \in \mathcal{A}\}$, which has size at most LA , the exact optimum can be found by sorting these candidates and using binary search over the feasibility test. Thus the dynamic programme is polynomial in the number of admissible layers A and in $L = p^s$.

Corollary 3. Assume that the cyclotomic cosets in $\mathbb{Z}_{n,r}$ have exactly two sizes $\ell < m$. Write each layer as $a_t = (a_{\ell,t}, a_{m,t})$, $0 \leq a_{\ell,t} \leq N_\ell$, $0 \leq a_{m,t} \leq N_m$. Then the optimization problem is the finite integer programme

$$\max \min_{0 \leq t < p^s} P_t \delta(a_{\ell,t}, a_{m,t})$$

subject to $a_{\ell,t+1} \leq a_{\ell,t}$, $a_{m,t+1} \leq a_{m,t}$ for $0 \leq t < p^s - 1$, and $\sum_{t=0}^{p^s-1} (\ell a_{\ell,t} + m a_{m,t}) \leq n p^s - K$. Here $\delta(a_{\ell,t}, a_{m,t})$ denotes either the exact distance of the simple-root layer or a certified lower bound.

Proof. In the two-size case, a half-cyclotomic layer is completely determined by the pair $(a_{\ell,t}, a_{m,t})$. The nesting condition is exactly the componentwise condition $a_{\ell,t+1} \leq a_{\ell,t}$, $a_{m,t+1} \leq a_{m,t}$. The weight of the layer is $|T_t| = \ell a_{\ell,t} + m a_{m,t}$. Substituting these expressions into the general optimization formulation gives the displayed finite integer programme. $\square$

Corollary 3 is the explicit integer form used in the two-size computations. In the two-coset-size case, $A = (N_\ell + 1)(N_m + 1)$. Hence the fixed-threshold dynamic-programming test has time complexity $O(p^s A^2)$, with $A = (N_\ell + 1)(N_m + 1)$.

Remark 4. In the two-coset-size case, many half-cyclotomic lower bounds are controlled by the number $a_{m,t}$ of selected large cosets in the layer T_t . Thus, as a first approximation, the repeated-root lower bound behaves like $\min_{0 \leq t < p^s} P_t a_{m,t}$. The factors P_t are usually irregular in t , because they depend on the base- p digits of t . A good multiplicity profile should therefore allocate more protection to layers where P_t is small and may allow less protection where P_t is large. The nesting condition $T_{t+1} \subseteq T_t$ imposes a strong constraint on this balancing process, since the number of selected cosets cannot increase as t increases. This interaction between the p -adic factors P_t and the nested half-cyclotomic layers is one of the main new design features in the repeated-root setting.

This principle is used in the numerical comparisons in Section 8.

8. Numerical Illustrations and Comparison

This section illustrates the practical advantage of layered multiplicity profiles. All comparisons use the same simple-root half-cyclotomic data, so any gain in dimension at the same certified lower bound comes from the repeated-root layering itself. Compared with the simple-root construction of [26], the present method does not choose only one defining set. It distributes half-cyclotomic defining sets across the p^s repeated-root levels. Compared with the distance reduction method of [27], the novelty is that the torsion layers are not treated merely as a consequence of a given repeated-root ideal; they are used as design variables.

We use the terminology “optimal”, “best-known”, and “record-breaking” in the precise sense fixed later in Definition 10: optimality refers to meeting a proved upper bound, while the latter two terms refer to comparison with specified tables or databases. The purpose of the examples is therefore structural: they show how multiplicity profiles steer dimension while keeping a certified repeated-root distance bound. We do not claim that every displayed code is optimal or best-known in Grassl’s database; when such a comparison is available, it is stated explicitly.

For a fixed repeated-root index $e_0 = p^s$, we compare four profiles:

- (i) The *uniform upper profile*, in which every layer is H^+ ;
- (ii) The *uniform lower profile*, in which every layer is H^- , whenever the lower half-cyclotomic distance bound is available;
- (iii) The *balanced mixed profile*, in which half of the layers are H^+ and half are H^- ;
- (iv) The *naive single-layer lift*, in which only the first layer is half-cyclotomic and the remaining layers are empty.

