

Optimal Locally Recoverable Codes With Hierarchy From Nested F -Adic Expansions

Austin Dukes, Giacomo Micheli^{ID}, and Vincenzo Pallozzi Lavorante^{ID}

Abstract—In this paper we construct new optimal hierarchical locally recoverable codes. Our construction is based on a combination of the ideas of Ballantine et al., (2019) and Sasidharan et al., (2015) with an algebraic number theoretical approach that allows to give a finer tuning of the minimum distance of the intermediate code (allowing larger dimension of the final code), and to remove restrictions on the arithmetic properties of q compared with the size of the locality sets in the hierarchy. In turn, we manage to obtain codes with a wider set of parameters both for the size q of the base field, and for the hierarchy size, while keeping the optimality of the codes we construct.

Index Terms—Hierarchical LRCs, Galois theory, Chebotarev density theorem.

I. INTRODUCTION

VARIOUS classes of locally recoverable codes have received great attention in recent times due to their applications to cloud and distributed storage systems [3], [4], [5], [6], [7], [8], [9], [10], [11]. In this paper we produce new optimal hierarchical locally recoverable codes (HLRCs). HLRCs are suitable solutions that address the problem of recovering lost information in a distributed storage system, and they have been widely studied in [1], [2], [12], [13], and [14].

Standard $[n, k, r, d]_q$ locally recoverable codes (LRCs) are linear codes over \mathbb{F}_q of dimension k , length n and distance d that allow to recover a single erasure in a codeword by looking at a maximum of r other components of the codeword, and to recover up to $d - 1$ simultaneous erasures by looking at a maximum of k other components. The set of other components that one needs to look at to recover a single erasure is said to be a locality set. It is clear from this definition that one desires r to be small and d to be large, while having relatively large rate k/n . In other words, one desires locality sets to be as small as possible but the global distance of the code to be large. The Singleton bound for locally recoverable codes relates these quantities and shows that at best d can be as

Manuscript received 17 August 2022; revised 8 March 2023; accepted 10 July 2023. Date of publication 24 July 2023; date of current version 20 October 2023. This work was supported by the National Science Foundation under Grant 2127742. (Corresponding author: Giacomo Micheli.)

Austin Dukes and Vincenzo Pallozzi Lavorante are with the Department of Mathematics, University of South Florida, Tampa, FL 33647 USA (e-mail: austindukes@usf.edu; vincenzop@usf.edu).

Giacomo Micheli is with the Department of Mathematics and the Center for Cryptographic Research, University of South Florida, Tampa, FL 33647 USA (e-mail: gmiceli@usf.edu).

Communicated by I. Tamo, Associate Editor for Coding and Decoding.

Digital Object Identifier 10.1109/TIT.2023.3298401

large as $n - k - \lceil k/r \rceil + 2$, showing that the smaller the locality is, the smaller the distance will be (for fixed k and n). From a general computer science perspective, LRCs deal with the most common scenario (one erasure) in an optimal way (looking only at a few other nodes) and with the unlikely scenario (multiple erasures) in an acceptable way (using the standard erasure recovery procedure). To motivate the research in this paper, consider the scenario in which there are λ simultaneous erasures, with λ satisfying $2 \leq \lambda \leq d - 2$. In particular, we set $\lambda = 2$ here for simplicity. If these 2 erasures happen in different locality sets, then one would need to look at $2r$ other nodes to recover the missing symbols, which is consistent with the fact that one recovers a single erasure by looking at r other nodes. If the erasures happen in the same locality set, then we are in the situation in which one needs (in principle) to use the same procedure as in the worst case scenario of $d - 1$ erasures, and this is inefficient because the case of 2 simultaneous erasures is by far more likely than the case of $d - 1$ erasures, when d is larger than λ .

Hierarchical locally recoverable codes allow the recovery of certain patterns of erasures by gradually looking at more components depending on the number of erasures that occurred. One can then design codes that recover one erasure by looking at a maximum of $b > 0$ other components; λ erasures by looking at $a > b$ other components; and $d - 1$ erasures by looking at a maximum of $k > a > b$ components, where k is the dimension of the code. This is impactful from a practical perspective, as one can deal with the most likely scenario (one erasure) in the optimal way, with the less likely scenario (λ erasures) in an acceptable way, and still be able to recover $d - 1$ erasures by accessing k nodes. Tuning these parameters in an efficient way depends on the reliability of the servers and the required efficiency of the system in terms of node retrieval. One of the features that one would desire from this kind of code is that λ is quite small, as the second-most likely scenario is the failure of only a few other nodes. We address this problem by writing a sharper Singleton bound for this regime of parameters and then constructing codes that achieve the bound.

A. Definitions

In the rest of the paper we will consider the occurrence of either one, λ , or $d - 1$ erasures, as these arise most commonly from applications (instead of the more general setting where one allows λ_1, λ_2 , or $d - 1$ erasures). Let n, k, b be positive integers with $k \leq n$. A locally recoverable code (LRC)

\mathcal{C} having parameters $[n, k, b]$ is an \mathbb{F}_q -subspace of \mathbb{F}_q^n of dimension k such that if one erases a single component of any $v \in \mathcal{C}$, that component can be recovered by accessing at most b other components of v . If d is the minimum distance of the code, we will write that \mathcal{C} is an $[n, k, d, b]$ LRC.

We now give the following definition which will be useful in the rest of the paper.

Definition 1.1: Let n be a positive integer, $\mathcal{C} \subseteq \mathbb{F}_q^n$ be a linear code, and S be a subset of the set of indices $\{1, \dots, n\}$. We say that \mathcal{C} can tolerate x erasures on S if, whenever there are x erasures on components of a codeword with indices belonging to S , the missing components can be recovered by looking at $|S| - x$ other coordinates with indices in S .

In this paper we construct new locally recoverable codes with hierarchy of locality sets. Our Definition 1.2 is equivalent to the one of hierarchical codes in [1] but we find it slightly easier to employ ours for practical situations, as we keep direct track of the size of the “hierarchy”.

Definition 1.2: Let n, k, d, b, a, λ be positive integers with $n > k$ and $2 \leq \lambda \leq b$. An $[n, k, d, b, a, \lambda]$ hierarchical locally recoverable code (HLRC) is an $[n, k, d]$ -linear code such that

- $(a + \lambda) \mid n$,
- $(b + 1) \mid (a + \lambda)$,
- the codeword indices are partitioned into $\ell \geq 1$ distinct sets A_i , each of size $a + \lambda$, such that \mathcal{C} tolerates λ erasures on A_i for every $i \in \{1, \dots, \ell\}$, and
- each A_i can be partitioned into $B_{i,j}$, each of size $b + 1$, such that \mathcal{C} tolerates 1 erasure on each $B_{i,j}$ for every $i \in \{1, \dots, \ell\}$ and every $j \in \{1, \dots, (a + \lambda)/(b + 1)\}$.

