

THE GEOMETRY OF COLLUSION LEAKAGE IN INNER-PRODUCT FUNCTIONAL ENCRYPTION

FERDINANDO ZULLO

ABSTRACT. Inner-product functional encryption (IPFE) allows the holder of a functional key for a vector y to learn $\langle x, y \rangle$ from an encryption of x , and nothing more. Keys, however, accumulate: a coalition holding sufficiently many independent keys recovers the plaintext. We model each institution by the span $W_i \leq \mathbb{F}_q^n$ of its authorized key vectors and show that a coalition I determines the plaintext exactly modulo $(\sum_{i \in I} W_i)^\perp$; for a uniform plaintext it learns exactly $\dim \sum_{i \in I} W_i$ field symbols. The resulting leakage function is a representable and entropic polymatroid. We then study authorization spaces of dimension r whose pairwise intersections have dimension at most h . We prove that some coalition of size $c \geq 2$ always learns at least $\min\{n, 2r - h + c - 2\}$ dimensions, and we give a Vandermonde-type construction in which every coalition of size c learns exactly this amount, for all c simultaneously. Hence the largest achievable reconstruction threshold is $n - 2r + h + 2$, which gives an exact tradeoff between functional diversity and collusion resistance. The optimal configuration is a sunflower, an equidistant subspace code and a generalized arc; its orthogonal complements form a generalized dual arc, and for $r = 1$ it reduces to Reed–Solomon secret sharing.

1. INTRODUCTION

Functional encryption provides a cryptographic mechanism in which a secret key is associated with a function rather than with the entire plaintext. Given an encryption of a message x and a functional key for a function f , the holder of the key learns $f(x)$ but, ideally, nothing else about x .

A particularly important and practically motivated special case is inner-product functional encryption (IPFE). In this setting, messages are vectors x and functional keys are associated with vectors y ; decryption returns

$$\langle x, y \rangle.$$

Inner-product functional encryption was introduced in [1], and a substantial body of work has since developed more efficient constructions, stronger security guarantees, multi-input variants, and mechanisms for fine-grained access control; for a survey of functional encryption, including inner-product schemes, see [19].

A fundamental difficulty arises when functional keys accumulate. Suppose that $x \in \mathbb{F}_q^n$ and that a user obtains keys corresponding to n linearly independent vectors $y_1, \dots, y_n$. The values

$$\langle x, y_1 \rangle, \dots, \langle x, y_n \rangle$$

determine x uniquely. Thus, even if each individual key reveals only one authorized linear statistic, a sufficiently rich collection of legitimate functional keys may reveal the entire plaintext. This limitation has been explicitly recognized in the IPFE literature and is one of the motivations for introducing additional access-control mechanisms; see, for example, [2, 7]. Alternative approaches to IPFE have also been investigated. In particular, La Scala and Meneghetti [18] studied an IPFE protocol based on the UOV multivariate framework, analyzing its efficiency and security and proposing variants aimed at reducing the decryption cost.

2020 *Mathematics Subject Classification.* Primary 94A60; Secondary 51E20, 05B35, 94A62, 94B27.

Key words and phrases. Inner-product functional encryption, collusion, subspace arrangements, polymatroids, sunflowers, generalized arcs.

In this paper we take a complementary point of view. We do not propose a new cryptographic construction. Instead, we study the unavoidable information revealed by collections of authorized inner products and ask how such authorizations should be distributed among several institutions in order to control leakage under collusion.

The central object is an *authorization space*. If institution i is entitled to evaluate inner products corresponding to the vectors $y_{i,1}, \dots, y_{i,s_i}$, we associate with it the subspace

$$W_i = \langle y_{i,1}, \dots, y_{i,s_i} \rangle \leq \mathbb{F}_q^n.$$

Even if functional keys themselves cannot be algebraically combined, their outputs can: knowledge of the inner products associated with a basis of W_i determines the inner product associated with every vector of W_i . Hence W_i , rather than the particular list of issued keys, is the natural object measuring the information made available to institution i .

For a coalition I of institutions, the effective authorization space is

$$W_I = \sum_{i \in I} W_i.$$

Our first observation is that two plaintexts x, x' are indistinguishable from the point of view of the authorized linear computations of the coalition precisely when $x - x' \in W_I^\perp$. Thus a coalition determines x exactly up to the affine space $x + W_I^\perp$, and the number of plaintexts compatible with its observations is $q^{n - \dim W_I}$, whatever the distribution of the plaintext.

This leads naturally to the leakage function

$$\rho(I) = \dim W_I.$$

The function ρ is normalized, nondecreasing, and submodular, and hence defines a representable polymatroid on the set of institutions. For a uniform plaintext it is, up to the factor $\log_2 q$, the joint entropy of the institutions' views, so it is also an entropic polymatroid in the information-theoretic sense. Since polymatroids are less common in the functional-encryption literature, we recall the necessary background in Section 2, together with a dictionary between the polymatroid axioms and statements about collusion. In this way, the problem of controlling cumulative leakage in IPFE becomes a problem about arrangements of subspaces.

The structure of coalition authorization spaces is closely related to that of *linear secret sharing schemes* and *span programs* [6, 14, 4]. In a span program over $\mathbb{F}_q$, each party is assigned a subspace of an ambient vector space, and an authorized coalition is one whose subspaces span a prescribed target. When the target is the whole space and the secret is a vector in $\mathbb{F}_q^n$ rather than a single field element, the model coincides exactly with ours. Our contribution is to impose uniform dimension and pairwise intersection constraints and to obtain a sharp extremal result with an explicit construction, rather than to study the access structure from a purely combinatorial or information-theoretic perspective.