The first two profiles are uniform repeated-root lifts of the simple-root construction. The mixed profile is different: it uses multiplicity levels to steer the dimension while keeping the same certified layer-distance bound. The naive lift is included to show that high dimension alone is misleading, since empty layers may force the repeated-root distance bound to collapse.

Example 1. We give a compact detailed example based on the best-known simple-root constacyclic codes in Example 5.4 and Table 5.1 of [26]. Let $q = 4$, $p = 2$, $s = 2$, $n = 17$, $r = 3$, and $\lambda_0 \in \mathbb{F}_4^*$ of order 3. Then $\lambda = \lambda_0^4 = \lambda_0$, so $x^{68} - \lambda = (x^{17} - \lambda_0)^4$. The 4-cyclotomic cosets in $Z_{17,3} \subseteq \mathbb{Z}/51\mathbb{Z}$ are

$$\begin{aligned} C_1 &= \{1, 4, 13, 16\}, \quad C_7 = \{7, 10, 28, 40\}, \quad C_{19} = \{19, 25, 43, 49\}, \\ C_{22} &= \{22, 31, 37, 46\}, \quad C_{34} = \{34\}. \end{aligned}$$

Thus sizes are 1 and 4, with $N_1 = 1$, $N_4 = 4$. Define $H^+ = C_1 \cup C_7 \cup C_{34}$, $H^- = C_1 \cup C_7$, so $|H^+| = 9$, $|H^-| = 8$. The source gives $C(H^+) = [17, 8, 8]_4$, $C(H^-) = [17, 9, 7]_4$ (both best-known; $C(H^+)$ is non-BCH). Take the nested chain $T_0 = T_1 = H^+$, $T_2 = T_3 = H^-$. Let M_i be the irreducible factor of $x^{17} - \lambda_0$ for coset C_i . The multiplicity profile is $e_1 = e_7 = 4$, $e_{34} = 2$, $e_{19} = e_{22} = 0$. Thus $g(x) = M_1^4 M_7^4 M_{34}^2$, $\deg g = 4 \cdot 4 + 4 \cdot 4 + 2 \cdot 1 = 34$, so $\dim_{\mathbb{F}_4} C = 68 - 34 = 34$. The 2-adic factors are $(P_0, P_1, P_2, P_3) = (1, 2, 2, 4)$. Using exact layer distances 8 and 7,

$$d_H(C) = \min\{1 \cdot 8, 2 \cdot 8, 2 \cdot 7, 4 \cdot 7\} = 8.$$

Hence we obtain the exact code $[68, 34, 8]_4$. The uniform upper lift gives $[68, 32, 8]_4$ (losing 2 dimensions), while the naive single-layer lift has dimension 59 but distance collapses to 2 (Table 1).

Table 1. Repeated-root profiles from the $q = 4$, $n = 17$, $r = 3$ best-known simple-root source codes.

Profile	Layers	Dimension	Distance	Parameter	Comment
Uniform upper	$4H^+$	32	8	$[68, 32, 8]_4$	uniform lift
Balanced mixed	$2H^+ + 2H^-$	34	8	$[68, 34, 8]_4$	+2 dimensions
Uniform lower	$4H^-$	36	7	$[68, 36, 7]_4$	higher dimension
Naive single layer	$H^+ + 3\emptyset$	59	2	$[68, 59, 2]_4$	distance collapses

The numerical computations were verified with MAGMA (V2.29-8). The verification constructs the relevant q -cyclotomic cosets, forms the defining sets H^+ and H^- , computes generator degrees and dimensions from the multiplicity profiles, and evaluates the repeated-root distance bounds using Theorem 4. In Example 1, the simple-root distances were computed exactly, so the resulting repeated-root distance is exact. In the remaining larger examples, the displayed distances are certified lower bounds unless exactness is explicitly stated.