B. Motivation

Let us now briefly explain the motivation behind codes with hierarchical locality. Let T be the time needed to replace a failed node. Suppose that a second node fails in the same locality set as the first node during the time T . An $[n, k, d, b]$ LRC will still need to access k information symbols, as the 1-locality procedure is not guaranteed to work anymore. However, an $[n, k, d, b, a, \lambda]$ HLRC only requires accessing at most a information symbols. Since the failure of only a few nodes, say $\lambda < d - 1$, is significantly more likely than the failure of $d - 1$ nodes in the span of time T , it is convenient to have a code which addresses separately the case in which only λ nodes fail. The codes in [1] address this issue, but they are restricted to certain λ ’s, as we explain in subsection III-E. Moreover, in many cases they require restrictions on the arithmetic properties of q and on the size of the hierarchy (see for example the case of power functions in [1, Section IV.A, Example 1]).

C. Our Contribution

In this paper we provide new constructions of optimal codes with hierarchical locality and an improved bound for HLRCs for a special set of parameters. Our construction is based on the ideas in [1] combined with powerful techniques from algebraic number theory, allowing us to remove arithmetic restrictions on the size of the hierarchy compared with q or $q - 1$.

Structure of the paper:

- In Section I and its subsections we explain the basic coding theoretical definitions and provide the practical motivations for the study of such codes.
- In Section II, for some regime of parameters, we provide a stronger Singleton bound than the one already present in the literature for HLRCs [2]. Our bound beats the previous bound for an infinite set of parameters (see for example Remark 1).
- In Section III we achieve our new bound with a new construction of HLRCs that covers a set of parameters that are not available using previous constructions (see subsection III-E). In subsection III-F we construct one of our codes and provide its generator matrix.
- In Section IV we show that our codes are constructible without requiring arithmetic restrictions on $q, q - 1$, the locality parameters, and the sizes of the sets in the hierarchy.
- Using the existential results provided in Section IV, in Section V we provide some practical choices of parameters for codes with large length.

II. AN IMPROVED BOUND FOR HIERARCHICAL LOCALLY RECOVERABLE CODES

A. The Singleton Bound for $[n, k, d, b, a, \lambda]$ HLRCs With $\lambda \leq b$

Let $M_{m \times n}(q)$ denote the set of all matrices of dimension $m \times n$ defined over \mathbb{F}_q . The following is a well-known proposition, but we include a proof for completeness.

Proposition 2.1: Let \mathcal{C} be an $[n, k, d]_q$ linear code with generator matrix $\mathcal{G} \in M_{k \times n}(q)$ and let $S \in M_{k \times t}(q)$ be a submatrix of \mathcal{G} having rank $\text{rk}(S) \leq k - 1$. Then $t \leq n - d$.

Proof: Let $S = [S_1, \dots, S_t]$, where S_i is a column of \mathcal{G} for $i \in \{1, \dots, t\}$. Define $\tilde{S}: \mathbb{F}_q^k \rightarrow \mathbb{F}_q^t$ such that $x \mapsto \tilde{S}(x) = xS = [xS_1, \dots, xS_t]$. Since $\text{rk}(S) \leq k - 1$ and we can write $\tilde{S}(x) = x_1R_1 + \dots + x_kR_k$, where R_i are the rows of S , there exists a non-zero $x' \in \mathbb{F}_q^k$ such that $\tilde{S}(x') = 0$. Assuming without loss of generality that S consists of the first t columns of \mathcal{G} , since x' is non-zero, there exists a non-zero codeword $c = (\tilde{S}(x'), y_{t+1}, \dots, y_n)$ whose weight is at most $n - t$. Hence $d \leq n - t$. ■

To help the reader understand the more complex bound we propose on HLRCs, we include here a proof of the standard Singleton bound for LRCs.

Corollary 2.2: Let \mathcal{C} be a $[n, k, d, b]$ LRC. Then

$$k - 1 + \left\lfloor \frac{k - 1}{b} \right\rfloor \leq n - d.$$

Proof: We know that every set of $b + 1$ columns in each repair group has rank b by the locality condition. This means that we can choose a set S of $\left\lfloor \frac{k-1}{b} \right\rfloor (b+1) + \left\{ \frac{k-1}{b} \right\} b$ columns in such a way that $\text{rk}(S) \leq k - 1$ (here $\{x\}$ denotes the fractional part of x , i.e. $\{x\} = x - \lfloor x \rfloor$). Thus, by applying Proposition 2.1, we have the

following:

$$\left(\left\lfloor \frac{k-1}{b} \right\rfloor + \left\{ \frac{k-1}{b} \right\}\right) b + \left\lfloor \frac{k-1}{b} \right\rfloor = k-1 + \left\lfloor \frac{k-1}{b} \right\rfloor \leq n-d.$$

■

Notice that the above is equivalent to the well-known bound $d \leq n - k - \lceil k/b \rceil + 2$.

We aim to generalize the bound in Corollary 2.2 when \mathcal{C} is an $[n, k, d, b, a, \lambda]$ HLRC. The key observation is that one can partition the columns of the generator matrix into ℓ sets of $a + \lambda$ columns so that each set has rank less than a , and each set of $a + \lambda$ columns can be partitioned further into sets of $b + 1$ columns so that each set has rank at most b . To see this, notice that for any $i \in \{1, \dots, \ell\}$, each set S_i of $a + \lambda$ columns (corresponding to A_i) can be divided into $\beta = (a + \lambda)/(b + 1)$ sets, say $S_{i,j}$ for $j \in \{1, \dots, \beta\}$, of $b + 1$ columns (corresponding to $B_{i,j}$) with rank at most b by the definition of the code. Now, in the first set $S_{i,1}$ we have λ columns which are in the span of the other a columns in S_i . This means that we can choose $b + 1 - \lambda$ columns from $S_{i,1}$ and b columns from each of the other $S_{i,j}$, with $j \neq 1$, and be able to reconstruct any λ of the $a + \lambda$ columns in S_i . Therefore, the rank of each S_i is at most

$$\rho := \underbrace{\left[\frac{(a + \lambda)/(b + 1) - 1}{\beta - 1} \right]}_{\beta - 1} b + (b + 1 - \lambda) \leq a. \quad (1)$$

Theorem 2.3: Let \mathcal{C} be a $[n, k, d, b, a, \lambda]$ HLRC with $\lambda \leq b$, and let $\rho = b(a + \lambda)/(b + 1) - (\lambda - 1)$. Then

$$\left\lfloor \frac{k-1}{\rho} \right\rfloor (a + \lambda) + k_1 + \left\lfloor \frac{k_1}{b} \right\rfloor \leq n - d, \quad (2)$$

where k_1 is defined by $k - 1 \equiv k_1 \pmod{\rho}$ and $0 \leq k_1 < \rho$.