We focus on the following uniform model. We assume

$$\dim W_i = r$$

for every institution and impose a diversity constraint

$$\dim(W_i \cap W_j) \leq h \quad (i \neq j).$$

The parameter h measures how similar the permissions of two institutions may be. Small h forces institutions to possess substantially different functional capabilities, whereas large h permits greater overlap.

Since larger overlaps can only help privacy, one may ask why a designer would impose such a constraint at all. In practice, authorizations are dictated by the roles of the institutions and not by privacy alone. Two agencies with distinct mandates, say a tax authority and a health authority querying the same encrypted registry, must each be able to compute statistics that the other cannot, and separation-of-duties policies require distinct roles to differ substantially. The condition $\dim(W_i \cap W_j) \leq h$, equivalently $\dim W_i - \dim(W_i \cap W_j) \geq r - h$, says precisely that every institution can evaluate at least $r - h$ independent statistics that are not available

to any other single institution. The question is then how much collusion resistance must be sacrificed to guarantee this functional diversity.

For a coalition size c , define

$$L_c(\mathcal{W}) = \max_{\substack{I \subseteq [m] \\ |I|=c}} \dim \left(\sum_{i \in I} W_i \right).$$

We prove the universal lower bound

$$L_c(\mathcal{W}) \geq \min\{n, 2r - h + c - 2\}, \quad c \geq 2.$$

The bound has a simple interpretation. Two different institutions already reveal at least $2r - h$ independent linear dimensions. After that, as long as the entire plaintext space has not been reached, there is always an additional institution contributing at least one new dimension.

We then show that the bound is sharp. Under natural conditions on the parameters, we construct authorization spaces for which

$$\dim \left(\sum_{i \in I} W_i \right) = \min\{n, 2r - h + |I| - 2\}$$

for every coalition with at least two members. The construction uses graphs of a family of linear maps whose parameters lie on a Vandermonde curve.

The resulting exact reconstruction threshold is

$$\kappa_{\max} = n - 2r + h + 2.$$

Equivalently, no authorization architecture satisfying the above constraints can guarantee privacy against $n - 2r + h + 2$ colluding institutions, while our construction guarantees that every smaller coalition still leaves some plaintext uncertainty.

The geometric formulation connects the problem with several established objects in finite geometry. The authorization spaces form an equidistant constant-dimension code when equality $\dim(W_i \cap W_j) = h$ holds, and in our construction they form a sunflower with center of dimension h ; see [17, 10]. In fact, the optimal configuration is a generalized arc in the sense of [16], and its orthogonal complements form a generalized dual arc; for constructions of such objects see [13, 16]. Subspace arrangements also naturally produce representable polymatroids; see, for example, [15].

Our aim is therefore not to introduce a new finite-geometric object, but to identify an exact geometric model for collusion leakage in IPFE and to determine the optimal privacy profile in a basic but natural authorization model.

The paper is organized as follows. Section 2 introduces authorization spaces, computes the exact leakage of a coalition, and develops the polymatroid point of view. Section 3 proves the universal lower bound, and Section 4 gives the optimal construction. Section 5 translates these results into privacy profiles and reconstruction thresholds, Section 6 interprets the optimal configuration in finite geometry, and Section 7 concludes with open problems.

2. INNER-PRODUCT AUTHORIZATIONS AND UNAVOIDABLE LEAKAGE

Throughout the paper, q is a prime power and $V = \mathbb{F}_q^n$. We fix a nondegenerate bilinear form

$$\langle \cdot, \cdot \rangle : V \times V \longrightarrow \mathbb{F}_q,$$

and for a subspace $W \leq V$ we put

$$W^\perp = \{x \in V : \langle x, y \rangle = 0 \text{ for all } y \in W\}.$$

Since the form need not be symmetric, we always use this (left) orthogonal. For the standard inner product all statements below have their usual coordinate interpretation.

2.1. Authorization spaces. Consider an IPFE system in which a ciphertext encrypts a vector $x \in V$, and m institutions, indexed by $[m] = \{1, \dots, m\}$. Suppose that institution i possesses functional keys corresponding to vectors

$$y_{i,1}, \dots, y_{i,s_i}.$$

The institution can recover the values

$$\langle x, y_{i,1} \rangle, \dots, \langle x, y_{i,s_i} \rangle.$$

Definition 2.1. *The authorization space of institution i is*

$$W_i = \langle y_{i,1}, \dots, y_{i,s_i} \rangle_{\mathbb{F}_q},$$

and we write $\mathcal{W} = (W_1, \dots, W_m)$ for the family of authorization spaces. For a coalition $I \subseteq [m]$, its coalition authorization space is

$$W_I = \sum_{i \in I} W_i.$$

The use of a span is forced by linearity. If $y = \sum_j \lambda_j y_{i,j}$, then

$$\langle x, y \rangle = \sum_j \lambda_j \langle x, y_{i,j} \rangle.$$

Thus the evaluations associated with any basis of W_i reveal the evaluation associated with every vector of W_i .

