

LOCALLY RECOVERABLE CODES WITH AVAILABILITY $t \geq 2$ FROM FIBER PRODUCTS OF CURVES

BETH MALMSKOG, GRETCHEN MATTHEWS, AND KATHRYN HAYMAKER

ABSTRACT. We generalize the construction of locally recoverable codes on algebraic curves given by Barg, Tamo and Vladut (CoRR, 2016) to those with arbitrarily many recovery sets by exploiting the structure of fiber products of curves. Employing maximal curves, we create several new families of locally recoverable codes with multiple recovery sets, including codes with two recovery sets from the generalized Giulietti and Korchmáros (GK) curves and the Suzuki curves, and a new locally recoverable code with many recovery sets based on the Hermitian curve, using a fiber product construction of Van der Geer and Van der Vlugt. In addition, we consider the relationship between local error recovery and global error correction, and the availability required to locally recover any pattern of a fixed number of erasures.

1. INTRODUCTION

Codes for local recovery were introduced in the context of distributed storage systems where there is a need to repair a single erasure or small number of erasures by accessing a few coordinates of the received word, rather accessing than the entire received word. A linear code C of length n over a finite field \mathbb{F} is a locally recoverable code, or LRC, with locality r if and only if for all $c \in C$ and for all $j \in [n] := \{1, \dots, n\}$ there exists $A_j \subseteq [n] \setminus \{j\}$, $|A_j| = r$, and $c_j = \phi_j(A_j)$ for some function $\phi_j : A_j \rightarrow \mathbb{F}$. The idea is that the codeword symbol c_j can be recovered from A_j without access to the other coordinates of the received word. The set A_j is called a recovery set. If one of its elements becomes unavailable, local recovery may not be possible. This leads to what is known as the availability problem and the need for multiple recovery sets.

We say that a code $C \subseteq \mathbb{F}^n$ has availability t with locality (r_1, \dots, r_t) provided for all $j \in [n]$ there exists $A_{1j}, \dots, A_{tj} \subseteq [n] \setminus \{j\}$ with $|A_{ij}| = r_i$, $A_{ij} \cap A_{kj} = \emptyset$ for $i \neq k$, and for all $c \in C$, $c_j = \phi_{ij}(A_{ij})$ for some function $\phi_{ij} : A_{ij} \rightarrow \mathbb{F}$. Such a code is called an LRC(t), or a locally recoverable code with availability t , to emphasize that each coordinate has t disjoint recovery sets.

In [3], the authors construct LRC(2)s based on fiber products of curves and propose a group-theoretic perspective on the construction, whereby a curve can be sometimes be expressed as a fiber product of its quotient curves by certain subgroups of the automorphism group of the curve. In particular, given a curve \mathcal{X} with automorphism group $Aut(\mathcal{X})$ that is a semi-direct product of two subgroups, an LRC(2) can be formed by considering the subfields of the field of functions on \mathcal{X} that are fixed by the subgroups. It is remarked that both perspectives can be extended in a straightforward way to provide multiple recovery sets, meaning $t > 2$.

In Section 2, we carry out the generalization to general LRC(t) for $t \geq 2$, employing a slightly different approach to bounding the code parameters than in [3]. This results in new, potentially sharper bounds on minimum distance. Properties of multiple local recovery sets




FIGURE 1. Curves for locally recoverable codes with locality t

are considered in Section 3. In Section 4, we employ generalized GK curves in the fiber product construction. While the construction from GK curves is quite explicit, the recovery sets are unbalanced, meaning there is a large difference between the cardinalities of the two recovery sets associated with a particular coordinate; this may adversely impact other code parameters. We illustrate the automorphism group perspective in Section 5, constructing theoretical codes on the Suzuki curve from its automorphism groups that provide balanced recovery sets. However, this section highlights the difficulty of using the automorphism group construction to explicitly construct LRCs: even with knowledge of the quotient curves, constructing the necessary maps can be quite difficult. Finally, in Section 6, we obtain locally recoverable codes on the Hermitian curve \mathcal{H}_{p^t} with availability t for $t \geq 2$ which are explicit in nature and balanced, using the full power of the LRC(t) construction. Our findings are summarized in Section 7.

2. FIBER PRODUCT CONSTRUCTION FOR LRC(t)

Let $\mathcal{X}, \mathcal{Y}, \mathcal{Y}_1, \mathcal{Y}_2, \dots, \mathcal{Y}_t$ be smooth projective absolutely irreducible algebraic curves over a finite field \mathbb{F} with rational, separable maps $h_i : \mathcal{Y}_i \rightarrow \mathcal{Y}$ and projection maps $g_i : \mathcal{X} \rightarrow \mathcal{Y}_i$ for $1 \leq i \leq t$ and a rational, separable map $g : \mathcal{X} \rightarrow \mathcal{Y}$ of degree d_g , so that the diagram in Figure 1 commutes.

Fix $i, 1 \leq i \leq t$. Notice that g, h_i , and g_i give rise to maps

$$\begin{aligned} g^* : \mathbb{F}(\mathcal{Y}) &\rightarrow \mathbb{F}(\mathcal{X}) \\ f &\mapsto f \circ g, \\ h_i^* : \mathbb{F}(\mathcal{Y}) &\rightarrow \mathbb{F}(\mathcal{Y}_i) \\ f &\mapsto f \circ h_i, \end{aligned}$$

and

$$\begin{aligned} g_i^* : \mathbb{F}(\mathcal{Y}_i) &\rightarrow \mathbb{F}(\mathcal{X}) \\ f &\mapsto f \circ g_i. \end{aligned}$$

One may consider g^*, h_i^* and g_i^* as embeddings so that we have

$$\mathbb{F}(\mathcal{Y}) \rightarrow h_i^*(\mathbb{F}(\mathcal{Y})) \hookrightarrow \mathbb{F}(\mathcal{Y}_i) \rightarrow g_i^*(\mathbb{F}(\mathcal{Y}_i)) \hookrightarrow \mathbb{F}(\mathcal{X}).$$

Since the diagram commutes, \mathcal{X} is a cover of the fiber product $\mathcal{Y}_1 \times_{\mathcal{Y}} \mathcal{Y}_2 \times_{\mathcal{Y}} \dots \times_{\mathcal{Y}} \mathcal{Y}_t$. If we assume that $\mathbb{F}(\mathcal{Y})$ is the compositum of the fields $g_i^*(\mathbb{F}(\mathcal{Y}_i))$ and

$$g^*(\mathbb{F}(\mathcal{Y})) = \bigcap_{i=1}^t g_i^*(\mathbb{F}(\mathcal{Y}_i)),$$

then we have $\mathcal{X} = \mathcal{Y}_1 \times_{\mathcal{Y}} \mathcal{Y}_2 \times_{\mathcal{Y}} \dots \times_{\mathcal{Y}} \mathcal{Y}_t$.