Proof: Given $\left\lfloor \frac{k-1}{\rho} \right\rfloor$ sets of $a + \lambda$ columns S_i corresponding to the larger locality sets for $i \in \{1, \dots, \left\lfloor \frac{k-1}{\rho} \right\rfloor\}$, denote by S the union of the S_i 's. Thus, $|S| = \left\lfloor \frac{k-1}{\rho} \right\rfloor (a + \lambda)$ and by the above discussion we have $\text{rk}(S) = \rho \left\lfloor \frac{k-1}{\rho} \right\rfloor \leq k - 1$. This allows us to add more columns to S until the rank equals $k - 1$ using a smaller locality set. More precisely, we can always choose a set of the remaining columns of \mathcal{G} , say S' , of size $\lfloor \frac{k_1}{b} \rfloor (b + 1) + \left\{ \frac{k_1}{b} \right\} b$, such that $\text{rk}(S') \leq k_1$ (explicitly, S' is the union of the columns in $B_{\lfloor \frac{k-1}{\rho} \rfloor + 1, j}$, for $j \in \{1, \dots, \lfloor \frac{k_1}{b} \rfloor\}$ and $\left\{ \frac{k_1}{b} \right\} b$ columns in $B_{\lfloor \frac{k-1}{\rho} \rfloor + 1, \lfloor \frac{k_1}{b} \rfloor + 1}$). Hence,

$$\text{rk}(S \cup S') \leq \left\lfloor \frac{k-1}{\rho} \right\rfloor \rho + k_1 = k - 1,$$

simply by the definition of k_1 . Applying Proposition 2.1 on the set of columns $S \cup S'$ we have that

$$\left\lfloor \frac{k-1}{\rho} \right\rfloor (a + \lambda) + \left\lfloor \frac{k_1}{b} \right\rfloor (b + 1) + \left\{ \frac{k_1}{b} \right\} b \leq n - d.$$

Now, since $\frac{k_1}{b} = \lfloor \frac{k_1}{b} \rfloor + \left\{ \frac{k_1}{b} \right\}$ we have

$$\left\lfloor \frac{k-1}{\rho} \right\rfloor (a + \lambda) + k_1 + \left\lfloor \frac{k_1}{b} \right\rfloor \leq n - d.$$

Definition 2.1: We say that an $[n, k, d, b, a, \lambda]$ HLRC is *optimal* if its minimum distance attains the upper bound in (2), i.e., if

$$d = n - \left(\left\lfloor \frac{k-1}{\rho} \right\rfloor (a + \lambda) + k_1 + \left\lfloor \frac{k_1}{b} \right\rfloor \right),$$

for $k - 1 \equiv k_1 \pmod{\rho}$ and $0 \leq k_1 < \rho$.

Remark 1: Note that our bound improves upon the bound in [2] for infinitely many parameters, but ours holds only for $\lambda \leq b$. In fact, for any length n , and for parameters $k = 6$, $a = 4$, $r_1 = \rho = 3$, $r_2 = b = 2$, $\delta_1 = \lambda + 1 = 3$ and $\delta_2 = 2$ [2, Theorem 2.1] gives $d \leq n - 8$ when instead our bound gives $d \leq n - 9$. The moral reasons for this are that we are taking into account a finer arithmetic of the parameters which involves the reduction of the dimension modulo the upper level hierarchical locality, and we are restricting to the case in which the number of nodes that we simultaneously erase is strictly smaller than the size of the smaller locality set.

III. OUR CONSTRUCTION OF OPTIMAL HLRCs USING NESTED F -ADIC EXPANSIONS

A. Main Tool for the Construction

Lemma 3.1: Let $f, h \in \mathbb{F}_q[X]$ be non-constant polynomials. Suppose there is some $t_0 \in \mathbb{F}_q$ such that $f(h(X)) - t_0$ splits completely (i.e., factors into $(\deg f)(\deg h)$ distinct factors) over \mathbb{F}_q . Then the set of roots of $f(h(X)) - t_0$, say A_0 , can be partitioned into sets $B_1, \dots, B_{\deg f} \subseteq \mathbb{F}_q$ which satisfy the following:

- $h(B_i) = c_i \in \mathbb{F}_q$ for each $1 \leq i \leq \deg f$,
- the cardinality of each B_i is $\deg h$, and
- $h(B_i) \neq h(B_j)$ whenever $i \neq j$.

Proof: By the hypothesis we may write $f(h(X)) - t_0 = \prod_{i=1}^{\deg f} (X - x_i)$ for distinct elements $x_1, \dots, x_{\deg f} \in \mathbb{F}_q$. Notice now that if $f(h(X)) - t_0$ splits completely, then $f(X) - t_0$ splits completely. If we let $\alpha_1, \dots, \alpha_{\deg f} \in \mathbb{F}_q$ be the (distinct) roots of $f(X) - t_0$, then we may also write $f(h(X)) - t_0 = \prod_{i=1}^{\deg f} (h(X) - \alpha_i) \in \mathbb{F}_q[X]$. Combining these two factorizations and relabeling the x_i appropriately yields

$$\prod_{i=1}^{\deg f} \prod_{j=1}^{\deg h} (X - x_{i,j}) = \prod_{i=1}^{\deg f} (h(X) - \alpha_i),$$

where $\prod_{j=1}^{\deg h} (X - x_{i,j}) = h(X) - \alpha_i$ for each $1 \leq i \leq \deg f$.

In particular, it follows that $\alpha_i \in \mathbb{F}_q$ for each i . Write $B_i = \{x_{i,j} : 1 \leq j \leq \deg h\}$. Then we have $h(B_i) = \alpha_i$ for each i , proving the first statement. The second and third statements both follow from the fact that the $x_{i,j}$'s are pairwise distinct and the B_i 's are pairwise disjoint. ■

Definition 3.1: For $f, h \in \mathbb{F}_q[X]$, we say that a set $A \subset \mathbb{F}_q$ is a *nest* for (f, h) if A is the set of preimages of $t_0 \in \mathbb{F}_q$ such that $f(h(X)) - t_0$ splits completely into distinct linear factors.

Furthermore, we say that $B \subset A$ is a *sub-nest* if h is constant on B and $|B| = \deg h$.

B. The Main Construction

We present a general method of constructing linear codes with the nested locality property. Later we will show that these codes are optimal in the sense of Section III. In line with the notion of (r, ℓ) -good polynomials in [15], we now begin defining our *nested* polynomials.

Definition 3.2 (ℓ -Nested): Let $f, h \in \mathbb{F}_q[X]$ and let ℓ be a positive integer. Then f and h are said to be ℓ -nested if $f(h(X)) - t_0$ splits completely over \mathbb{F}_q for at least ℓ elements $t_0 \in \mathbb{F}_q$.