Example 2. Let $q = 2^a$ with $a \geq 1$, let $p = 2$, let $r = 1$, and take $n = q^3 - 1$. This is the family $n = (q^p - 1)/(r\sigma)$ with $\rho = 3$ and $\sigma = 1$, of the type considered in [26]. The q -cyclotomic cosets in $Z_{n,1}$ have sizes 1 and 3. The numbers of such cosets are $N_1 = q - 1$ and $N_3 = (q^3 - q)/3$. Since q is even, N_1 is odd and N_3 is even. Hence

$$|H^+| = \left\lceil \frac{N_1}{2} \right\rceil + 3 \left\lceil \frac{N_3}{2} \right\rceil = \frac{q^3}{2}, \quad |H^-| = \left\lfloor \frac{N_1}{2} \right\rfloor + 3 \left\lfloor \frac{N_3}{2} \right\rfloor = \frac{q^3 - 2}{2}. \quad (41)$$

Both H^+ and H^- satisfy the same certified simple-root distance bound, namely

$$\Delta(q) = \left\lfloor \frac{qN_3}{2(q-1)} \right\rfloor = \left\lfloor \frac{q^2(q+1)}{6} \right\rfloor. \quad (42)$$

Let $e_0 = 2^s$. For $0 \leq u \leq e_0$, define a nested profile by $T_0 = \dots = T_{u-1} = H^+$ and $T_u = \dots = T_{e_0-1} = H^-$. Since $H^- \subseteq H^+$, this is a valid nested half-cyclotomic chain. The corresponding repeated-root code $\mathcal{C}(u) \subseteq \mathbb{F}_q[x]/\langle x^{e_0n} - 1 \rangle$ has length $N = e_0(q^3 - 1)$ and dimension

$$\dim_{\mathbb{F}_q} \mathcal{C}(u) = N - u|H^+| - (e_0 - u)|H^-| = \frac{N}{2} + \frac{e_0}{2} - u.$$

Moreover, each active layer has distance at least $\Delta(q)$, and $\min_t P_t = 1$. Therefore

$$d_H(\mathcal{C}(u)) \geq \Delta(q) = \left\lfloor \frac{q^2(q+1)}{6} \right\rfloor.$$

Thus, as u varies, the dimension moves in unit steps through $\frac{N}{2} - \frac{e_0}{2} \leq \dim_{\mathbb{F}_q} \mathcal{C}(u) \leq \frac{N}{2} + \frac{e_0}{2}$, while the same certified distance lower bound is retained. The uniform upper lift is obtained when $u = e_0$, the uniform lower lift when $u = 0$, and the exact half-rate mixed code when $u = e_0/2$.

The preceding example shows the main extra flexibility beyond a direct simple-root construction [26]: the simple-root method gives one defining set and hence one dimension, while the repeated-root profile gives a whole interval of dimensions at the same certified distance. It also gives more flexibility than a uniform repeated-root lift because the dimension can be increased without weakening the layer-distance certificate.

Table 2 should be read at fixed simple-root input. For each value of q , all three non-naive profiles have the same ambient length and the same certified lower bound. The uniform upper lift is the most conservative repeated-root lift. The balanced profile keeps exact rate $1/2$ and gains two dimensions over this uniform upper lift when $e_0 = 4$. The

uniform lower profile gains two further dimensions, but its rate is above $1/2$. Hence the balanced profile gives the exact half-rate member of a family that cannot be obtained by choosing only one simple-root half-cyclotomic defining set. The comparison is internal to the proposed construction and should not be interpreted as a claim of optimality against all linear codes of the same length and dimension.

Table 2. Dimension steering for the family $q = 2^a$, $n = q^3 - 1$, and $e_0 = p^s = 4$. The three non-naive repeated-root codes use the same half-cyclotomic layer-distance bound $\Delta(q)$ but different multiplicity profiles.

q	n	$N = 4n$	$\Delta(q)$	Uniform H^+ Dim	Balanced Dim	Uniform H^- Dim	Gain Over H^+
2	7	28	2	12	14	16	+2
4	63	252	13	124	126	128	+2
8	511	2044	96	1020	1022	1024	+2
16	4095	16,380	725	8188	8190	8192	+2
32	32,767	131,068	5632	65,532	65,534	65,536	+2

Example 3. Take $q = 4$, $p = 2$, $s = 2$, $e_0 = 4$, and $n = 4^3 - 1 = 63$. The repeated-root length is $N = 252$. The coset-size distribution is $N_1 = 3$ and $N_3 = 20$. Thus $|H^+| = 2 + 3 \cdot 10 = 32$ and $|H^-| = 1 + 3 \cdot 10 = 31$. The certified simple-root layer bound is $\Delta(4) = \lfloor 4 \cdot 20 / (2(4 - 1)) \rfloor = 13$. The 2-adic repeated-root factors for $s = 2$ are $(P_0, P_1, P_2, P_3) = (1, 2, 2, 4)$.