Let $x_i \in \mathbb{F}(\mathcal{Y}_i)$ so that

$$\begin{array}{c} \mathbb{F}(\mathcal{Y}_i) = h_i^*(\mathbb{F}(\mathcal{Y}))(x_i) \\ \left| \begin{array}{c} d_{h_i} \\ \hline \end{array} \right. \\ \mathbb{F}(\mathcal{Y}) \cong h_i^*(\mathbb{F}(\mathcal{Y})) \end{array}$$

where x_i is the root of a degree d_{h_i} polynomial $b_i(t) \in g_i^*(\mathbb{F}(\mathcal{Y}))[t]$. For convenience, denote $g_i^*(x_i) \in \mathbb{F}(\mathcal{X})$ by x_i^* . Let D_i be the divisor of the function $x_i^* \in \mathbb{F}(\mathcal{X})$. Let $D_i = D_{i,+} - D_{i,-}$, where $D_{i,+}$ and $D_{i,-}$ are both effective; that is, $D_{i,+} = (x_i^*)_0$ is the zero divisor of x_i^* , and $D_{i,-} = (x_i^*)_\infty$ is the pole divisor of x_i^* . Let $\deg(D_{i,-}) = d_{x_i}$ be the degree of the function x_i from $\mathcal{Y}_i \rightarrow \mathbb{P}^1$. Then, if x_i^* is viewed as a function from $\mathcal{X} \rightarrow \mathbb{P}^1$, its degree is $d_{g_i} d_{x_i}$.

We now have that

$$g^*(\mathbb{F}(\mathcal{Y}))(x_1^*, x_2^*, \dots, x_t^*) = \mathbb{F}(\mathcal{X}).$$

Define the natural maps

$$\begin{array}{l} \tilde{g}_i : \mathcal{X} \rightarrow \mathcal{Y}_1 \times_{\mathcal{Y}} \cdots \times_{\mathcal{Y}} \mathcal{Y}_{i-1} \times_{\mathcal{Y}} \mathcal{Y}_{i+1} \times_{\mathcal{Y}} \cdots \times_{\mathcal{Y}} \mathcal{Y}_t \\ P \mapsto (g_1(P), \dots, g_{i-1}(P), g_{i+1}(P), \dots, g_t(P)) \end{array}$$

and

$$\begin{array}{l} \tilde{h}_i : \mathcal{Y}_1 \times_{\mathcal{Y}} \cdots \times_{\mathcal{Y}} \mathcal{Y}_{i-1} \times_{\mathcal{Y}} \mathcal{Y}_{i+1} \times_{\mathcal{Y}} \cdots \times_{\mathcal{Y}} \mathcal{Y}_t \rightarrow \mathcal{Y} \\ (P_1, P_2, \dots, P_{i-1}, P_{i+1}, \dots, P_t) \mapsto h_j(P_j) \end{array}$$

for any $j \neq i$. Then $\mathcal{X} = \mathcal{Y}_i \times_{\mathcal{Y}} \tilde{g}_i(\mathcal{X})$. The degree of \tilde{g}_i must be equal to the degree of h_i , denoted d_{h_i} .

Let $\tilde{D} = \sum_{i=1}^t h_i(D_i)$. Let D be any effective divisor on \mathcal{Y} with $\text{supp}(D) \subset \text{supp}(\tilde{D})$. Let $\deg(D) = l$. Now, set $S = \{P_1, \dots, P_s\} \subset \mathcal{Y}(\mathbb{F})$ such that

$$|g^{-1}(P_i) \cap \mathcal{X}(\mathbb{F})| = d_g,$$

for all $i > 0$, i.e. so that g is unramified above all elements of S and the entire fiber is contained in $\mathcal{X}(\mathbb{F})$, and $\{P_1, \dots, P_s\} \cap \text{Supp}(D) = \emptyset$. We require that $l < s$ so that for all $f \in \mathcal{L}(D)$, there exists some $P_i \in S$ with $f(P_i) \neq 0$. Then set

$$B = g^{-1}(S).$$

Note that $n = |B| = d_g s$. For $P \in B$, set

$$B^{(i)}(P) = \tilde{g}_i^{-1}(\tilde{g}_i(P))$$

and

$$A^{(i)}(P) = B^{(i)}(P) \setminus \{P\}.$$

Then let

$$V = \text{Span}\{f_j^* x_1^{*e_1} \cdots x_t^{*e_t} : 0 \leq e_i \leq d_{h_i} - 2 \text{ for all } i, 1 \leq j \leq m\}.$$

Set $C(D, B) := \text{Im}(ev_B)$ where

$$\begin{array}{l} ev_B : V \rightarrow \mathbb{F}^n \\ f \mapsto (f(P))_{P \in B}. \end{array}$$

Given an erasure in the coordinate associated with P , each $A^{(i)}(P)$ acts as a recovery set, because on the set $A^{(i)}(P)$ the function f is constant except in x_i^* , so on $A^{(i)}(P)$ it acts as $\tilde{f}(x_i^*)$, a polynomial of degree less than or equal to $d_{h_i} - 2$. It must also be that there are no

two points $P_a \neq P_b$ in $A^{(i)}(P)$ where $x_i^*(P_a) = x_i^*(P_b)$. The $d_{h_i} - 1$ points of $A^{(i)}(P)$ therefore give rise to $d_{h_i} - 1$ distinct pairs $(x_i^*(P), \tilde{f}(x_i^*(P)))$. Since any polynomial of this degree is determined by its values on $d_{h_i} - 1$ points, these pairs are sufficient to determine the value of \tilde{f} on P .

This construction gives rise to the following theorem.

Theorem 1. *Let D, S , and B be as described above, where $l = \deg(D) \leq |S|$, and $m = \ell(D)$. Then the code $C(D, B)$ is an LRC(t) with length $n = |B|$, dimension $m(d_{h_1} - 1)(d_{h_2} - 1) \cdots (d_{h_t} - 1)$, minimum distance $d \geq n - ld_g - \sum_{i=1}^t (d_{h_i} - 1)(d_{g_i} d_{x_i})$, and locality $((d_{h_1} - 1), (d_{h_2} - 1), \dots, (d_{h_t} - 1))$.*

Proof. The length of $C(D, B)$ is given by definition. The minimum distance is bounded by considering the number of zeros that a function $f \in V$ can have on the curve \mathcal{X} . Note that

$$f_j^* x_1^{*e_1} x_2^{*e_2} \cdots x_t^{*e_t} \in \mathcal{L} \left(g^{-1}(D) + \sum_{i=0}^t (d_{h_i} - 2) g_i^{-1}(D_{i,-}) \right) \subset \mathbb{F}(\mathcal{X}),$$

where D has degree l and $g^{-1}(D)$ has degree $l(d_g)$ and $g_i^{-1}(D_{i,-})$ has degree equal to $d_{g_i} d_{x_i}$. Thus the function f can have at most $ld_g + \sum_{i=1}^t (d_{h_i} - 2)d_{g_i} d_{x_i}$ poles on \mathcal{X} (counted with multiplicity), so f may have at most this many zeros on \mathcal{X} (counted with multiplicity). The dimension is as stated because the evaluation map is injective for e_i in the ranges described. Given an erasure in the coordinate associated with P , each of $A^{(i)}(P)$ act as recovery sets, where the missing coordinate can be determined as described above by polynomial interpolation on the appropriate x_i^* .