Remark 2: Note that if f and h are ℓ -nested, then from Lemma 3.1 there exist A_1, \dots, A_ℓ distinct nests for (f, h) such that

- for any $i \in \{1, \dots, \ell\}$, $f(h(A_i)) = \{t_i\}$ for some $t_i \in \mathbb{F}_q$,
- $|A_i| = \deg f \deg h$,
- $A_i \cap A_j = \emptyset$ for any $i \neq j$, and
- each A_i can be partitioned into sub-nests $B_{i,j}$ for (f, h) .

Those properties will be the key of our next construction.

Construction 3.2 (Nested HLRCs): Let $f, h \in \mathbb{F}_q[X]$ be ℓ -nested, with $3 \leq \deg h = b+1$ and $\deg f(X) = (a+\lambda)/(b+1)$ for some integer $2 \leq \lambda \leq b$, and let $\mathcal{A} = \cup_{i=1}^\ell A_i$, where $\{A_1, \dots, A_\ell\}$ is a set of nests for (f, h) .

For a positive integer $s \geq 1$, consider the set \mathcal{M} of polynomials of the form $m(X)$ equals to

$$\sum_{i=0}^s \left[\left(\sum_{j=0}^{\deg f-2} g_{i,j}(X)h(X)^j \right) + \tilde{g}_i(X)h(X)^{\deg f-1} \right] f(h(X))^i, \quad (3)$$

where $g_{i,j} \in \mathbb{F}_q[X]_{\leq \deg h-2}$ and $\tilde{g}_i \in \mathbb{F}_q[X]_{\leq \deg h-\lambda-1}$.

Let $n = (\deg f \deg h)\ell$ and let k be the dimension of \mathcal{M} as an \mathbb{F}_q -vector space. Once an ordering on \mathcal{A} is fixed, we can define

$$\mathcal{C} := \{(m(x), x \in \mathcal{A}) \mid m \in \mathcal{M}\}. \quad (4)$$

We will prove that \mathcal{C} is an optimal $[n, k, b, a, \lambda]$ HLRC over \mathbb{F}_q .

C. Locality

Since we evaluate at n distinct points of \mathbb{F}_q , we need $q \geq n$. Write $n = (a+\lambda)(s+1)$ and recall that $b+1$ divides $a+\lambda$.

Take $\mathbf{a} \in \mathbb{F}_q^k$ and write

$$\text{Enc}_{\mathcal{C}}(\mathbf{a}) = \mathbf{c} = (c_{i,j_1,j_2})_{1 \leq i \leq s+1, 1 \leq j_1 \leq (a+\lambda)/(b+1), 1 \leq j_2 \leq b+1}.$$

Note that the index i determines a nest A_i , the index j_1 determines a sub-nest B_{i,j_1} and the index j_2 determines an element of the sub-nest being considered, which is then denoted by c_{i,j_1,j_2} . We begin by showing that the code \mathcal{C} described in Construction 3.2 allows one to recover a single missing component of \mathbf{c} by accessing at most b other components of \mathbf{c} .

Fix $\ell \geq 1$ and let $f, h \in \mathbb{F}_q[X]$ be the ℓ -nested polynomials from which \mathcal{C} is obtained. Write $\mathcal{A} = \{A_1, \dots, A_\ell\}$ with $A_i =$

$\bigcup_{j=1}^{\deg f} B_{i,j_1}$ and $B_{i,j_1} = \{x_{i,j_1,j_2} : 1 \leq j_2 \leq b+1\}$ as in Remark 2.

Without loss of generality, assume that the missing component is $c_{1,1,b+1} = m_{\mathbf{a}}(x_{1,1,b+1})$, where $m_{\mathbf{a}} \in \mathcal{M}$. Observe immediately that because both of $f \circ h$ and h are constant on $B_{1,1}$, the restriction $m_{\mathbf{a}}|_{B_{1,1}}$ can be written as a polynomial of degree $\max\{\deg h - 2, \deg h - \lambda - 1\} = \deg h - 2 = b - 1$. Since $x_{1,1,j_2} \in B_{1,1}$ for each j_2 , we have that $m_{\mathbf{a}}|_{B_{1,1}}(x_{1,1,j_2}) = m_{\mathbf{a}}(x_{1,1,j_2}) = c_{1,1,j_2}$. Using Lagrange interpolation on the points $(x_{1,1,j_2}, c_{1,1,j_2})$ for $1 \leq j_2 \leq b$, we obtain a polynomial $\Delta^{B_{1,1}}$ of degree $b-1$ which agrees with $m_{\mathbf{a}}|_{B_{1,1}}$ at b distinct points, so the two polynomials must be equal. Thus we can recover $c_{1,1,b+1}$ by evaluating $\Delta^{B_{1,1}}$ at the element $x_{1,1,b+1}$.

Let us now consider the case of λ erasures. Among these λ erasures, the erasures which are isolated in locality sets $B_{i,j}$ can be recovered by using the 1-locality, so the interesting case is when multiple erasures occur in the same $B_{i,j}$. Let us assume that $\lambda \geq 2$ erasures occur in the same locality set $B_{i,j}$. In this case, since $f \circ h$ is constant on A_i , the restriction $m_{\mathbf{a}}|_{A_i}$ is a polynomial of degree $\deg f \deg h - \lambda - 1 = a + \lambda - \lambda - 1 = a - 1$. Thus Lagrange interpolation on a set of at most a points of A_i on which no erasure occurred yields a polynomial Δ^{A_i} which agrees with $m_{\mathbf{a}}$ on all of A_i . Hence the missing components can be obtained by evaluating Δ^{A_i} at each of the corresponding locations in A_i .

D. Optimality of the Code

We dedicate this subsection to proving the optimality of our code \mathcal{C} . Therefore, we will be computing the values of k and d .

Lemma 3.3: Let \mathcal{C} be the code in (4). Then

$$k = (s+1)((\deg f - 1)(\deg h - 1) + \deg h - \lambda).$$

Proof: Since in particular $\deg(g_{i,j}h^j + \deg \tilde{g}_i h^{\deg f-1}) \leq \deg(f \circ h)$ and $\deg g_{i,j}, \deg \tilde{g}_i \leq \deg h$, by uniqueness of F -adic expansion both for $F = f \circ h$ and for $F = h$, we have

$$k = \dim_{\mathbb{F}_q} \mathcal{M} = (s+1)((\deg f - 1)(\deg h - 1) + (\deg h - \lambda)),$$

as we wanted to prove. \blacksquare

Lemma 3.4: Let \mathcal{C} be the code in (4). Then $d \geq n - \delta$, for

$$\delta = (s+1) \deg h \deg f - \lambda - 1.$$

Proof: A lower bound for the minimum distance is obtained by subtracting δ from n , where δ is the upper bound for the maximum number of zeros of $m \in \mathcal{M}$. We compute

$$\begin{aligned} \delta &= ((\deg f \deg h)s + (\deg f - 1) \deg h + \deg h - \lambda - 1) \\ &= (s+1) \deg h \deg f - \lambda - 1, \end{aligned} \quad (5)$$

and this proves the claim. \blacksquare

Theorem 3.5: Let \mathcal{C} be the code obtained by using Construction 3.2. Then \mathcal{C} is an optimal $[n, k, b, a, \lambda]$ HLRC.