Proposition 2.2 (Exact residual uncertainty). *Let $I \subseteq [m]$ and let $x, x' \in V$. The two plaintexts give the same inner-product evaluations for every authorized vector of the coalition I if and only if $x - x' \in W_I^\perp$. Consequently, after learning all authorized evaluations, the coalition determines the plaintext precisely up to the coset $x + W_I^\perp$.*

Proof. The coalition obtains the same evaluations on x and x' precisely when

$$\langle x, y \rangle = \langle x', y \rangle \quad \text{for every } y \in W_I.$$

Equivalently,

$$\langle x - x', y \rangle = 0 \quad \text{for every } y \in W_I,$$

which is exactly the condition

$$x - x' \in W_I^\perp.$$

□

We now quantify the residual uncertainty described in Proposition 2.2. Let $d = \dim W_I$ and fix a basis $y_1, \dots, y_d$ of W_I . The *observation map* of the coalition is

$$\phi_I : V \longrightarrow \mathbb{F}_q^d, \quad x \longmapsto (\langle x, y_1 \rangle, \dots, \langle x, y_d \rangle).$$

By linearity of the form, $\phi_I(x)$ determines $\langle x, y \rangle$ for every $y \in W_I$, and conversely. Hence, up to a bijective relabelling that does not affect any entropy, the complete view of the coalition is the random variable

$$Z_I = \phi_I(X).$$

A different choice of basis changes ϕ_I by an invertible linear map of $\mathbb{F}_q^d$, so nothing below depends on that choice.

Lemma 2.3. *The map ϕ_I is surjective and $\ker \phi_I = W_I^\perp$. Consequently, every fiber $\phi_I^{-1}(z)$, $z \in \mathbb{F}_q^d$, is a coset of $W_I^\perp$ and has exactly q^{n-d} elements.*

Proof. A vector x lies in $\ker \phi_I$ if and only if $\langle x, y_j \rangle = 0$ for all j , that is, if and only if $x \in W_I^\perp$. Since the form is nondegenerate, the map $x \mapsto \langle x, \cdot \rangle$ is an isomorphism $V \rightarrow V^*$, and $W_I^\perp$ is the preimage of the annihilator of W_I , which has dimension $n - d$. By rank-nullity, $\dim \phi_I(V) = n - (n - d) = d$, so ϕ_I is surjective. Each fiber is then a coset of the kernel and has cardinality $|W_I^\perp| = q^{n-d}$. □

We use two uncertainty measures, both in bits. For random variables X and Z on finite sets, $H(X | Z)$ is the conditional Shannon entropy, and

$$\tilde{H}_\infty(X | Z) = -\log_2 \mathbb{E}_{z \leftarrow Z} \left[\max_x \Pr[X = x | Z = z] \right]$$

is the average conditional min-entropy of [9]. The latter is the standard cryptographic measure: $2^{-\tilde{H}_\infty(X|Z)}$ is the probability that an adversary who sees Z guesses X correctly in a single attempt.

Corollary 2.4. *Let X be uniformly distributed on V , let $Z_I = \phi_I(X)$, and put $d = \dim W_I$. Then:*

- (i) Z_I is uniformly distributed on $\mathbb{F}_q^d$;
- (ii) for every $z \in \mathbb{F}_q^d$, the conditional distribution of X given $Z_I = z$ is uniform on the coset $\phi_I^{-1}(z)$;
- (iii) $H(X | Z_I) = \tilde{H}_\infty(X | Z_I) = (n - d) \log_2 q$;
- (iv) $I(X; Z_I) = d \log_2 q$.

In particular, the best single-guess success probability of the coalition is exactly $q^{-(n-d)}$.

Proof. By Lemma 2.3, every $z \in \mathbb{F}_q^d$ has exactly q^{n-d} preimages, so

$$\Pr[Z_I = z] = \frac{q^{n-d}}{q^n} = q^{-d},$$

which proves (i). For $x \in \phi_I^{-1}(z)$,

$$\Pr[X = x | Z_I = z] = \frac{\Pr[X = x]}{\Pr[Z_I = z]} = \frac{q^{-n}}{q^{-d}} = q^{-(n-d)},$$

and the conditional probability vanishes outside $\phi_I^{-1}(z)$; this is (ii). A uniform distribution on q^{n-d} points has Shannon entropy $(n - d) \log_2 q$ and maximal point probability $q^{-(n-d)}$. Both values are independent of z , so averaging over z gives (iii), together with the guessing probability. Finally,

$$I(X; Z_I) = H(X) - H(X | Z_I) = n \log_2 q - (n - d) \log_2 q = d \log_2 q,$$

which is (iv). □

Statement (iv) is the precise sense in which $\dim W_I$ is the *leakage* of the coalition: it learns exactly $\dim W_I$ symbols of $\mathbb{F}_q$ worth of information about the plaintext. Note that only the dimension of W_I matters, not which particular vectors were issued as keys.

Example 2.5. Let $n = 3$, $q = 5$, and use the standard inner product. If $W_I = \langle e_1, e_2 \rangle$, the coalition learns x_1 and x_2 , and each observation is consistent with the 5 choices of x_3 . If instead $W_I = \langle e_1 + e_2, e_2 + e_3 \rangle$, the coalition learns $x_1 + x_2$ and $x_2 + x_3$, and each observation is consistent with the 5 points of a coset of $W_I^\perp = \langle (1, -1, 1) \rangle$. In both cases $H(X | Z_I) = \log_2 5 = (3 - 2) \log_2 5$.