□

Remark For $\rho > 1$ an integer, let

$$V_\rho = \text{Span}\{f_j^* x_1^{*e_1} \cdots x_t^{*e_t} : 0 \leq e_i \leq d_{h_i} - 1 - \rho \text{ for all } i, 1 \leq j \leq m\}.$$

As observed in [3], applying the same construction with V_ρ in place of V allows up to ρ erasures to be recovered by each recovery set, increasing the local distance of $C(D, B)$. However, this also reduces the dimension of the code to $m(d_{h_1} - \rho)(d_{h_2} - \rho) \cdots (d_{h_t} - \rho)$, so this must be seen as a tradeoff between dimension and recovery. Throughout the rest of the paper, we use the maximal dimension construction from Theorem 1, but the modification above may be made to recover more erasures if desired.

There are (at least) two natural ways in which the setting in Theorem 1 arises, and they are described in the next two results.

Corollary 2. *Let $\mathcal{Y}, \mathcal{Y}_1, \dots, \mathcal{Y}_t$ be smooth curves with separable maps $h_i : \mathcal{Y}_i \rightarrow \mathcal{Y}$. Then Theorem 1 applies, giving locally recoverable codes from the fiber product $\mathcal{X} = \mathcal{Y}_1 \times_{\mathcal{Y}} \cdots \times_{\mathcal{Y}} \mathcal{Y}_t$.*

Corollary 3. *Suppose that \mathcal{X} is a curve such that $\text{Aut}(\mathcal{X})$ has subgroups T_1, \dots, T_t so that the subgroup generated by these groups is an associative semi-direct product within $\text{Aut}(\mathcal{X})$. Then Theorem 1 applies to give locally recoverable codes from \mathcal{X} by taking \mathcal{Y}_i to be the associated quotient curves \mathcal{Y}/T_i and $\mathcal{Y} := \mathcal{X}/(T_1 \rtimes \cdots \rtimes T_t)$.*

Remark For $t > 2$, the requirement that the semi-direct product is associative is far from trivial. Corollaries 2 and 3 are applied in Sections 4 and 5, as we construct LRC(2)s from generalized GK curves and Suzuki curves. One difficulty with the ‘top down’ approach of Corollary 3 is explicitly writing down equations for \mathcal{Y}_i , the functions x_i , and the spaces $L(D)$

when starting with a given model of \mathcal{X} . The ‘bottom-up’ perspective of Corollary 2 is also applied in Section 6 to explicitly construct codes with arbitrarily many recovery sets by building a fiber product of curves to obtain the top curve.

The curves considered in this paper are maximal, meaning their numbers of rational points meet the upper Hasse-Weil bound,

$$q + 1 - 2g_C\sqrt{q} \leq |\mathcal{C}(\mathbb{F}_q)| \leq q + 1 + 2g_C,$$

where g_C is the genus of the curve \mathcal{C} . Maximal curves have played an important role in coding theory, especially in the construction of algebraic geometry codes as they support the construction of long codes. Their utility in the construction of locally recoverable codes comes from this as well, specifically if $|\mathcal{X}(\mathbb{F})|$ is large, then the hope is that one can find appropriate map $g : \mathcal{X} \rightarrow \mathcal{Y}$ so that $|\mathcal{B}|$ is large. Moreover, some families of maximal curves (such as Hermitian and Suzuki) have large automorphism groups which may allow for more choices when constructing recovery sets (equivalently the maps $g_i : \mathcal{Y}_i \rightarrow \mathcal{Y}$). However, we see in Section 5 that the quotient group construction is very difficult, even for a well-studied curve like the Suzuki curve where models of quotient curves are known.

3. LRC(t) CODES, ERASURE RECOVERY, AND ERROR CORRECTION

As described in Section 1, locally recoverable codes have been developed with the primary goal of facilitating convenient recovery of erasures, potentially created by server failure in distributed storage systems. The construction for $C(D, B)$ described in Section 2 results in t transverse recovery sets for each location, meaning if any one location is erased there are t disjoint recovery sets that would allow recovery of the erasure. Moreover, since recovery sets are defined by fibers over points of the curves $\tilde{\mathcal{Y}}_i$, the point P is in the i th recovery set for the point Q if and only if Q is in the i th recovery set for the point P ; i.e. for all $i \in [n]$,

$$P \in A^{(i)}(Q) \Leftrightarrow Q \in A^{(i)}(P).$$

In fact, there are $P'_1, \dots, P'_{s'} \in \mathcal{X}(\mathbb{F})$ such that $\{B^{(i)}(P) : P \in \{P'_1, \dots, P'_{s'}\}\}$ defines a partition of $[n]$. We say that such recovery sets are symmetric. This section addresses two questions that arise from the recovery procedure for symmetric recovery sets. First, what benefit is provided by having a large number of recovery sets? Second, how does the number of recovery sets relate to minimum distance and global error correction?

First, we note that for erasure of a small fraction of locations in known positions, it can be much more efficient to use local recovery over global error correction, even if multiple local recovery sets are required. Consider a locally recoverable code C of length n and availability t with symmetric recovery sets. Say that f erasures occur. For global error correction, we assume that all $n - f$ known locations must be consulted. In general, we see that if a code has locality t and symmetric recovery sets, then the fact that recovery sets are transverse implies that $\text{lcm}\{(r_i + 1) : 1 \leq i \leq t\}$ divides n . Let $c_t(f)$ be the maximum number of positions that need to be consulted to recover f erasures using local recovery. If there is a pattern of f erasures that may not be locally recoverable by a code of availability t , then $c_t(f) = \infty$.

If only one position P has been erased and C admits a single local recovery set of size r_1 for each position, then r_1 positions may be consulted to recover P , and $c_1(1) \leq r_1$. We may assume $n = s(r_1 + 1)$ for some integer $s > 1$, so local recovery saves at least $(s - 1)(r_1 + 1)$ consultations.