Proof: Let $\rho = b(a+\lambda)/(b+1) - (\lambda-1)$ and $k_1 = k-1 - \left\lfloor \frac{k-1}{\rho} \right\rfloor \rho$. Moreover, we recall that $a+\lambda = \deg f \deg h$

and $\deg h = b + 1$. Let d' denote the optimal distance, such that

$$\delta' = n - d' = \left(\left\lfloor \frac{k-1}{\rho} \right\rfloor (a + \lambda) + k_1 + \left\lfloor \frac{k_1}{b} \right\rfloor \right).$$

Note that

$$\left\lfloor \frac{k-1}{\rho} \right\rfloor = \left\lfloor \frac{1}{-\deg f \deg h + \deg f + \lambda - 1} + s + 1 \right\rfloor = s,$$

since $\lambda \leq \deg h - 1$, and

$$\left\lfloor \frac{k_1}{b} \right\rfloor = \left\lfloor \deg f - \frac{\lambda}{\deg h - 1} \right\rfloor = \deg f - \left\lceil \frac{\lambda}{\deg h - 1} \right\rceil,$$

in fact $k_1 = \deg f(-(b+1)s + \deg h(s+1) - 1) - \lambda = \deg f(\deg h - 1) - \lambda$. By using the results of Lemma 3.3, 3.4, we have

$$\begin{aligned} \delta - \delta' &= (s+1) \deg h \deg f - \lambda - 1 - s \deg f \deg h \\ &\quad + -k_1 - \left\lfloor \frac{k_1}{b} \right\rfloor \\ &= \deg f \deg h - \lambda - 1 - (\deg f(\deg h - 1) - \lambda) \\ &\quad + -\deg f + \left\lceil \frac{\lambda}{\deg h - 1} \right\rceil \\ &= \left\lceil \frac{\lambda}{\deg h - 1} \right\rceil - 1, \end{aligned} \quad (6)$$

and since $\left\lceil \frac{\lambda}{\deg h - 1} \right\rceil - 1 = 0$ for $\lambda \leq \deg h - 1$, the code is optimal. \blacksquare

E. Comparison With the Optimal Hierarchical HLRCs in [1]

A construction of optimal HLRCs for a certain set of parameters is presented in [1, Proposition IV.2]. Let us fix the parameters for which that construction exists, i.e., $r_1 = sr_2$ (we note that we do not require such a constraint, but that even in this scenario we show that we can construct codes that are not available from [1, Proposition IV.2]). The set of parameters of the codes in [1, Proposition IV.2], given also in our notation, is as follows:

- the length of the codes in both settings is n ,
- each small locality set (at the bottom level of the hierarchy) has size $r_2 + 1$, so in our case each has size $b + 1$,
- their ν is our $a + \lambda$,
- the middle code has distance $r_2 + 3$ and hence can tolerate $r_2 + 2$ erasures, so their $r_2 + 2$ corresponds to our λ ,
- their r_1 is our ρ ,
- the code is optimal, with distance $d = n - t(r_1 + r_2 + 1 + s) + r_2 + 3$, for some t, s ,
- the two-level hierarchy has locality parameters $(r_1, r_2 + 3)$ and $(r_2, 2)$.

Therefore, using the notation of [1] for an $[n, k, d, \{r_1, \delta_1\}, \{r_2, 2\}]$ HLRC, we constructed an $[n, k, d, \{\rho, \lambda + 1\}, \{b, 2\}]$ HLRC, where ρ is defined as in (1). This shows immediately that our class of codes is different from the codes in [1, Proposition IV.2]. In fact, the optimality of our codes strongly relies on the assumption $\lambda \leq r_2$, which is not the case in the construction in

[1, Proposition IV.2], in which instead they deal with a complementary case $\lambda = r_2 + 2$. It follows that our class of codes contains codes which are not covered by this construction, as we can construct optimal hierarchical codes with two-level hierarchy having locality parameters $(r_1, \lambda + 1)$ and $(r_2, 2)$, for any $\lambda \leq r_2$, such as for $\lambda = 2$.

We emphasize that in [2] it is necessary to set a fixed $\lambda = r_2 + 2$ since in this way one can reach optimality using the bound in [2, Theorem 2.1], while, using our improved bound and enhancing the construction in [1], one obtains more flexibility as we explained. Moreover, we will see in detail how to construct our codes without the arithmetic restrictions appearing in the examples which use monomials or linearized functions (see Section IV).

For a better comparison and to simplify the understanding, in the next paragraph, we will still use monomials for a tutorial example, even if it is not a requirement as we explain in Section IV.

F. Tutorial Example

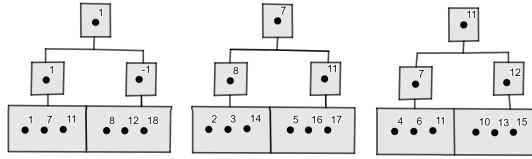
Suppose one desires a code over \mathbb{F}_{19} of dimension 6 which can recover 1, 2, and 8 lost nodes by accessing at most 2, 4, and 6 other nodes, respectively (i.e., the distance of the code equal 9). This is not possible using the standard Tamo-Barg construction since, to recover more than 1 node, one would need to access as many nodes as the dimension of the code, that is, 6 nodes. Let now C_n denote the cyclic group with n elements (in this framework we use the multiplicative notation, i.e., $C_n := \{1, g, \dots, g^{n-1}\}$ with $g^n = 1$). Another option is to consider codes with availability using an orthogonal partition of the multiplicative group of \mathbb{F}_{19} that includes C_3 (as one wants the locality to be 3). But this does not work in this case either as the only other option is C_9 and $C_3 \subseteq C_9$ (since \mathbb{F}_q^* is cyclic for any prime power q). Moreover [2, Proposition IV.2] does not hold for $\lambda = 2$.

Our construction instead provides a code that allows these recovery capabilities and is information theoretically optimal in the sense of the Singleton bound in Subsection II-A.