Taking the balanced chain $T_0 = T_1 = H^+$ and $T_2 = T_3 = H^-$ gives $\dim_{\mathbb{F}_4} \mathcal{C}_{\text{bal}} = 252 - (32 + 32 + 31 + 31) = 126$. Since every active layer has distance at least 13, the distance reduction theorem gives

$$d_H(\mathcal{C}_{\text{bal}}) \geq \min\{1 \cdot 13, 2 \cdot 13, 2 \cdot 13, 4 \cdot 13\} = 13.$$

Hence the construction gives the certified repeated-root code $[252, 126, \geq 13]_4$. The two natural uniform endpoints are $\mathcal{C}_+ = [252, 124, \geq 13]_4$ and $\mathcal{C}_- = [252, 128, \geq 13]_4$. Therefore the balanced profile gains two dimensions over the uniform upper lift without lowering the certified distance, while remaining exactly half-rate.

By contrast, the naive single-layer lift $T_0 = H^+$ and $T_1 = T_2 = T_3 = \emptyset$ has dimension $252 - 32 = 220$, but the empty higher layers are full spaces with distance 1. Hence the repeated-root distance reduction of [27] gives

$$d_H \geq \min\{1 \cdot 13, 2 \cdot 1, 2 \cdot 1, 4 \cdot 1\} = 2.$$

Thus the large dimension of the naive lift is obtained at the cost of a very small certified repeated-root distance. The balanced and uniform profiles avoid this collapse by keeping nontrivial half-cyclotomic layers active at all multiplicity levels.

Table 3 makes the comparison transparent. The naive lift has the largest dimension but loses the repeated-root distance certificate. The uniform upper lift preserves the certified distance but is below half-rate. The balanced profile preserves the same certified distance and reaches exact half-rate. This is the structural flexibility obtained from the multiplicity-profile method.

Table 3. Comparison of four repeated-root profiles over $\mathbb{F}_4$ with $n = 63$ and $e_0 = 4$. The balanced code is the exact half-rate member.

Profile	Layers	Dimension	Certified d	Comment
Naive single layer	$H^+, \emptyset, \emptyset, \emptyset$	220	≥ 2	distance collapses
Uniform upper	H^+, H^+, H^+, H^+	124	≥ 13	conservative lift
Balanced mixed	H^+, H^+, H^-, H^-	126	≥ 13	exact half-rate
Uniform lower	H^-, H^-, H^-, H^-	128	≥ 13	higher-rate endpoint

For comparison with Grassl's database, Example 1 gives the exact repeated-root code $[28, 14, 4]_2$, while Grassl's database records a $[28, 14, 8]_2$ code and the matching upper bound 8. Hence that small example is not best-known as a general binary linear code. Its role is to verify the repeated-root profile mechanism in a case where the full generator polynomial and exact distance can be computed explicitly.

Example 4. Fix $q = 4$ and $n = 63$, and let $e_0 = 2^s$. The layer sizes remain $|H^+| = 32$ and $|H^-| = 31$, and the certified layer distance remains 13. For the balanced profile, $e_0/2$ layers are chosen equal to H^+ and $e_0/2$ layers equal to H^- . Therefore $\dim C_{\text{bal}} = 63e_0/2$ and $d_H(C_{\text{bal}}) \geq 13$. The uniform upper lift has dimension $63e_0/2 - e_0/2$, so the balanced profile gains $e_0/2$ dimensions at the same certified lower bound.

Table 4 illustrates a feature that is absent from the simple-root construction of [26]. In the simple-root case there is only one defining set. In the repeated-root case there are p^s multiplicity levels, and the profile can be tuned across these levels. Thus multiplicity is not merely a complication of repeated-root constacyclic codes, as in the classification problems studied for special lengths in [29]; it becomes a useful design variable.

Table 4. Fixed $q = 4$, $n = 63$: increasing the repeated-root index enlarges the dimension-steering interval while the certified lower bound remains 13.