Now consider the possibility that two positions P and Q were erased. If C has only one recovery set for each position and the recovery set of P contains Q (and vice versa, by symmetry), then P and Q are not recoverable using that recovery set. However, if C has availability 2, then C has 2 transverse recovery sets for each position, and the fact that the first recovery set of P contains Q implies that the second recovery set does not (and similarly for Q). If C has locality (r_1, r_2) , then P and Q can both be recovered with at most $r_1 + r_2$ consultations. We have saved at least $(s(r_1 + 1) - 2) - (r_1 + r_2)$ consultations.

Without loss of generality, assume $r_1 = \max_{1 \leq i \leq t} \{r_i\}$. For a general code of availability t , assuming that t is sufficiently large to recover f erasures, recovery will require at most f recovery sets be consulted, resulting in a total number of

$$c_t(f) \leq fr_1$$

consultations in local recovery, in comparison with at least $s(r_1 + 1) - f$ consultations for global error correction. This also assumes that the minimum distance of the code is sufficiently large to correct f errors, which is not necessarily implied by the construction. Clearly, if s is close in size to f , then there is little difference in the numbers of consultations required for local recovery and for error correction. However, our construction generally results in relatively large s . In all our examples, we have that $\prod_{i=1}^t (r_i + 1)$ divides n , so the savings are significant for moderate size f .

In fact, the interplay between parameters becomes complicated as the number of potential failures grows. For example, consider the following lemma. Since the recovery sets are symmetric, we may denote the i th recovery set for a coordinate position P as $A^{(i)}(P)$, and define $B^{(i)}(P) = A^{(i)}(P) \cup \{P\}$.

The following lemma demonstrates that more than 2 recovery sets become necessary rather quickly.

Lemma 4. *Let $b(f)$ be the availability required for a locally recoverable code C with symmetric recovery sets to be capable of locally recovering any pattern of f erasures. Then*

- $b(1) = 1$,
- $b(2) = 2$,
- $b(3) = 2$, and
- $b(4) > 3$.

Proof. It is not difficult to see that $b(1) = 1$ and $b(2) = 2$ based on the situations described above. To see that $b(3) = 2$, assume that positions P , Q , and R have been erased. Availability $t = 1$ is clearly not sufficient, because it could be that all $B^{(1)}(P) = B^{(1)}(Q) = B^{(1)}(R)$. However, $t = 2$ is sufficient; if $B^{(1)}(P) = B^{(1)}(Q) = B^{(1)}(R)$, then $B^{(2)}(P) \neq B^{(2)}(Q)$ and $B^{(2)}(P) \neq B^{(2)}(R)$. Thus, P may be recovered from the recovery set $A^{(2)}(P)$, leaving 2 points, which can be recovered with 2 recovery sets since $b(2) = 2$. If $B^{(1)}(P) \neq B^{(1)}(Q)$ but $B^{(1)}(P) = B^{(1)}(R)$, then Q can be recovered from $B^{(1)}(Q)$, again leaving 2 points.

To see that $b(4) > 3$, consider the following scenario. Positions $\{P_1, P_2, P_3, P_4\}$ have been erased, and we have that

- $B^{(1)}(P_1) = B^{(1)}(P_2)$ and $B^{(1)}(P_3) = B^{(1)}(P_4)$,
- $B^{(2)}(P_1) = B^{(2)}(P_3)$ and $B^{(2)}(P_2) = B^{(2)}(P_4)$,
- $B^{(3)}(P_1) = B^{(3)}(P_4)$ and $B^{(3)}(P_2) = B^{(3)}(P_3)$.

Clearly, none of the three recovery sets can be used to recover any of the erased positions. Therefore $b(4) > 3$. □

The values of $b(i)$ are for $i > 4$ are unknown. Preliminary computations indicate that $b(4) = b(5) = b(6) = b(7) = 4$ and $b(8) = 8$.

Note that the recovery procedure and all discussion to this point assumes that the positions of the erased locations are known. This is not the case in general error correction, so a code that is capable of restoring f erasures may not be capable of correcting f global errors. Minimum distance bounds give some indication of the global error correction capacity of the code. The minimum distance bounds in Theorem 1 are based on the divisors of certain functions. For theoretical purposes, however, we would like to bound minimum distance without explicitly constructing the given field extensions as in Section 2. More generally, one might wish to know if an LRC(t) is actually an error-correcting code, regardless of its construction. It is interesting to consider how the presence of recovery sets influences the potential error-correcting capability of a locally recoverable code, independent of how the code itself is defined. With this in mind, some very modest bounds on minimum distance can be derived by considering that, by construction, there exists an algorithm for correcting a certain number of erasures.

Lemma 5. *Let C be a locally recoverable code of availability t with symmetric recovery sets and minimum distance d . Then*

- if $t \geq 1$, then $d \geq 2$,
- if $t \geq 2$, then $d \geq 3$.

Proof. Say that $t = 1$. Then, assume a codeword w has one error, in position P , resulting in the word w' . On the each recovery set, the entries in the codeword w corresponds to the evaluation of a polynomial function of degree $\leq (r - 1)$. Checking each recovery set for the word w' , one of these recovery sets, $B^{(1)}(P)$, will fail to be the evaluation of such a polynomial. If this were not true, and the error position had actually been erased, the recovery would not have been unique. Therefore any one error can be detected. Therefore the minimum distance of C must be at least 2.

Say that $t = 2$. Then, assume a codeword w has one error, in position P , resulting in the word w' . Since there is only one error in w' , the set $B^{(1)}(P)$ contains one error, and thus there is no polynomial of degree $\leq (r_1 - 1)$ which evaluates to the entries in these locations. Similarly, $B^{(2)}(P)$ contains one error. Therefore the error must be in a position in $B^{(1)}(P) \cap B^{(2)}(P)$, which is just $\{P\}$ by the transversality of recovery sets. Since the position of the error is known, the error can be erased and recovered using either recovery set. Since any one error can be corrected, it must be that the minimum distance of C is at least 3. □

Define $e(t)$ to be the maximum number of erasures that can be recovered in the worst case erasure configuration, using a code with availability t and symmetric transverse recovery sets. The function $e(t)$ is closely related to $b(f)$, although they are not inverses. The following values of $e(t)$ are based on Lemma 4 and the remarks following it:

$$e(1) = 1, e(2) = e(3) = 3, e(4) = e(5) = e(6) = e(7) = 4, e(8) = 8.$$

Using this function, we obtain a general version of Lemma 5.

Lemma 6. *Let C be a locally recoverable code of availability t with symmetric recovery sets and minimum distance d . Then $d \geq e(t) + 1$.*




FIGURE 2. Generalized GK curve as a fiber product.