Suppose we choose $f = X^2$ and $h = X^3$ (so $b = 2$ and $a = 4$). A general information polynomial is given by

$$m(X) = \sum_{i=0}^1 \left[g_i(X) + \tilde{g}_i(X)X^3 \right] x^{6i},$$

for $g_i \in \mathbb{F}_q[X]_{\leq 1}$ and $\tilde{g}_i \in \mathbb{F}_q[X]_{\leq 0}$. In particular the $\tilde{g}_i(X) = \tilde{g}_i$ are constants (notice that the internal sum in j in (3) disappears as $\deg(f) = 2$). Therefore, by evaluating the messages at the preimage of the 3 totally split places of $x^6 = f \circ h$, we get a code of length $n = 18$, dimension $k = 6$, $b = 2$ and $a = 4$. Notice that this code can recover 1 erasure by looking at $b = 2$ other nodes. Moreover, if two erasures occur, we have two possibilities: either the erasures occurred in different nests, or the erasures occurred in the same nest for (f, h) . In the first case, one can use twice the locality (that is 2) to recover each node. In the second case one needs to access just 3 other nodes by carefully considering all the linear dependencies of the nodes in the nest. Since we are evaluating polynomials of degree at most 9, the distance of

Fig. 1. The nests for the pair (x^2, x^3) and their evaluations.

the code is $18 - 9 = 9$ and therefore one has a fault tolerance of 8 erasures. Practically, given those 18 nodes, we are looking at the following matrix:

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 7 & 11 & 8 & 12 & 18 & 2 & 3 & 14 & 5 & 16 & 17 & 4 & 6 & 9 & 10 & 13 & 15 \\ 1 & 1 & 1 & 18 & 18 & 18 & 8 & 8 & 8 & 11 & 11 & 11 & 7 & 7 & 7 & 12 & 12 & 12 \\ 1 & 1 & 1 & 1 & 1 & 1 & 7 & 7 & 7 & 7 & 7 & 11 & 11 & 11 & 11 & 11 & 11 & 11 \\ 1 & 7 & 11 & 8 & 12 & 18 & 14 & 2 & 3 & 16 & 17 & 5 & 6 & 9 & 4 & 15 & 10 & 13 \\ 1 & 1 & 1 & 18 & 18 & 18 & 18 & 18 & 18 & 1 & 1 & 1 & 1 & 1 & 18 & 18 & 18 & 18 \end{pmatrix}$$

where the rows correspond to (the evaluation of) the basis $\{1, x, x^3, x^6, x^7, x^9\}$ and the columns to the elements of \mathbb{F}_{19}^* ordered as in Figure 1. This means that to check the locality of each set one just needs to check the rank of the corresponding set of columns in the above matrix. For example, suppose we want to recover the third column, which corresponds to the symbol 11. We can do that using only the first two columns, since the matrix

$$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 7 & 11 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 7 & 11 \\ 1 & 1 & 1 \end{pmatrix}$$

has rank equal to 2. Similarly, we can recover any two lost symbols using either 3 (if they belong to the same nest) or 4 (otherwise, if they belong to distinct nests) other symbols.

IV. EXISTENTIAL RESULTS VIA CHEBOTAREV DENSITY THEOREM

In this section we explain how to apply Chebotarev Density Theorem to count the places $t_0 \in \mathbb{F}_q$ such that $f(h) - t_0$ is totally split, i.e. splits completely into $\deg(f(h))$ linear factors. A lower bound on this quantity determines directly a lower bound on the size of the hierarchy in our construction. This determines completely the range of parameters of our hierarchical codes, and in turn it shows that they always exist for q large enough, without arithmetic restrictions on the localities and the size of the base field.

A. Background on Galois Theory

We begin by recalling a few preliminary definitions. Let K and M be fields. We will write $K[X]$ to denote the polynomial ring in the indeterminate X over K . The field extension $K \subseteq M$ will be written as M/K , and its degree, that is, the dimension of M as a K -vector space, as $[M : K]$. For q a power of a prime, let \mathbb{F}_q be the finite field with q elements and let $\mathbb{F}_q^* = \mathbb{F}_q \setminus \{0\}$ be its (cyclic) multiplicative subgroup. Let t be transcendental over \mathbb{F}_q and denote by $\mathbb{F}_q(t)$ the rational function field in t over \mathbb{F}_q .

We follow closely the notation and terminology in [16] throughout this section, and we provide the essential notions here. A finite-dimensional extension K of $\mathbb{F}_q(t)$ is said to be a (global) function field over \mathbb{F}_q . A valuation ring of a function field M is a ring \mathcal{O} such that $K \subsetneq \mathcal{O} \subsetneq M$ and which contains at least one of z or z^{-1} for every $z \in M$. A place P of M is the unique maximal ideal of some valuation ring \mathcal{O} of M , and the degree of P is defined to be $\deg P = [\mathcal{O}/P : \mathbb{F}_q]$. In particular, P is said to be a place of degree one (or equivalently, a rational place) of M if $[\mathcal{O}/P : \mathbb{F}_q] = 1$. There is a one-to-one correspondence between places of M and valuation rings \mathcal{O} of M , so we will write \mathcal{O}_P to denote the valuation ring whose maximal ideal is P . We will write \mathbb{P}_M to denote the set of all places of M and $\mathbb{P}_M^1 \subseteq \mathbb{P}_M$ to denote the set of all degree one (rational) places of M . Let $K \subseteq M$ be an extension of function fields. For places $P \in \mathbb{P}_K$ and $Q \in \mathbb{P}_M$, we say that Q lies above P (and write $Q | P$) if $P \subseteq Q$. We denote the ramification index and relative degree of the extension of places $Q | P$ by $e(Q | P)$ and by $f(Q | P)$, respectively. Let \mathbb{P}_K^1 be the set of places of degree 1 of K . Also, we define

$$\text{Ram}^1(M/K) = \{P \in \mathbb{P}_K^1 : e(Q | P) \geq 2, \text{ for some } Q \text{ place of } M \text{ lying above } P\}.$$

The automorphism group of M/K , that is, the group of all automorphisms of M which fix K element-wise, is denoted by $\text{Aut}(M/K)$. When $|\text{Aut}(M/K)| = [M : K]$, we say that the extension M/K is Galois with Galois group $\text{Gal}(M/K) = \text{Aut}(M/K)$. For a polynomial $g \in K[x]$ with splitting field M we write $\text{Gal}(g | K) = \text{Gal}(M/K)$. We say that a polynomial $f \in \mathbb{F}_q[X]$ is separable over \mathbb{F}_q if $f \notin \mathbb{F}_q[X^p]$, where $p = \text{char } \mathbb{F}_q$, and for such an f , the polynomial $f - t$ is seen to be a separable irreducible polynomial over $\mathbb{F}_q(t)$. We will write M_f to denote the splitting field of $f - t$ over $\mathbb{F}_q(t)$. Equivalently, M_f denotes the Galois closure of the extension $\mathbb{F}_q(x)/\mathbb{F}_q(t)$, where x is any root of $f(X) - t$ in the algebraic closure $\overline{\mathbb{F}_q(t)}$ of $\mathbb{F}_q(t)$. The field of constants of M_f will be denoted by k_f , and we note that it is possible to have $k_f \supsetneq \mathbb{F}_q$. Let G_f be the monodromy group of f (sometimes called the arithmetic Galois group of $f - t$), that is, the Galois group of the extension $M_f/\mathbb{F}_q(t)$. Let $N_f = \text{Gal}(M_f/k_f(t)) \cong \text{Gal}(\overline{\mathbb{F}_q M_f}/\overline{\mathbb{F}_q(t)})$.