$e_0 = p^s$	Length $N = 63e_0$	H^+ Dim	Balanced Dim	H^- Dim	Gain Over H^+	Certified d
2	126	62	63	64	+1	≥ 13
4	252	124	126	128	+2	≥ 13
8	504	248	252	256	+4	≥ 13
16	1008	496	504	512	+8	≥ 13

Example 5. Let $q = 3$, $p = 3$, $r = 1$, and $n = 3^3 - 1 = 26$. Then the cyclotomic cosets in $Z_{26,1}$ have sizes 1 and 3, with $N_1 = 2$ and $N_3 = 8$. Here $|H^+| = |H^-| = 1 + 3 \cdot 4 = 13$, so the upper and lower half-cyclotomic layers have the same size. The certified layer bound is $\Delta(3) = \lfloor 3 \cdot 8 / (2(3 - 1)) \rfloor = 6$. For $e_0 = 3^s$, the constant profile gives $\dim_{\mathbb{F}_3} C = 26e_0 - 13e_0 = 13e_0$, $d_H(C) \geq 6$. Thus the method still gives exact half-rate repeated-root codes with certified distance. However, because $|H^+| = |H^-|$, there is no dimension interval to exploit. This baseline clarifies where the advantage of the mixed-layer construction comes from: it is strongest when two admissible half-cyclotomic layers have the same certified distance but different sizes.

The first row of Table 5 gives the certified repeated-root parameter $[78, 39, \geq 6]_3$. Grassl's database gives lower and upper bounds $18 \leq d \leq 25$ for general ternary linear codes with parameters $[78, 39]$. Thus this entry is not competitive with the best-known general linear-code bounds. This comparison again emphasizes that the present examples are designed to demonstrate the repeated-root profile mechanism rather than to claim universal improvements over the best-known tables.

Table 5. Odd-characteristic baseline for $q = 3$, $n = 26$. Upper and lower half-cyclotomic layers have equal size, so the profile gives half-rate codes but no upper/lower dimension steering.

$e_0 = p^s$	Length $N = 26e_0$	Layer Size	Dimension	Certified d
3	78	13	39	≥ 6
9	234	13	117	≥ 6
27	702	13	351	≥ 6

8.1. Discussion

The examples show that the main role of the repeated-root multiplicity profile is not merely to increase the length, but to distribute simple-root defining sets across the p^s torsion layers. A uniform lift assigns the same multiplicity to all selected irreducible factors, while a mixed profile allows different protected layers to appear at different repeated-root levels. This gives a controlled way to adjust the dimension while preserving the strongest available layer-distance certificate.

The detailed example in Table 1 illustrates this effect clearly. The uniform upper lift has parameters $[68, 32, 8]_4$, whereas the balanced mixed profile gives $[68, 34, 8]_4$. Thus, the mixed profile gains two dimensions without decreasing the exact minimum distance. By contrast, the naive single-layer lift has much larger dimension, but its higher layers are full-space layers of distance 1, so the repeated-root distance collapses to 2.

This also explains the comparison with the simple-root construction of [26]. The simple-root construction provides good half-cyclotomic defining sets at length n . The present construction uses these sets as layer inputs and distributes them across the repeated-root length np^s . The distance formula of [27] then converts the layer distances into an exact repeated-root distance formula or into a certified lower bound. Hence, the layers are not only analyzed after the code is constructed; they are part of the design.

The method has natural limitations. If the available upper and lower half-cyclotomic layers have the same size, then mixing them does not create a dimension-steering interval, although it may still give half-rate repeated-root codes with certified distance. Also, heuristic distance estimates are not used as certificates; the tables use only exact distances or proved lower bounds. The construction is most effective when there are nested simple-root layers with comparable certified distances but different sizes, since this is precisely when multiplicity profiles can increase dimension without forcing a distance collapse.