Remark 2.6 (Nonuniform plaintexts). *Uniformity of X is used only in Corollary 2.4; the combinatorial statement of Proposition 2.2 holds for every source. For an arbitrary distribution of X , since Z_I is a function of X taking at most q^d values,*

$$I(X; Z_I) = H(Z_I) \leq d \log_2 q,$$

with equality if and only if $\phi_I(X)$ is uniform on $\mathbb{F}_q^d$. Similarly, by [9, Lemma 2.2],

$$\tilde{H}_\infty(X | Z_I) \geq H_\infty(X) - d \log_2 q.$$

Thus $\dim W_I \cdot \log_2 q$ is always an upper bound on the information leaked to coalition I , and it is attained for uniform plaintexts. For structured sources the actual leakage may be smaller, while the residual uncertainty may also be smaller than $(n - d) \log_2 q$, because prior knowledge about X can combine with the observations.

This motivates the following definition.

Definition 2.7. *The linear leakage of a coalition I is*

$$\rho_{\mathcal{W}}(I) = \dim \left(\sum_{i \in I} W_i \right).$$

Its residual privacy dimension is

$$\pi_{\mathcal{W}}(I) = n - \rho_{\mathcal{W}}(I).$$

2.2. The authorization polymatroid. Polymatroids are the natural language for functions that measure “how much” a subset of a finite set of objects reveals. They are a standard tool in information theory and in secret sharing, where they are used, for instance, to derive lower bounds on share sizes; see [8, 4]. Since they are less common in the functional-encryption literature, we recall the basic notions; for background see [11, 21].

Definition 2.8. *Let Ω be a finite set. A function $f : 2^\Omega \rightarrow \mathbb{R}_{\geq 0}$ is a polymatroid (rank function) on Ω if*

- (P1) $f(\emptyset) = 0$ (normalization);
- (P2) $f(I) \leq f(J)$ whenever $I \subseteq J$ (monotonicity);
- (P3) $f(I) + f(J) \geq f(I \cup J) + f(I \cap J)$ for all $I, J \subseteq \Omega$ (submodularity).

It is an integer polymatroid if it takes integer values, and a matroid if, in addition, $f(\{e\}) \leq 1$ for every $e \in \Omega$.

The axiom that carries the content is submodularity. It is equivalent to the following *diminishing-returns* property: for all $I \subseteq J \subseteq \Omega$ and $e \in \Omega \setminus J$,

$$(1) \quad f(I \cup \{e\}) - f(I) \geq f(J \cup \{e\}) - f(J).$$

In words, the gain from adding e can only shrink as the set grows. (That (P3) implies (1) follows by applying (P3) to the sets $I \cup \{e\}$ and J , whose union is $J \cup \{e\}$ and whose intersection is I ; the converse is an easy induction, see [21].)

Two classes of polymatroids are relevant here.

Definition 2.9. *A polymatroid f on Ω is $\mathbb{F}$ -representable (or linear) if there are a finite-dimensional $\mathbb{F}$ -vector space U and subspaces $U_e \leq U$, $e \in \Omega$, such that*

$$f(I) = \dim \left(\sum_{e \in I} U_e \right) \quad \text{for every } I \subseteq \Omega.$$

It is entropic if there are jointly distributed random variables $(Z_e)_{e \in \Omega}$ on a finite probability space such that $f(I) = H((Z_e)_{e \in I})$ for every I .

When every U_e is a line spanned by a vector u_e , one recovers the familiar vector matroid of the configuration $\{u_e\}$, whose rank function is $I \mapsto \text{rk}\{u_e : e \in I\}$. Representable polymatroids are the generalization in which each element of Ω is a subspace of arbitrary dimension rather than a single vector. The fact that the joint entropy of any family of random variables is a polymatroid is due to Fujishige [11]: axiom (P3) for $f(I) = H(Z_I)$ is exactly the nonnegativity of the conditional mutual information

$$(2) \quad f(I) + f(J) - f(I \cup J) - f(I \cap J) = I(Z_{I \setminus J}; Z_{J \setminus I} \mid Z_{I \cap J}) \geq 0,$$

where $Z_K = (Z_e)_{e \in K}$. Not every polymatroid is representable. Representable ones satisfy additional linear rank inequalities, the best known being the Ingleton inequality [12], which can fail for entropic polymatroids; see [15] for further inequalities of this kind.

Proposition 2.10. *The linear leakage*

$$\rho_{\mathcal{W}} : 2^{[m]} \longrightarrow \mathbb{Z}_{\geq 0}, \quad I \longmapsto \dim \left(\sum_{i \in I} W_i \right)$$

is an $\mathbb{F}_q$ -representable integer polymatroid on the set $[m]$ of institutions, and $\rho_{\mathcal{W}}([m]) = \dim(W_1 + \dots + W_m)$.

Proof. Representability holds by definition, taking $U = V$ and $U_i = W_i$; it remains to check (P1)–(P3). Clearly $\rho_{\mathcal{W}}(\emptyset) = 0$, and monotonicity follows from inclusion of sums.

Let

$$A = \sum_{i \in I} W_i, \quad B = \sum_{j \in J} W_j.$$

Then

$$A + B = \sum_{k \in I \cup J} W_k,$$

whereas

$$\sum_{k \in I \cap J} W_k \leq A \cap B.$$

Hence

$$\begin{aligned} \rho_{\mathcal{W}}(I) + \rho_{\mathcal{W}}(J) &= \dim A + \dim B \\ &= \dim(A + B) + \dim(A \cap B) \\ &\geq \rho_{\mathcal{W}}(I \cup J) + \rho_{\mathcal{W}}(I \cap J). \end{aligned}$$

□

Each polymatroid axiom has a direct meaning for coalitions of institutions.

Normalization and monotonicity. The empty coalition learns nothing, and admitting new members never decreases what a coalition learns.