Proof. By definition, C can recover $e(t)$ erasures. A code that can recover any pattern of $e(t)$ erasures must have minimum distance at least $e(t) + 1$. Indeed, if d were equal to $e(t)$, then there would be a codeword of weight $e(t)$ which would be indistinguishable from the all-zeros codeword upon the erasure of all of its nonzero entries. Therefore $d > e(t)$. \square

4. LRC(2)S ON GENERALIZED GIULIETTI-KORCHMAROS CURVES

Let $q = p^h$ for p a prime, and let $N \geq 3$ be an odd natural number. We consider the family of generalized Giulietti-Korchmaros (GK) curves \mathcal{C}_N , which are maximal over the field $\mathbb{F}_{q^{2N}}$ [5, 6]. The curve \mathcal{C}_N is the normalization of the intersection of two surfaces \mathcal{H}_q and \mathcal{X}_N in \mathbb{P}^3 , defined by the following affine equations:

$$\mathcal{H}_q : x^q + x = y^{q+1}$$

$$\mathcal{X}_N : y^{q^2} - y = z^{\frac{q^N+1}{q+1}}.$$

The intersection of these surfaces has single point at infinity, denoted by ∞ , which is a cusp singularity for $N > 3$ and is smooth when $N = 3$. The curve is smooth elsewhere. The number of $\mathbb{F}_{q^{2N}}$ -rational points on \mathcal{C}_N is

$$\#\mathcal{C}_N(\mathbb{F}_{q^{2N}}) = q^{2N+2} - q^{N+3} + q^{N+2} + 1.$$

The curve \mathcal{C}_N can also be defined as the normalized fiber product over \mathbb{P}^1 of the two curves in \mathbb{P}^2 given by the same equations. As curves, both \mathcal{H}_q and \mathcal{X}_N are both maximal over the field $\mathbb{F}_{q^{2N}}$ [1], each with a single point at infinity, denoted by $\infty_{\mathcal{H}_q}$ and $\infty_{\mathcal{X}_N}$ respectively. Let ∞_y denote the single point at infinity on \mathbb{P}_y^1 . Define $h_1 : \mathcal{X}_N \rightarrow \mathbb{P}_y^1$ to be the natural degree $\frac{q^N+1}{q+1}$ projection map onto the y coordinate for affine points, with $\infty_{\mathcal{X}_N} \mapsto \infty_y$. Similarly, let $h_2 : \mathcal{H}_q \rightarrow \mathbb{P}_y^1$ be the natural degree q projection map onto the y coordinate for affine points, with $\infty_{\mathcal{H}_q} \mapsto \infty_y$. We then have the fiber product construction depicted in Figure ??.

Then $\mathcal{C}_N = \widetilde{\mathcal{X}_N \times \mathcal{H}_q}$, the normalization of the fiber product described above. The map $g : \mathcal{C}_N \rightarrow \mathbb{P}_y^1$ has degree $d_g := \frac{q(q^N+1)}{q+1}$ and is ramified above ∞_y and a_y with $a \in \mathbb{F}_{q^2}$, where a_y denotes the point on \mathbb{P}_y^1 with $y = a$ [?, 5]. Notice that $d_{g_1} = d_{h_2} = q$ and $d_{g_2} = d_{h_1} = \frac{q^N+1}{q+1}$.

To construct an LRC(2), we use the commutative diagram in Figure ?? and the construction detailed in Section 2. The degree of the function $x : \mathcal{H}_q \rightarrow \mathbb{P}^1$ is $d_x = q + 1$. The degree of the function $z : \mathcal{X}_N \rightarrow \mathbb{P}^1$ is $d_z = q^2$. We take the divisor Q to be ∞_y and choose a parameter l so that $D = l\infty_y$.

Theorem 7. *The locally recoverable code $C(l\infty_y, B)$ constructed from the generalized GK curve C_N as described above is an (n, k, d) code over $\mathbb{F}_{q^{2N}}$ with availability 2 and locality $\left(q, \frac{q^N+1}{q+1}\right)$ where*

$$\begin{aligned} n &= q^{2N+2} - q^{N+3} + q^{N+2} - q^3, \\ k &= \left(\frac{q^N+1}{q+1} - 1\right)(q-1)(l+1), \\ d &\geq n - l \left(\frac{q(q^N+1)}{q+1}\right) - \left((q^N+1)(q-2) + q^3 \left(\frac{q^N+1}{q+1} - 2\right)\right), \end{aligned}$$

and l is any positive integer yielding $0 < k < n$, $l < q^{N+2} + q^{N+1} - q - 1$, and $d > 0$.

Proof. The set B of evaluation points consists of points on the curve C_N that are not above ramification points in the ground curve, \mathbb{P}_y^1 , over the field $\mathbb{F}_{q^{2N}}$. This means we let $S = g(B)$. Since

$$|g^{-1}(\{\infty_y, a_y : a \in \mathbb{F}_{q^2}\})| = q^3 + 1,$$

we calculate the size of the evaluation set B :

$$|B| = \#\mathcal{C}_N - (q^3 + 1) = q^{2N+2} - q^{N+3} + q^{N+2} + 1 - (q^3 + 1) = q^{2N+2} - q^{N+3} + q^{N+2} - q^3 = sd_g,$$

where $s = |S| = q^2(q^{N-1} - 1)(q+1) = q^{N+2} + q^{N+1} - q - 1$.

Let $D = l\infty_y$. Since the genus of $\mathcal{Y} = \mathbb{P}^1$ is 0, we know by the Riemann-Roch theorem that $\ell(D) = l + 1$, and we can realize these functions as polynomials in y of degree bounded by l . The set of evaluation functions for the code is denoted by V , where

$$V = \text{Span} \left\{ x^i z^j y^\kappa : 0 \leq i \leq \frac{q^N+1}{q+1} - 2, 0 \leq j \leq q-2, 0 \leq \kappa \leq l, \text{ and } i, j, \kappa \in \mathbb{Z} \right\}.$$

Then, by the construction in Theorem 1, we obtain a code with the claimed attributes. \square

Given the choice of l , we can construct codes with rate close to one, where the tradeoff is low minimum distance:

$$R_{N,l} \geq \frac{\left(\frac{q^N+1}{q+1} - 1\right)(q-1)(l+1)}{|B|},$$

where for maximal l and increasing N we have

$$\lim_{N \rightarrow \infty} R_{N,l} = 1.$$

Example Taking $N = 3$, we consider codes from the curves C_3 over the field \mathbb{F}_{q^6} . We obtain blocklength $|B| = q^3(q-1)(q^2-q+1)(q+1)^2 = q^8 - q^6 + q^5 - q^3$, as above, and the following bounds on the dimension and minimum distance:

$$\begin{aligned} k &= (q-1)(q^2-q)(l+1), \\ d &\geq n - lq(q^2-q+1) - (q^3+1)(q-2) - q^3(q^2-q-1). \end{aligned}$$

In Table 1, we consider $q = 3$ and provide bounds on the code parameters for different values of l . \square

l	$k \geq$	$d \geq$
270	3252	215
260	3132	425
250	3012	635
240	2892	845
230	2772	1055
220	2652	1265
210	2532	1475

TABLE 1. The generalized GK curves \mathcal{C}_3 over \mathbb{F}_{729} produce LRC(2)s of length $n = 6048$, with $N = 3$, $q = 3$, $r_1 = 6$, $r_2 = 2$, and $D = l\infty_y$.