8.2. MAGMA Verification Code

For reproducibility, we include a concise MAGMA script verifying Example 1 and Table 1. The remaining numerical tables are obtained by the same substitutions in the dimension and layer-bound formulas.

```

F<a> := GF(4);
K<b> := ext<F | 4>;
R<x> := PolynomialRing(F);

beta := b^((#K-1) div 51);
lambda0 := F!(beta^17);

function PFactor(t,p,s)
A := Intseq(t,p);
while #A lt s do Append(~A,0); end while;
return &*[A[i]+1 : i in [1..s]];
end function;

```

```

function CodeFromGenerator(g,N)
R := Parent(g);
F := BaseRing(R);
x := R.1;
G := Matrix(F,[[Coefficient(x^j*g,i) : i in [0..N-1]]
: j in [0..N-Degree(g)-1]]);
return LinearCode(G);
end function;

// Simple-root factors for q=4, n=17, r=3.
M1 := R!MinimalPolynomial(beta);
M7 := R!MinimalPolynomial(beta^7);
M34 := R!MinimalPolynomial(beta^34);

Factorization(x^17-lambda0);

// H^+ = C_1 union C_7 union C_34,
// H^- = C_1 union C_7.
CHplus := CodeFromGenerator(M1*M7*M34,17);
CHminus := CodeFromGenerator(M1*M7,17);

<Length(CHplus), Dimension(CHplus), MinimumDistance(CHplus)>;
<Length(CHminus), Dimension(CHminus), MinimumDistance(CHminus)>;

// Mixed repeated-root profile:
// e_1=e_7=4, e_34=2.
Crep := CodeFromGenerator(M1^4*M7^4*M34^2,68);
<Length(Crep), Dimension(Crep), MinimumDistance(Crep)>;

// Distance reduction check.
P := [PFactor(t,2,2) : t in [0..3]];
Min([P[i]*[8,8,7,7][i] : i in [1..4]]);

// Table comparison.
N := 68;
Hplus := 9;
Hminus := 8;

uniform_upper_dim := N-4*Hplus;
mixed_dim := N-2*Hplus-2*Hminus;
uniform_lower_dim := N-4*Hminus;
naive_dim := N-Hplus;

uniform_upper_d := 8;
mixed_d := 8;
uniform_lower_d := 7;
naive_d := Min([P[1]*8,P[2],P[3],P[4]]);

<"Uniform upper", uniform_upper_dim, uniform_upper_d>;

```

```

<"Balanced mixed", mixed_dim, mixed_d>;
<"Uniform lower", uniform_lower_dim, uniform_lower_d>;
<"Naive", naive_dim, naive_d>;

```

9. Conclusions

We introduced half-cyclotomic repeated-root constacyclic codes of length np^s through multiplicity profiles on the cyclotomic cosets of $x^n - \lambda_0$. These profiles produce nested simple-root layers, allowing the repeated-root distance formula [27] to be combined with half-cyclotomic simple-root bounds [26,30]. The main results are the layer reconstruction and dimension formulas, the exact distance reduction theorem, the layerwise lower bounds, the conditional half-rate family theorem, and explicit constructions for the three arithmetic length families considered in the paper.

The main conceptual point is that repeated-root multiplicities are used as active design variables rather than only as repeated algebraic factors. By choosing different defining sets at different multiplicity levels, the construction controls dimension through the total size of the layers and controls distance through the weakest P_t -weighted simple-root layer distance. This separates the theoretical mechanism from the computational verification: the framework is algebraic, while exact computations are used only to certify particular layer distances or examples.

The numerical examples illustrate this mechanism. In particular, the detailed mixed-profile example gives $[68, 34, 8]_4$, while the corresponding uniform upper lift gives $[68, 32, 8]_4$. Thus, the mixed profile gains two dimensions without decreasing the exact minimum distance. The examples also show why naive single-layer lifts may fail: once higher layers become full-space layers of distance 1, the repeated-root distance certificate can collapse.

The construction is not claimed to produce new optimal codes beyond Grassl's database unless an exact comparison for the same length and dimension is available. In the examples, we therefore state exact parameters or certified lower bounds, and any best-known or record-breaking claim requires separate verification.

Computationally, the method reduces repeated-root distance analysis to simple-root layer computations and the finite minimum in Theorem 4. Thus, instead of searching directly in the repeated-root ambient space, one constructs nested layers, verifies or certifies their simple-root distances, and evaluates the corresponding P_t -weighted minimum.

The framework may also extend to repeated-root quasi-cyclic or quasi-twisted codes by using several interacting layer chains. A skew constacyclic analogue is another possible direction, but it would require suitable noncommutative factorization tools and a skew version of the repeated-root distance reduction formula.