Marginal leakage and diminishing returns. For a coalition I and an institution $e \notin I$, the dimension formula gives the *marginal leakage*

$$(3) \quad \Delta_e(I) := \rho_{\mathcal{W}}(I \cup \{e\}) - \rho_{\mathcal{W}}(I) = \dim W_e - \dim(W_e \cap W_I).$$

This is the part of e 's authorization that the coalition cannot already compute by itself. If $I \subseteq J$, then $W_e \cap W_I \leq W_e \cap W_J$, so $\Delta_e(I) \geq \Delta_e(J)$. This is the diminishing-returns form (1) of submodularity, obtained here directly: recruiting a new member is worth the most to a small coalition. In particular, the bound $L_2(\mathcal{W}) \geq 2r - h$ and the inductive step in the proof of Theorem 3.5 below are both statements about marginal leakages.

Merging coalitions. Submodularity can be read as

$$\rho_{\mathcal{W}}(I \cup J) \leq \rho_{\mathcal{W}}(I) + \rho_{\mathcal{W}}(J) - \rho_{\mathcal{W}}(I \cap J).$$

Two coalitions that pool their information learn at most what they learn separately, minus what their common members already contributed to both. In particular, leakage is subadditive: $\rho_{\mathcal{W}}(I \cup J) \leq \rho_{\mathcal{W}}(I) + \rho_{\mathcal{W}}(J)$.

Closure: simulable institutions. The *closure* of a coalition I is

$$\text{cl}(I) = \{e \in [m] : \Delta_e(I) = 0\} = \{e \in [m] : W_e \leq W_I\}.$$

An institution $e \in \text{cl}(I)$ adds nothing to I : every inner product that e is authorized to evaluate can already be computed by I from its own outputs. In this sense, the coalition I can *simulate* every member of its closure.

Reconstruction. When $W_1 + \dots + W_m = V$, the family

$$\Gamma(\mathcal{W}) = \{I \subseteq [m] : \rho_{\mathcal{W}}(I) = n\}$$

of coalitions that recover the plaintext is monotone (closed under supersets). In secret-sharing terminology, it is the access structure of the configuration. The reconstruction threshold introduced in Section 3 is the minimum size of a member of $\Gamma(\mathcal{W})$. Unlike perfect secret sharing, coalitions outside $\Gamma(\mathcal{W})$ are not ignorant: they learn $\rho_{\mathcal{W}}(I)$ dimensions. The present setting is therefore closer to ramp schemes [5], in which partial information is revealed gradually as coalitions grow, and the whole function $\rho_{\mathcal{W}}$, not only $\Gamma(\mathcal{W})$, is the object of interest.

The polymatroid $\rho_{\mathcal{W}}$ coincides, up to the factor $\log_2 q$, with the entropy function of the institutions' views.

Proposition 2.11. *Let X be uniform on V , and let $Z_i = \phi_{\{i\}}(X)$ be the view of institution i . For every $I \subseteq [m]$, writing $Z'_I = (Z_i)_{i \in I}$ for the joint view of the members of I ,*

$$H(Z'_I) = \rho_{\mathcal{W}}(I) \cdot \log_2 q.$$

Hence $\rho_{\mathcal{W}} \cdot \log_2 q$ is an entropic polymatroid, and submodularity of $\rho_{\mathcal{W}}$ is the instance (2) of the nonnegativity of conditional mutual information.

Proof. The joint view Z'_I consists of the values $\langle X, y \rangle$ for $y \in W_i$, $i \in I$. By linearity these determine $\langle X, y \rangle$ for every $y \in W_I$, that is, they determine $\phi_I(X)$. Conversely, $\phi_I(X)$ determines each Z_i . Thus Z'_I and $\phi_I(X)$ are functions of each other and have the same entropy. By Corollary 2.4(i), $\phi_I(X)$ is uniform on $\mathbb{F}_q^{\rho_{\mathcal{W}}(I)}$, so its entropy is $\rho_{\mathcal{W}}(I) \log_2 q$. $\square$

For a uniform plaintext, therefore, $H(Z'_I) = I(X; Z'_I)$: since the views are functions of X , the entropy of the joint view is exactly the information it carries about the plaintext. Proposition 2.11 makes the leakage function a familiar object for cryptographers. The linear structure of IPFE adds information that a general entropic function does not have. For example, $\rho_{\mathcal{W}}$ always satisfies the Ingleton inequality. This constraint on coalition leakage holds for every IPFE authorization configuration, but it can fail for the entropy functions of general, nonlinear, leakage mechanisms.

Example 2.12. *Let $n = 2$ and consider three institutions with*

$$W_1 = \langle e_1 \rangle, \quad W_2 = \langle e_2 \rangle, \quad W_3 = \langle e_1 + e_2 \rangle.$$

Every single institution has leakage 1, while every pair, and the whole set, has leakage 2: this is the uniform matroid $U_{2,3}$. The marginal leakage of institution 3 is

$$\Delta_3(\emptyset) = 1, \quad \Delta_3(\{1\}) = 1, \quad \Delta_3(\{1, 2\}) = 0,$$

illustrating diminishing returns. Indeed $3 \in \text{cl}(\{1, 2\})$: the coalition $\{1, 2\}$ learns x_1 and x_2 and can compute $\langle x, e_1 + e_2 \rangle = x_1 + x_2$ without holding a key for $e_1 + e_2$. The access structure Γ consists of all coalitions of size at least two, as in a 2-out-of-3 threshold scheme.

3. THE EXTREMAL AUTHORIZATION PROBLEM

We now restrict to authorization spaces of fixed dimension.