B. The Number of Totally Split Places t_0 of $f(h) - t$

We will use the Chebotarev density theorem as in Proposition 3.1 of [15] since this formulation is the most convenient for our purposes. We provide a full exposition in this section, but we briefly describe in the next paragraph the general procedure and ideas.

For polynomials $f, h \in \mathbb{F}_q[X]$, consider the composition $f(h)$. By the lower bound in [15, Proposition 3.1] on the number ℓ of $t_0 \in \mathbb{F}_q$ such that $f(h) - t_0$ splits into linear factors over \mathbb{F}_q , we have that for large enough q it is guaranteed to have a large number of totally split places of degree 1 of $\mathbb{F}_q(x)/\mathbb{F}_q(t)$ when $f(h)$ is chosen correctly. Now, we may assume that the field of constants $k_{f(h)}$ of $M_{f(h)}$ is trivial since otherwise there cannot be a totally split place of degree 1. Since we want ℓ to be as large as possible,

one quickly sees from the lower bound in [15, Proposition 3.1] that minimizing the size of the monodromy group $G_{f(h)}$ of $f(h)$ guarantees a large lower bound for ℓ . Thus our construction always effectively results in an optimal code as long as the size of the alphabet verifies a certain lower bound.

For the extension $M_{f(h)}/\mathbb{F}_q(t)$, let $G_{f(h)} = \text{Gal}(M_{f(h)}/\mathbb{F}_q(t))$ be its arithmetic Galois group and let $N_{f(h)}$ be its geometric Galois group. Since we are interested in the number ℓ of places $P \subseteq \mathbb{F}_q(t)$ of degree 1 which are totally split in $M_{f(h)}$, by Proposition 3.4 of [15] we may assume that $k_{f(h)} = M_{f(h)} \cap \mathbb{F}_q = \mathbb{F}_q$ is the field of constants of the extension $M_{f(h)}/\mathbb{F}_q(t)$ since otherwise $\ell = 0$. Hence $G_{f(h)} = N_{f(h)}$.

Lemma 4.1: Let $f, h \in \mathbb{F}_q[X]$ be nonzero polynomials having positive degrees and assume that $k_{f(h)} = \mathbb{F}_q$. Define $G_f = \text{Gal}(f(X) - t \mid \mathbb{F}_q(t))$ and similarly $G_h = \text{Gal}(h(X) - t \mid \mathbb{F}_q(t))$. Then the number of $t_0 \in \mathbb{F}_q$ such that $f(h(X)) - t_0$ splits completely into distinct (linear) factors over \mathbb{F}_q is at least $\frac{1}{|G_h|^{\deg f}|G_f|}q + O(\sqrt{q})$, where the implied constant can be made explicit and is independent of q .

Proof: Denoting the number of $t_0 \in \mathbb{F}_q$ we are considering by $|T_{\text{split}}^1(f(h))|$, from Proposition 3.1(ii) of [15] we immediately have

$$|T_{\text{split}}^1(f(h))| \geq \frac{q+1-2g\sqrt{q}}{|G_{f(h)}|} - \frac{\#\text{Ram}^1(M_{f(h)}/\mathbb{F}_q(t))}{2}, \quad (7)$$

where g is the genus of $M_{f(h)}$. We proceed by proving an upper bound on the size of $G_{f(h)}$, which in turn gives the wanted lower bound for $|T_{\text{split}}^1(f(h))|$.

Let \mathcal{T} be the rooted tree of height 2 with $\deg f$ branches and $\deg h$ leaves on each branch. One can easily see that $G_{f(h)}$ is a subgroup of the wreath product $\underbrace{(G_h \times \cdots \times G_h)}_{\deg f} \rtimes G_f$,

because Galois automorphisms have to preserve adjacency of the nodes of \mathcal{T} : in fact, if $\alpha \in \overline{\mathbb{F}_q(t)}$ is a root of f , then $h - \alpha$ is a factor of $f(h)$ and therefore also $h - \gamma(\alpha)$ is a factor of $f(h)$ for any $\gamma \in G_{f(h)}$. It follows that $|G_{f(h)}| \leq |G_h|^{\deg f}|G_f|$.

Combining (7) with the bound on $|G_{f(h)}|$, we obtain

$$\ell \geq \frac{q+1-2g\sqrt{q}}{|G_h|^{\deg f}|G_f|} - \frac{\#\text{Ram}^1(M_{f(h)}/\mathbb{F}_q(t))}{2}.$$

Note that the previous theorem implies that

$$\ell \geq \frac{(q+1)-2g\sqrt{q}}{|G_h|^{\deg f}|G_f|} - (\deg f)(\deg h)/2.$$

Proposition 4.2: Let $f, h \in \mathbb{F}_q[X]$ be separable polynomials such that $f - t$, $h - t$, and $f(h) - t$ have Galois groups G_f , G_h , and $G_{f(h)}$, respectively. Suppose that the splitting field $M_{f(h)}$ of $f(h) - t$ has constant field equal to \mathbb{F}_q . Then there exists an optimal HLRC, with parameters $[\deg(f(h))\ell, k, d, \deg(h) - 1, \deg(f(h)) - \lambda, \lambda]$ for any $\lambda < \deg(h)$ and

$$\ell \geq \frac{q}{|G_h|^{\deg f}|G_f|} + O(\sqrt{q}),$$

where the implied constant can be made explicit and independent of q , and the dimension k (resp. the distance d) is as in Lemma 3.3 (resp. Lemma 3.4).

Remark 3: Notice that the condition of having trivial constant field extension is automatic once there is a single totally split place, and this situation is generic if the polynomials f, h are chosen at random.

Proof: Since $M_{f(h)}$ has trivial constant field \mathbb{F}_q , Lemma 4.1 guarantees that there exist at least

$$\ell \geq \frac{1}{|G_h|^{\deg f}|G_f|}q + O(\sqrt{q})$$

totally split places, i.e. elements t_0 of \mathbb{F}_q such that $f(h) - t_0$ is totally split. Let us denote this set of t_0 's by T . Now construct the code by evaluating the polynomials in (3) at the subset \mathcal{A} of preimages of T via $f(h)$, i.e. $\mathcal{A} = (f \circ h)^{-1}(T)$, which has size $\deg(f(h))\ell$ and is a nest for the pair (f, h) by Lemma 3.1. The hierarchy is now given by the nest structure in the sense of Remark 2, and the parameters obtained from Subsection III-D. ■

V. PRACTICAL CHOICE OF PARAMETERS TO CONSTRUCT OPTIMAL HLRC

The construction we presented in the previous sections allows us to exhibit some interesting examples of HLRCs. To begin with, we consider the field \mathbb{F}_{64} . Choosing $f = h = x^3$ and $\ell = 7$, our construction gives rise to a $(63, k, d, 2, 7, 2)$ HLRC, where the values of k and d depend on the choice of s in Construction 3.2 (which is flexible). In fact, the first locality b equals $\deg h - 1 = 2$, whereas the second locality ($\rho = 5$) can be computed by following the passages of Section III-C and using the linear dependencies of the generator matrix arising from the first locality. This means that we are able to recover 1 (resp. 2) lost node(s) by looking at 2 (resp. 5) other nodes at most. We point out that the Tamo-Barg construction for availability over the field of size 64, under the same first locality assumption ($b = 2$), forces to have length 21 (with locality sets of size 3 and 7), whereas ours permits to have length 63, leading to a much better minimum distance and a larger number of servers allowed. More precisely, the Tamo-Barg construction requires the use of two orthogonal partitions, and this can be achieved by using 21 symbols corresponding to the action of x^3 and of x^7 on $\mathbb{F}_{64} \setminus \{0\}$. Note further that their construction has a larger second locality: 7, against our better parameter $\rho = 5$.