From a practical viewpoint, the construction is most useful when good simple-root constacyclic layers are already available and one wants to tune the repeated-root dimension without losing the verified layer-distance certificate. Future refinements can focus on optimizing admissible layered profiles and inserting exact layer distances into Theorem 4 to obtain sharper repeated-root parameters.

Author Contributions: Conceptualization, A.A.A. and S.H.S.; methodology, A.A.A. and S.H.S.; validation, A.A.A. and S.H.S.; formal analysis, A.A.A. and S.H.S.; writing—original draft preparation, A.A.A. and S.H.S.; writing—review and editing, A.A.A. and S.H.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the Ongoing Research Funding program, (ORF-2026-871), King Saud University, Riyadh, Saudi Arabia.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding authors.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- MacWilliams, F.J.; Sloane, N.J.A. *The Theory of Error-Correcting Codes*; North-Holland: Amsterdam, The Netherlands, 1977. Available online: <https://www.sciencedirect.com/bookseries/north-holland-mathematical-library/vol/16/suppl/C> (accessed on 16 June 2026).
- Huffman, W.C.; Pless, V. *Fundamentals of Error-Correcting Codes*; Cambridge University Press: Cambridge, UK, 2003. <https://doi.org/10.1017/CBO9780511807077>
- Ding, C.; Li, C. BCH cyclic codes. *Discret. Math.* **2024**, *347*, 113918. <https://doi.org/10.1016/j.disc.2024.113918>
- Aydin, N.; Murphree, J.M. New linear codes from constacyclic codes. *J. Franklin Inst.* **2014**, *351*, 1691–1699. <https://doi.org/10.1016/j.jfranklin.2013.11.019>
- Chen, B.; Fang, W.; Xia, S.T.; Fu, F.W. Constructions of optimal (r, δ) locally repairable codes via constacyclic codes. *IEEE Trans. Commun.* **2019**, *67*, 5253–5263. <https://doi.org/10.1109/TCOMM.2019.2916085>
- Kai, X.; Zhu, S.; Li, P. Constacyclic codes and some new quantum MDS codes. *IEEE Trans. Inform. Theory* **2014**, *60*, 2080–2086. <https://doi.org/10.1109/TIT.2014.2308180>
- Krishna, A.; Sarwate, D.V. Pseudocyclic maximum-distance-separable codes. *IEEE Trans. Inform. Theory* **1990**, *36*, 880–884. <https://doi.org/10.1109/18.53751>
- Dahl, C.; Pedersen, J.P. Cyclic and pseudo-cyclic MDS codes of length $q + 1$. *J. Combin. Theory Ser. A* **1992**, *59*, 130–133. [https://doi.org/10.1016/0097-3165\(92\)90104-3](https://doi.org/10.1016/0097-3165(92)90104-3)
- Wang, X.; Ding, C.; Liu, H.; Zheng, D. MDS constacyclic codes of length $q + 1$ over $\text{GF}(q)$. *Cryptogr. Commun.* **2024**, *16*, 21–48. <https://doi.org/10.1007/s12095-022-00624-0>
- Akre, D.; Aydin, N.; Harrington, M.; Pandey, S. A generalization of cyclic code equivalence algorithm to constacyclic codes. *Des. Codes Cryptogr.* **2023**, *91*, 763–777. <https://doi.org/10.1007/s10623-022-01124-1>
- Aydin, N.; Siap, I.; Ray-Chaudhuri, D.K. The structure of 1-generator quasi-twisted codes and new linear codes. *Des. Codes Cryptogr.* **2001**, *24*, 313–326. <https://doi.org/10.1023/A:1011283523000>
- Liu, H.; Li, C.; Ding, C. Five infinite families of binary cyclic codes and their related codes with good parameters. *Finite Fields Appl.* **2023**, *91*, 102270. <https://doi.org/10.1016/j.ffa.2023.102270>
- Sun, Z.; Ding, C.; Wang, X. Two classes of constacyclic codes with variable parameters $[(q^m - 1)/r, k, d]$. *IEEE Trans. Inform. Theory* **2024**, *70*, 93–114. <https://doi.org/10.1109/TIT.2023.3322990>
- Xie, C.; Chen, H.; Ding, C.; Sun, Z. Self-dual negacyclic codes with variable lengths and square-root-like lower bounds on the minimum distances. *IEEE Trans. Inform. Theory* **2024**, *70*, 4879–4888. <https://doi.org/10.1109/TIT.2024.3380613>
- Li, L.; Shi, M.; Tao, S.; Sun, Z.; Zhu, S.; Kim, J.-L.; Solé, P. A generalization of the Tang–Ding binary cyclic codes. *Discret. Math.* **2025**, *348*, 114390. <https://doi.org/10.1016/j.disc.2024.114390>
- Lu, H.; Qian, L.; Shi, M. An infinite family of binary cyclic codes and two infinite families of ternary cyclic codes with good parameters. *Cryptogr. Commun.* **2025**, *17*, 1407–1425. <https://doi.org/10.1007/s12095-025-00803-9>
- Lu, H.; Wang, X.; Shi, M. A tighter bound on the minimum distances for an infinite family of binary BCH codes and its generalization. *Finite Fields Appl.* **2025**, *108*, 102628. <https://doi.org/10.1016/j.ffa.2025.102628>
- Grassl, M. Bounds on the Minimum Distance of Linear Codes and Quantum Codes. Available online: <http://www.codetables.de> (accessed on 16 June 2026).
- Wang, D.; Cao, X.; Mi, J. Some arithmetical properties of cyclotomic cosets and their applications. *Discrete Math.* **2020**, *343*, 111971. <https://doi.org/10.1016/j.disc.2020.111971>
- Washington, L.C. *Introduction to Cyclotomic Fields*, 2nd ed.; Springer: New York, NY, USA, 1997. <https://doi.org/10.1007/978-1-4612-1934-7>
- Martinez-Perez, C.; Willems, W. Is the class of cyclic codes asymptotically good? *IEEE Trans. Inform. Theory* **2006**, *52*, 696–700. <https://doi.org/10.1109/TIT.2005.862123>
- Tang, C.; Ding, C. Binary $[n, (n + 1)/2]$ cyclic codes with good minimum distances. *IEEE Trans. Inform. Theory* **2022**, *68*, 7842–7849. <https://doi.org/10.1109/TIT.2022.3193715>
- Chen, T.; Ding, C.; Li, C.; Sun, Z. Four infinite families of ternary cyclic codes with a square-root-like lower bound. *Finite Fields Appl.* **2023**, *92*, 102308. <https://doi.org/10.1016/j.ffa.2023.102308>
- Chen, T.; Sun, Z.; Xie, C.; Chen, H.; Ding, C. Two classes of constacyclic codes with a square-root-like lower bound. *IEEE Trans. Inform. Theory* **2024**, *70*, 8734–8745. <https://doi.org/10.1109/TIT.2024.3464630>