Definition 3.1. *Let $0 \leq h < r < n$. An $(n, r, h; m)_q$ authorization configuration is a family*

$$\mathcal{W} = \{W_1, \dots, W_m\} \subseteq \text{Gr}_q(r, n)$$

satisfying

$$W_1 + \dots + W_m = V$$

and

$$\dim(W_i \cap W_j) \leq h \quad (i \neq j).$$

The spanning condition excludes information coordinates that can never be recovered by any collection of institutions.

Remark 3.2. *Every $(n, r, h; m)_q$ authorization configuration satisfies $m \geq 2$ and $n \geq 2r - h$. Indeed, a single r -dimensional subspace with $r < n$ cannot span V , and any two r -dimensional subspaces of V intersect in dimension at least $2r - n$, so that $2r - n \leq h$.*

For $1 \leq c \leq m$, define

$$L_c(\mathcal{W}) = \max_{\substack{I \subseteq [m] \\ |I|=c}} \rho_{\mathcal{W}}(I).$$

Thus $L_c(\mathcal{W})$ is the largest unavoidable leakage of a coalition of c institutions.

For $c = 1$ we have

$$L_1(\mathcal{W}) = r.$$

The first nontrivial case is $c = 2$.

Proposition 3.3. *For every two distinct members of an $(n, r, h; m)_q$ authorization configuration,*

$$\dim(W_i + W_j) \geq 2r - h.$$

Consequently,

$$L_2(\mathcal{W}) \geq 2r - h.$$

Proof. The dimension formula gives

$$\dim(W_i + W_j) = 2r - \dim(W_i \cap W_j) \geq 2r - h.$$

□

Remark 3.4. *The extremal case $h = 0$ represents maximum functional diversity: any two institutions have trivially intersecting authorization spaces, so any pair already leaks $2r$ independent dimensions. In this case the largest achievable reconstruction threshold, determined in Section 5, equals $n - 2r + 2$. Reducing h toward 0 therefore lowers the largest achievable reconstruction threshold, reflecting the unavoidable cost of demanding highly diverse functional permissions.*

The pairwise estimate propagates to arbitrary coalitions, yielding a universal collusion bound.

Theorem 3.5. *Let $\mathcal{W}$ be an $(n, r, h; m)_q$ authorization configuration. Then, for every $2 \leq c \leq m$,*

$$L_c(\mathcal{W}) \geq \min\{n, 2r - h + c - 2\}.$$

Proof. Choose two distinct spaces W_{i_1} and W_{i_2} . Their sum has dimension at least $2r - h$.

Suppose inductively that distinct indices $i_1, \dots, i_s$ have been chosen, where $s < c$, and put

$$S_s = W_{i_1} + \dots + W_{i_s}.$$

If $S_s = V$, then any extension to a coalition of size c has dimension n .

Assume therefore that $S_s \neq V$. Since the entire family spans V , there exists some $j \notin \{i_1, \dots, i_s\}$ such that

$$W_j \not\subseteq S_s.$$

Hence $\dim(S_s + W_j) \geq \dim S_s + 1$.

Starting from dimension at least $2r - h$ after two institutions and repeating this argument gives a coalition of size c whose span has dimension at least

$$\min\{n, 2r - h + c - 2\}.$$

□

Definition 3.6. *The reconstruction threshold of $\mathcal{W}$ is*

$$\kappa(\mathcal{W}) = \min \left\{ |I| : \sum_{i \in I} W_i = V \right\}.$$

Corollary 3.7. *Every $(n, r, h; m)_q$ authorization configuration $\mathcal{W}$ satisfies*

$$\kappa(\mathcal{W}) \leq \min\{m, n - 2r + h + 2\}.$$

Proof. Since the whole family spans V , we have $\kappa(\mathcal{W}) \leq m$. Put $t = n - 2r + h + 2$. If $m \geq t$, then $t \geq 2$ by Remark 3.2, and Theorem 3.5 with $c = t$ gives $L_t(\mathcal{W}) \geq \min\{n, t\} = n$; hence some coalition of size t spans V , and $\kappa(\mathcal{W}) \leq t$. □

The quantity $r - h$ has a natural interpretation. A single institution reveals r dimensions, whereas a second distinct institution forces the leakage to increase by at least $r - h$. Thereafter, the worst-case leakage must increase by at least one dimension for every additional institution until the whole message space is recovered.

4. AN OPTIMAL VANDERMONDE CONSTRUCTION

We now show that the preceding bounds are sharp.

Let $0 \leq h < r < n$ with $n \geq 2r - h$; by Remark 3.2, the last condition is necessary for any authorization configuration to exist. Put

$$k = r - h, \quad t = n - 2r + h + 2,$$

so that $k \geq 1$ and $t \geq 2$. We shall construct m authorization spaces whenever $t \leq m \leq q$.

Choose a direct-sum decomposition

$$V = H \oplus A \oplus C \oplus E$$

with

$$\dim H = h, \quad \dim A = k, \quad \dim C = k - 1, \quad \dim E = t - 1.$$

Note that $h + k + (k - 1) + (t - 1) = n$.