In this section, we employed generalized GK curves to obtain LRC(2)s over $\mathbb{F}_{q^{2N}}$ with recovery sets of sizes q and $\frac{q^N+1}{q+1}$. While this addresses the availability problem, it leads to recovery sets of very different sizes. In the next section, we construct LRC(2)s with recovery sets which are more balanced in size.

5. LRC(2)S ON SUZUKI CURVES

Let $a \in \mathbb{N}$, $q_0 = 2^a$, and $q = 2q_0^2$. The Suzuki group $Sz(q)$ can be realized as a subgroup of $GL_4(q)$ as follows. Let $a, c, d \in \mathbb{F}_q$, $d \neq 0$, and define

$$T_{a,c} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ a & 1 & 0 & 0 \\ c & a^{2q_0} & 1 & 0 \\ a^{2q_0+2} + ac + c^{2q_0} & a^{2q_0+1} + c & a & 1 \end{pmatrix}, \quad M_d = \begin{pmatrix} d^{-q_0-1} & 0 & 0 & 0 \\ 0 & d^{-q_0} & 0 & 0 \\ 0 & 0 & d^{q_0} & 0 \\ 0 & 0 & 0 & d^{q_0+1} \end{pmatrix},$$

and

$$W = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}.$$

Let $T = \{T_{a,c} : a, c \in \mathbb{F}_q\}$ and $M = \{M_d : d \in \mathbb{F}_q^*\}$. Then $Sz(q) = \langle M, T, W \rangle$.

The Suzuki curve \mathcal{S}_q is the Deligne-Lusztig curve with automorphism group $Sz(q)$. The curve \mathcal{S}_q has a model in \mathbb{P}^2 with affine equation

$$y^q + y = x^{q_0}(x^q + x).$$

Note that this model is singular. The genus of \mathcal{S}_q is $q_0(q-1)$ and it has q^2+1 points over \mathbb{F}_q , making it optimal over the field. Over \mathbb{F}_{q^4} , \mathcal{S}_q is maximal, attaining the upper Weil bound [8]. Smooth models for \mathcal{S}_q in higher dimensional spaces have been determined [2, 4]. As in [7], a convenient model in \mathbb{P}^4 can be defined by the affine equations

$$\begin{aligned} y^q + y &= x^{q_0}(x^q + x), \\ z &= x^{2q_0+1} + y^{2q_0}, \end{aligned}$$

and

$$w = xy^{2q_0} + z^{2q+1}.$$




FIGURE 3. Suzuki curve and its quotients used for constructing LRC(2) with balanced recovery sets

Let $[U : X : Y : Z : W]$ be a set of projective coordinates for \mathbb{P}^4 , where for all $U \neq 0$ we have affine coordinates given by

$$x = \frac{X}{U}, \quad y = \frac{Y}{U}, \quad z = \frac{Z}{U}, \quad w = \frac{W}{U}.$$

For any $(x, y) \in \mathbb{F}_q^2$ satisfying $y^q + y = x^{q_0}(x^q + x)$, let $P_{(x,y)}$ denote the point $[1 : x : y : z : w] \in \mathcal{S}_q(\mathbb{F}_q)$. Let $P_\infty = [0 : 0 : 0 : 0 : 1] \in \mathcal{S}_q(\mathbb{F}_q)$.

Theorem 8. *There is are locally recoverable codes $C(D, B)$ with availability 2, and locality $(q - 1, q - 2)$ on the Suzuki curve \mathcal{S}_q*

- (1) *over \mathbb{F}_q , with length $n = q(q - 1)$ and dimension $k \geq (q - 1)(q - 2)$; and*
- (2) *over \mathbb{F}_{q^4} , with length $n = q(q - 1)(q^2 + 2qq_0 + q + 1)$ and dimension $k \geq (q - 1)(q - 2)$.*

Proof. Let $T_0 = \{T_{0,c} : c \in \mathbb{F}_q\} \leq T$. It is straightforward to compute that $M_d T_{0,c} M_d^{-1} = T_{0,cd^{2q_0+1}}$, so M normalizes T_0 . Thus T_0 is a normal subgroup of $G = \langle M, T_0 \rangle$, and so $G = T_0 \rtimes M$. Thus $G = (T_0 \rtimes M) \leq Sz(q)$.

Let $\mathcal{Y}_1 = \mathcal{S}_q/T_0$. The accompanying natural map $g_1 : \mathcal{S}_q \rightarrow \mathcal{Y}_1$ is degree q and fully ramified at a single point $P_\infty \in \mathcal{S}_q(\mathbb{F}_q)$ [7, Theorem 6.1]. Then, by [7, Theorem 6.6], \mathcal{Y}_1 has genus 0, and has an affine model as given in [7].

Let $\mathcal{Y}_2 = \mathcal{S}_q/M$. By [7, Theorem 4.1], the natural map $g_2 : \mathcal{S}_q \rightarrow \mathcal{Y}_2$ is degree $q - 1$ and fully ramified below two points, $P_{(1,0)}$ and $P_\infty \in \mathcal{S}_2(\mathbb{F}_q)$. Then \mathcal{Y}_2 has genus q_0 , and has an affine model as given in [7].

Let $\mathcal{Y} = \mathcal{S}_q/G$. Since \mathcal{Y}_1 covers \mathcal{Y} , it must be that \mathcal{Y} has genus 0 as well. Let $g : \mathcal{S}_q \rightarrow \mathcal{Y}$ be the accompanying natural map. We then have the diagram of curves as shown in Figure 3.

In Case (1), let $B = \mathcal{S}_q(\mathbb{F}_q) \setminus \{P_\infty, g^{-1}(g(P_{(1,0)}))\}$. Then $n = |B| = |\mathcal{S}_q(\mathbb{F}_q)| - (q + 1) = q(q - 1)$.