We conclude the paper with an infinite family of examples of HLRCs with some specified localities and practical parameters that can be constructed for infinitely many q 's.

Theorem 5.1: There exists an optimal $[6\ell, 3s+3, d, 2, 4, 2]$ HLRC over \mathbb{F}_q for $\gcd(q, 6) = 1$, an integer $s \in \{1, \dots, \ell-2\}$,

$$\ell = \left\lceil \frac{q+1-2\sqrt{q}}{12} - 2 \right\rceil, \text{ and } d = 6\ell - 6s - 3.$$

Proof: Let α be a non-square in \mathbb{F}_q . Apply our construction for $h = x^3 - \alpha x$, $f = x^2$ and then use Theorem 3.14 of [15] to obtain the minimum number of nests for (f, h) . The optimality follows by specializing the parameters of Definition 2.1 to the above set of parameters. ■

Notice that this theorem also shows that one can fix a nested pair for which the general lower bound in Proposition 4.2 can be improved (for some special set of parameters) for infinitely many q 's, as $|G_{f(h)}| = 12$ in this case. The study of such nested pairs is an interesting direction for future work.

VI. CONCLUSION AND FUTURE WORK

In this paper we constructed HRLCs that attain the (improved) Singleton-like bound (2) from nested F -adic expansions of polynomials. Future research directions related to the family of codes studied here include determining and classifying good nested polynomial pairs $(f, h) \in \mathbb{F}_q[x]^2$, understanding which is the minimal Galois group of $f(h) - t \in \mathbb{F}_q(t)[x]$ when f, g varies among polynomials of certain fixed degree. Whenever this Galois group is small, this type of result, combined with the results of this paper, would provide a tool to construct optimal HRLCs with large length and dimension. Another interesting research direction is extending our construction to algebraic geometric codes arising from global function fields.

REFERENCES

- [1] S. Ballantine, A. Barg, and S. Vladut, "Codes with hierarchical locality from covering maps of curves," *IEEE Trans. Inf. Theory*, vol. 65, no. 10, pp. 6056–6071, Oct. 2019.
- [2] B. Sasidharan, G. K. Agarwal, and P. V. Kumar, "Codes with hierarchical locality," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2015, pp. 1257–1261.
- [3] A. Barg, K. Haymaker, E. W. Howe, G. L. Matthews, and A. Várilly-Alvarado, "Locally recoverable codes from algebraic curves and surfaces," in *Algebraic Geometry for Coding Theory and Cryptography*. Cham, Switzerland: Springer, 2017, pp. 95–127.
- [4] A. Barg, I. Tamo, and S. Vladut, "Locally recoverable codes on algebraic curves," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2015, pp. 1252–1256.
- [5] D. Bartoli, M. Montanucci, and L. Quoos, "Locally recoverable codes from automorphism group of function fields of genus $g \geq 1$," *IEEE Trans. Inf. Theory*, vol. 66, no. 11, pp. 6799–6808, Nov. 2020.
- [6] P. Gopalan, C. Huang, B. Jenkins, and S. Yekhanin, "Explicit maximally recoverable codes with locality," *IEEE Trans. Inf. Theory*, vol. 60, no. 9, pp. 5245–5256, Sep. 2014.
- [7] G. M. Kamath, N. Prakash, V. Lalitha, and P. V. Kumar, "Codes with local regeneration and erasure correction," *IEEE Trans. Inf. Theory*, vol. 60, no. 8, pp. 4637–4660, Aug. 2014.
- [8] J. Liu, S. Mesnager, and L. Chen, "New constructions of optimal locally recoverable codes via good polynomials," *IEEE Trans. Inf. Theory*, vol. 64, no. 2, pp. 889–899, Feb. 2018.
- [9] N. Silberstein, A. S. Rawat, O. O. Koyluoglu, and S. Vishwanath, "Optimal locally repairable codes via rank-metric codes," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2013, pp. 1819–1823.
- [10] I. Tamo and A. Barg, "A family of optimal locally recoverable codes," *IEEE Trans. Inf. Theory*, vol. 60, no. 8, pp. 4661–4676, Aug. 2014.
- [11] I. Tamo, A. Barg, and A. Frolov, "Bounds on the parameters of locally recoverable codes," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3070–3083, Jun. 2016.
- [12] R. Freij-Hollanti, T. Westerbäck, and C. Hollanti, "Locally repairable codes with availability and hierarchy: Matroid theory via examples," in *Proc. Int. Zurich Seminar Commun.* Zurich, Switzerland: ETH Zurich, 2016, pp. 45–49.
- [13] I. Tamo, D. S. Papailiopoulos, and A. G. Dimakis, "Optimal locally repairable codes and connections to matroid theory," *IEEE Trans. Inf. Theory*, vol. 62, no. 12, pp. 6661–6671, Dec. 2016.
- [14] G. Zhang and H. Liu, "Constructions of optimal codes with hierarchical locality," *IEEE Trans. Inf. Theory*, vol. 66, no. 12, pp. 7333–7340, Dec. 2020.
- [15] G. Micheli, "Constructions of locally recoverable codes which are optimal," *IEEE Trans. Inf. Theory*, vol. 66, no. 1, pp. 167–175, Jan. 2020.
- [16] H. Stichtenoth, *Algebraic Function Fields and Codes*, vol. 254. Berlin, Germany: Springer, 2009.

Austin Dukes is currently pursuing the Graduate degree in applied algebra and coding theory with the University of South Florida.

Giacomo Micheli received the degree from the University of Rome "La Sapienza" in July 2012 and the Ph.D. degree (Hons.) from the Zurich Graduate School in Mathematics in October 2015, under the supervision of Prof. Joachim Rosenthal. He is currently a tenure-track Assistant Professor with the University of South Florida and the Co-Director of the Center for Cryptographic Research, USF.

Vincenzo Pallozzi Lavorante received the degree from the University of Perugia in October 2018 and the Ph.D. degree (Hons.) from the University of Modena and Reggio Emilia in February 2022, under the supervision of Prof. Massimo Giulietti. He is currently a Post-Doctoral Researcher with the University of South Florida.