Fix bases $a_1, \dots, a_k$ of A and $c_1, \dots, c_{k-1}$ of C . Identify $E \cong \mathbb{F}_q^{t-1}$. For $\lambda \in \mathbb{F}_q$, define

$$v(\lambda) = (\lambda, \lambda^2, \dots, \lambda^{t-1}) \in E.$$

Define a linear map

$$T_\lambda : A \longrightarrow C \oplus E$$

by

$$T_\lambda(a_j) = \lambda c_j, \quad j = 1, \dots, k - 1,$$

and

$$T_\lambda(a_k) = v(\lambda).$$

Let $U_\lambda = \{(x, T_\lambda(x)) : x \in A\} \leq A \oplus C \oplus E$ be its graph, and put $W_\lambda = H \oplus U_\lambda$.

Clearly

$$\dim W_\lambda = h + k = r.$$

Lemma 4.1. *For distinct $\lambda, \mu \in \mathbb{F}_q$, $W_\lambda \cap W_\mu = H$.*

Proof. It suffices to show $U_\lambda \cap U_\mu = \{0\}$. The intersection of two graphs is naturally isomorphic to $\ker(T_\lambda - T_\mu)$.

For $j = 1, \dots, k - 1$,

$$(T_\lambda - T_\mu)(a_j) = (\lambda - \mu)c_j,$$

while

$$(T_\lambda - T_\mu)(a_k) = v(\lambda) - v(\mu).$$

The first $k - 1$ vectors span C , since $\lambda \neq \mu$, while the last vector is nonzero and lies in the complementary summand E . Since $C \oplus E$ is a direct sum, the images of $a_1, \dots, a_k$ are linearly independent; therefore $T_\lambda - T_\mu$ is injective. $\square$

Thus the construction satisfies the diversity constraint with equality:

$$\dim(W_\lambda \cap W_\mu) = h,$$

for any distinct λ and μ .

The following elementary Vandermonde observation is the key to controlling larger coalitions.

Lemma 4.2. *Let $\lambda_1, \dots, \lambda_c$ be distinct elements of $\mathbb{F}_q$, with $c \leq t$. Then*

$$v(\lambda_2) - v(\lambda_1), \dots, v(\lambda_c) - v(\lambda_1)$$

are linearly independent in E .

Proof. Suppose

$$\sum_{i=2}^c \alpha_i (v(\lambda_i) - v(\lambda_1)) = 0.$$

Set

$$\alpha_1 = -\sum_{i=2}^c \alpha_i.$$

Then

$$\sum_{i=1}^c \alpha_i = 0$$

and

$$\sum_{i=1}^c \alpha_i (1, \lambda_i, \lambda_i^2, \dots, \lambda_i^{t-1}) = 0.$$

Since $c \leq t$ and the λ_i are distinct, the corresponding Vandermonde vectors are linearly independent. Hence every α_i is zero. $\square$

Theorem 4.3. *Let $\lambda_1, \dots, \lambda_m$ be distinct elements of $\mathbb{F}_q$ and assume $t \leq m \leq q$. Then the family*

$$\mathcal{W}^* = \{W_{\lambda_1}, \dots, W_{\lambda_m}\}$$

satisfies, for every $I \subseteq [m]$ with $|I| = c \geq 2$,

$$\dim \left(\sum_{i \in I} W_{\lambda_i} \right) = \min\{n, 2r - h + c - 2\}.$$

In particular, any t members span V .

Proof. Suppose first that $2 \leq c \leq t$, and after relabelling write the coalition as

$$W_{\lambda_1}, \dots, W_{\lambda_c}.$$

We have

$$U_{\lambda_1} + \dots + U_{\lambda_c} = U_{\lambda_1} + \left(\{0\} \oplus \sum_{i=2}^c \text{Im}(T_{\lambda_i} - T_{\lambda_1}) \right).$$

Since a graph intersects the vertical space trivially,

$$\dim(U_{\lambda_1} + \dots + U_{\lambda_c}) = k + \dim \left(\sum_{i=2}^c \text{Im}(T_{\lambda_i} - T_{\lambda_1}) \right).$$

For $i \geq 2$,

$$\text{Im}(T_{\lambda_i} - T_{\lambda_1}) = C + \langle v(\lambda_i) - v(\lambda_1) \rangle.$$

Therefore

$$\sum_{i=2}^c \text{Im}(T_{\lambda_i} - T_{\lambda_1}) = C \oplus \langle v(\lambda_2) - v(\lambda_1), \dots, v(\lambda_c) - v(\lambda_1) \rangle.$$

By Lemma 4.2, the second summand has dimension $c - 1$. Hence

$$\dim(U_{\lambda_1} + \dots + U_{\lambda_c}) = k + (k - 1) + (c - 1) = 2k + c - 2.$$

Adding the common subspace H , we obtain

$$\begin{aligned} \dim(W_{\lambda_1} + \dots + W_{\lambda_c}) &= h + 2k + c - 2 \\ &= h + 2(r - h) + c - 2 \\ &= 2r - h + c - 2. \end{aligned}$$

For $c = t$, the latter quantity equals $2r - h + t - 2 = n$. Hence every t members span V . If $c > t$, any coalition of size c contains a subcoalition of size t , so its span is again V . $\square$

Theorem 4.4. Assume $0 \leq h < r < n$, $n \geq 2r - h$, and put $t = n - 2r + h + 2$. If $t \leq m \leq q$, then among all $(n, r, h; m)_q$ authorization configurations,

$$\min_{\mathcal{W}} L_c(\mathcal{W}) = \min\{n, 2r - h + c - 2\}$$

for every $2 \leq c \leq m$. Moreover, the configuration $\mathcal{W}^*$ of Theorem 4.3 satisfies $L_c(\mathcal{W}^*) = \min_{\mathcal{W}} L_c(\mathcal{W})$ for all $2 \leq c \leq m$ simultaneously.