In Case (2), let $B = \mathcal{S}_q(\mathbb{F}_{q^4}) \setminus \{P_\infty, g^{-1}(g(P_{(1,0)}))\}$. Then $n = |B| = |\mathcal{S}_q(\mathbb{F}_{q^4})| - (q + 1) = q(q - 1)(q^2 + 2qq_0 + q + 1)$. □

Remark In both Case (1) and Case (2), the bound on the minimum distance from Theorem 1 applies. However, without explicit generators for the given function field extensions, one cannot determine the degrees of the functions x_i to give an explicit bound on d . Hence, some work around is required. In the construction from Theorem 1, we may choose $l = 0$, so $\mathcal{L}(D)$ is simply the set of constant functions on \mathcal{Y} , giving $m = 1$ and yielding the stated dimension. Therefore we use the very modest bound from Lemma 5 to guarantee $d \geq 3$.

In Case (2), the construction from Theorem 1 allows a larger range of parameters. Since \mathcal{Y} has genus 0, we may let $D = lP_\infty$, and assume that functions in $L(D)$ are represented by polynomials of degree less than or equal to $q - 2$. Polynomials of larger degree will not be distinct when evaluated on points defined over \mathbb{F}_q . We then have maximal dimension given by $l = q - 2$. Then $m = \ell(D) = q - 1$, so $k = (q - 1)^2(q - 2)$ by Theorem 1. Again, we use the bound from Lemma 5 for distance to see that $d \geq 3$.

This illustrates a drawback of the ‘top-down’ construction presented in Corollary 3 in that while the codes $C(D, B)$ might have reasonable classical parameters n , k , and d , one might not have access to the information needed to provide a good estimate of the minimum distance d .

Remark There is no benefit to considering the code defined over \mathbb{F}_{q^2} because $|\mathcal{S}_m(\mathbb{F}_{q^2})| = |\mathcal{S}_m(\mathbb{F}_q)|$. The field \mathbb{F}_{q^4} is a good choice for increasing length and dimension of the code because, as mentioned above, \mathcal{S}_m is maximal over this field.

In this section, we used Suzuki curves to determine LRC(2)s with recovery sets which are more balanced in size than those constructed in Section 4. However, the quotient curve construction does not naturally provide explicit expressions for the bases of functions. In this case, equations and explicit realizations of their function fields are known for the quotient curves [7]. However, the necessary functions x_1 and x_2 that generate the function fields of the quotient curves are unknown, even for this very well-studied curve. This is a larger issue with the quotient curve construction—knowledge of the top curve and about the existence, genera, and even models of the quotient curves does not give full information about the functions that generate the associated function field extensions, and their degrees as maps to \mathbb{P}^1 . So even given a curve \mathcal{X} with many points and a large automorphism group, it is still difficult to generate useful codes from the quotients, and, as mentioned in Section 2, extending this construction to more than two subgroups of $\text{Aut}(\mathcal{X})$ faces additional obstacles. This motivates us to seek good general examples for LRC(t)s from fiber product constructions. In the next section, we consider a general fiber product construction that is both explicit and gives rise to balanced recovery sets. Moreover, it naturally leads to even more recovery sets for each position.

6. LRCs WITH MULTIPLE RECOVERY SETS FROM FIBER PRODUCTS OF ARTIN-SCHREIER CURVES

In Van der Geer and Van der Vlugt [9], the authors develop several constructions of fiber products of Artin-Schreier curves with many points. In particular, they construct maximal curves via the fiber product of several maximal curves, in both characteristic 2 and odd characteristic. These constructions are a very natural source of curves for locally recoverable codes with many recovery sets.

The simplest of these constructions is given in [9, Section 3, Method I]. Let p be prime, h an even natural number, and $q = p^h$ ([9] also gives a similar construction for h odd). Let $A = \{a \in \mathbb{F}_q : a^{p^{\frac{h}{2}}} + a = 0\}$. As the kernel of the \mathbb{F}_p -linear trace map $\mathbb{F}_q/\mathbb{F}_{\sqrt{q}}$, A is an $\frac{h}{2}$ -dimensional \mathbb{F}_p vector space. Let $\{a_1, a_2, \dots, a_{\frac{h}{2}}\}$ generate A over \mathbb{F}_p . Then the curves

$$C_{a_i} : y^p + y = a_i x^{\sqrt{q}+1}$$

each have genus $\frac{(p-1)\sqrt{q}}{2}$ and have $pq + 1$ points over \mathbb{F}_q , with one point $\infty_{C_{a_i}}$ at infinity.




FIGURE 4. Curves for locally recoverable codes with availability t

Let t be an integer with $1 \leq t \leq \frac{h}{2}$. Then consider the natural map $h_i : C_{a_i} \rightarrow \mathbb{P}_x^1$ given by projection onto the x coordinate, where ∞_x represents the point at infinity on the projective line \mathbb{P}_x^1 and $\infty_{C_{a_i}} \mapsto \infty_x$. These are all degree- p Artin-Schreier covers of \mathbb{P}_x^1 , fully ramified above ∞_x .

Define \mathcal{C} to be the fiber product of these curves C_{a_i} over \mathbb{P}_x^1 , i.e.

$$\mathcal{C} = C_{a_1} \times_{\mathbb{P}_x^1} C_{a_2} \times_{\mathbb{P}_x^1} \cdots \times_{\mathbb{P}_x^1} C_{a_t}.$$

The corresponding maps $g_i : \mathcal{C} \rightarrow C_{a_i}$ are degree p^{t-1} , ramified only above $\infty_{C_{a_i}}$. Let $\infty_{\mathcal{C}}$ be the single point above $\infty_{C_{a_i}}$ on \mathcal{C} .

As shown in [9, Theorem 3.1], the curve \mathcal{C} has genus $\frac{(p^t-1)\sqrt{q}}{2}$ and $|\mathcal{C}(\mathbb{F}_q)| = p^t q + 1$, making \mathcal{C} maximal over \mathbb{F}_q .

Theorem 9. *The code $C((q - tp^t)\infty_x, B)$ constructed from the fiber product \mathcal{C} of the Artin-Schreier curves above is a locally recoverable (n, k, d) code over \mathbb{F}_q with availability t and locality $(p - 1, p - 1, \dots, p - 1)$ where*

$$\begin{aligned} n &= p^t q, \\ k &= (q - tp^t + 1)(p - 1)^t, \text{ and} \\ d &\geq 2tp^t. \end{aligned}$$

Proof. The construction in Theorem 1 gives rise to a code of length $n = p^t q$, where $B = \mathcal{C}(\mathbb{F}_q) \setminus \{\infty_{\mathcal{C}}\}$. The set $S = g(B)$ has size $s = q$. To achieve positive minimum distance, we then take $l = q - tp$, where $g_{\mathcal{C}} = 0$ so $m = q - tp + 1$, giving a code of dimension $k \geq (q - tp + 1)(p - 1)^t$. Each coordinate has t disjoint recovery sets of size $(p - 1)$, given by the fibers $\tilde{g}_i^{-1}(P)$ for $P \in \tilde{\mathcal{C}}_{a_i}$ where $1 \leq i \leq t$. The minimum distance is bounded by $d \geq n - p^t(q - tp) - t(p - 2)p^t = 2tp^t$. \square