25. Liu, H.; Gan, C.; Li, C.; Shi, X. Constructions of binary cyclic codes with minimum weights exceeding the square-root lower bound. *Des. Codes Cryptogr.* **2025**, *93*, 2971–2992. <https://doi.org/10.1007/s10623-025-01621-z>
26. Chen, Z.; Sha, M. Constacyclic codes with best-known parameters. *IEEE Trans. Inform. Theory* **2026**, *72*, 4837–4849. <https://doi.org/10.1109/TIT.2026.3697372>.
27. Castagnoli, G.; Massey, J.L.; Schoeller, P.A.; von Seemann, N. On repeated-root cyclic codes. *IEEE Trans. Inform. Theory* **1991**, *37*, 337–342. <https://doi.org/10.1109/18.75249>
28. Chen, H.; Xie, C.; Ding, C. Repeated-root cyclic codes with optimal parameters or best parameters known. *IEEE Trans. Inform. Theory* **2026**, *72*, 1618–1628. <https://doi.org/10.1109/TIT.2026.3653915>
29. Dinh, H.Q.; Ha, H.V.; Tran, N.T.H.; Vo, T.N. On constacyclic codes of length $10p^s$ over $\mathbb{F}_{p^m}$ and their optimal codes. *Finite Fields Appl.* **2026**, *116*, 102860. <https://doi.org/10.1016/j.ffa.2026.102860>
30. Sun, Z.; Li, C.; Ding, C. An infinite family of binary cyclic codes with best parameters. *IEEE Trans. Inform. Theory* **2024**, *70*, 2411–2418. <https://doi.org/10.1109/TIT.2023.3307732>

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.