Explicitly, we can write $B = \{(x, y_1, y_2, \dots, y_t) \in \mathbb{F}_q^{t+1} : y_i^p + y_i = a_i x^{\sqrt{q}+1}\}$. The functions $g_i : \mathcal{C} \rightarrow C_{a_i}$ are given by $g_i(x, y_1, y_2, \dots, y_t) = (x, y_i)$ and the functions $\tilde{g}_i : \mathcal{C} \rightarrow \tilde{\mathcal{C}}_{a_i}$ are given by $\tilde{g}_i(x, y_1, y_2, \dots, y_t) = (x, y_1, y_2, \dots, y_{i-1}, y_{i+1}, \dots, y_t)$. The divisor D in question is $D = (q - tp^t)\infty_x$, meaning $\mathcal{L}(D)$ is the set of polynomials in x over \mathbb{F}_q of degree $\leq q - tp^t$. Then functions leading to codewords are

$$V = \text{Span}\{x^j y_1^{e_1} y_2^{e_2} \dots y_t^{e_t} : 0 \leq j \leq q - tp^t, 0 \leq e_i \leq p - 2\}.$$

Let $P = (\alpha, \beta_1, \beta_2, \dots, \beta_t) \in B$. Then $B^{(i)}(P) = \{(\alpha, y_1, y_2, \dots, y_t) \in B : y_j = \beta_j \forall j \neq i\}$. We then have $|B^{(i)}(P)| = p$, where, on $B^{(i)}(P)$, any function in V varies as a polynomial

in y_i of degree $\leq (p-2)$, which can therefore be interpolated by knowing its values on any $p-1$ points.

Remark Though it does not seem to have been remarked before, it is straightforward to see that when $t = \frac{h}{2}$, then the curve \mathcal{C} has genus $\frac{(\sqrt{q}-1)\sqrt{q}}{2}$ and is maximal over \mathbb{F}_q . Since the Hermitian curve $\mathcal{H}_{\sqrt{q}}$ is the only curve of this genus maximal over \mathbb{F}_q [10], we have that $\mathcal{C} \cong \mathcal{H}_{\sqrt{q}}$. Therefore the construction in Theorem 9 is yet another example of a code on the Hermitian curve with interesting properties.

Example We construct a locally recoverable code with three recovery sets of size two for each position. Consider the field

$$\mathbb{F}_{3^6} \cong \mathbb{F}_3[x]/\langle x^6 + 2x^4 + x^2 + 2x + 2 \rangle.$$

As the construction indicates, we need roots of the polynomial $x^{27} + x$ that generate a 3-dimensional vector space over \mathbb{F}_3 . We choose the following roots:

$$\begin{aligned} 2a^3 + a + 1 &= a^{350} := a_1 \\ a^4 + a^2 + 1 &= a^{210} := a_2 \\ 2a^5 + a^3 + a^2 + 1 &= a^{490} := a_3. \end{aligned}$$

Now the components of the fiber product are the curves $C_{a_1}, C_{a_2}, C_{a_3}$, given below:

$$\begin{aligned} y^3 + y &= a^{350} x^{28} \\ y^3 + y &= a^{210} x^{28} \\ y^3 + y &= a^{490} x^{28}. \end{aligned}$$

The code has length $n = 19683$, dimension $k = 5192$, and minimum distance $d \geq 162$.

7. CONCLUSION

In this paper, we detailed a construction of locally recoverable codes from fiber products of algebraic curves, building on the work of Barg et. al. [3]. This construction results in different bounds for minimum distance in the case $t = 2$ and allows for $t > 2$ recovery sets. This gives rise to several new families of locally recoverable codes, specifically those from generalized Giulietti and Korchmáros (GK) curves, Suzuki curves, and the maximal curves of Van der Geer and Van der Vlugt, including an LRC(t) for the Hermitian curve \mathcal{H}_{p^t} . The code construction from generalized GK curves is explicit, but the recovery sets have vastly different cardinalities (an advantage or disadvantage, depending on the perspective taken). In contrast, the codes from the Suzuki curves provide balanced recovery sets but the codes themselves are not easily explicitly constructed. The construction using the curves of Van der Geer and Van der Vlugt provides both features, meaning explicit code construction and balanced recovery sets, of which arbitrarily many are available.

Acknowledgements We would like to thank Bjorn Poonen and Rachel Pries for helpful comments. In addition, we would like to acknowledge the hospitality of IPAM and the organizers of its Algebraic Geometry for Coding Theory and Cryptography Workshop: Everett Howe, Kristin Lauter, and Judy Walker.

REFERENCES

- [1] M. Abdón, J. Bezerra, and Luciane Quoos, Further examples of maximal curves, *Journal of Pure and Applied Algebra*, **213** (2009), 1192–1196.
- [2] E. Ballico and A. Ravagnani, Embedding Suzuki curves in \mathbb{P}^4 , *Journal of Commutative Algebra*, **7** (2015), 149–166.
- [3] A. Barg, I. Tamo, and S. Vlăduț, Locally recoverable codes on algebraic curves, *Proceedings of the IEEE Int. Symp. Info. Theory*, (2015), 1252–1256.
- [4] A. Eid and I. Duursma, Smooth embeddings for the Suzuki and Ree curves, *Proceedings of the conference on Arithmetic, Geometry and Coding Theory (AGCT 2013)*, Contemporary Mathematics Series (AMS), **637**, 251–291.
- [5] A. Garcia, C. G'uner, and H. Stichtenoth, A generalization of the Giulietti-Korchmáros maximal curve, *Advances in Geometry*, **10** (2010), 427–434.
- [6] M. Giulietti and G. Korchmáros, A new family of maximal curves over a finite field, *Math. Ann.* **343** (2009), no. 1, 229–245.
- [7] M. Giulietti, G. Korchmáros, and F. Torres, Quotient curves of the Suzuki curve, *Acta Arithmetica*, **122** (2006), 245–274.
- [8] Deligne-Lusztig varieties and group codes, in *Coding Theory and Algebraic Geometry*, **1518** (2006), Lecture Notes in Mathematics, 63–81.
- [9] How to construct curves over finite fields with many points, in *Arithmetic geometry (Cortona, 1994)*, Symposia Mathematica XXXVII (1997), Cambridge: Cambridge University Press, 169–189.
- [10] H.-G. R'uck and H. Stichtenoth, A characterization of Hermitian function fields over finite fields, *Journal fr die reine und angewandte Mathematik*, **457** (1994), 185–188.