Proof. Since $m \leq q$, we may choose distinct elements $\lambda_1, \dots, \lambda_m \in \mathbb{F}_q$, and we let

$$\mathcal{W}^* = \{W_{\lambda_1}, \dots, W_{\lambda_m}\}$$

be the family of Theorem 4.3. Throughout, put

$$g(c) = \min\{n, 2r - h + c - 2\}.$$

Step 1: $\mathcal{W}^*$ is an $(n, r, h; m)_q$ authorization configuration. Each W_{λ_i} has dimension $h + k = r$, so $\mathcal{W}^* \subseteq \text{Gr}_q(r, n)$. By Lemma 4.1, $\dim(W_{\lambda_i} \cap W_{\lambda_j}) = \dim H = h$ for $i \neq j$; in particular, since $h < r$, the m spaces are pairwise distinct. Finally, since $m \geq t$, Theorem 4.3 shows that any t members of $\mathcal{W}^*$ span V , and a fortiori $W_{\lambda_1} + \dots + W_{\lambda_m} = V$.

Step 2: the value of the minimum. Fix $2 \leq c \leq m$. By Theorem 3.5, every $(n, r, h; m)_q$ authorization configuration $\mathcal{W}$ satisfies $L_c(\mathcal{W}) \geq g(c)$; hence $\min_{\mathcal{W}} L_c(\mathcal{W}) \geq g(c)$. Conversely, by Theorem 4.3, every coalition $I \subseteq [m]$ with $|I| = c$ satisfies $\rho_{\mathcal{W}^*}(I) = g(c)$, and taking the maximum over all such I gives $L_c(\mathcal{W}^*) = g(c)$. Since $\mathcal{W}^*$ is admissible by Step 1, $\min_{\mathcal{W}} L_c(\mathcal{W}) \leq L_c(\mathcal{W}^*) = g(c)$. Therefore $\min_{\mathcal{W}} L_c(\mathcal{W}) = g(c)$.

Step 3: simultaneous optimality. The configuration $\mathcal{W}^*$ was fixed before c was chosen, so it does not depend on c . By Step 2,

$$L_c(\mathcal{W}^*) = g(c) = \min_{\mathcal{W}} L_c(\mathcal{W}) \quad \text{for every } 2 \leq c \leq m.$$

Equivalently, for every $(n, r, h; m)_q$ authorization configuration $\mathcal{W}$,

$$L_c(\mathcal{W}^*) \leq L_c(\mathcal{W}) \quad \text{for all } 1 \leq c \leq m,$$

where the case $c = 1$ holds trivially because $L_1(\mathcal{W}) = r$ for every configuration. Thus the leakage profile $(L_1(\mathcal{W}^*), \dots, L_m(\mathcal{W}^*))$ is componentwise minimal. A priori, the minimum of L_c could be attained by different configurations for different values of c , so that protecting against small coalitions would force a worse profile against large ones, or conversely. The existence of $\mathcal{W}^*$ shows that no such tradeoff between coalition sizes occurs. $\square$

Remark 4.5. In the configuration $\mathcal{W}^*$ of Theorem 4.3, the leakage of a coalition depends only on its size: $\rho_{\mathcal{W}^*}(I) = g(|I|)$, where

$$g(0) = 0, \quad g(1) = r, \quad g(c) = \min\{n, 2r - h + c - 2\} \quad (c \geq 2).$$

Polymatroids with this property are called symmetric. In particular, in $\mathcal{W}^*$ the best and the worst coalitions of a given size leak exactly the same amount.

The shape of g is best seen through the marginal leakage $g(c) - g(c-1)$, that is, the number of new dimensions revealed when a coalition grows from $c-1$ to c members:

c	1	2	$3, \dots, t$	$t+1, \dots, m$
$g(c) - g(c-1)$	r	$r - h$	1	0

(the third column is empty if $t = 2$, the last one if $m = t$). These increments add up to $r + (r - h) + (t - 2) = n$. The leakage is therefore concentrated at the beginning. The first member reveals r dimensions; the second reveals $r - h$ more, which is unavoidable by Proposition 3.3. From then on, each new member reveals exactly one further dimension, the smallest positive amount possible, until the plaintext is determined at $c = t$. Since $0 \leq h < r$, the increments are nonincreasing; for a function that depends only on $|I|$, this is precisely the diminishing-returns property (1).

Example 4.6. Let $r = 1$, $h = 0$ and $n \geq 2$. Then $k = 1$, $t = n$, $H = C = \{0\}$, and $V = A \oplus E$ with $A = \langle a_1 \rangle$. In the coordinates given by a_1 together with the standard basis of $E \cong \mathbb{F}_q^{n-1}$, the space W_λ is the line spanned by

$$u_\lambda = (1, \lambda, \lambda^2, \dots, \lambda^{n-1}),$$

a point of the normal rational curve. With the standard inner product in these coordinates, and writing $x = (x_0, \dots, x_{n-1})$, institution λ learns

$$\langle x, u_\lambda \rangle = x_0 + x_1 \lambda + \dots + x_{n-1} \lambda^{n-1} = f_x(\lambda),$$

the value at λ of the polynomial whose coefficient vector is the plaintext. This is the share generation of Shamir's scheme [22], except that the whole polynomial, and not only its constant term, is secret; that is, it is the Reed–Solomon ramp scheme of [20, 5]. Any $c \leq n$ institutions learn exactly c dimensions, because distinct evaluation points give linearly independent vectors u_λ , and any n of them recover x by Lagrange interpolation. Hence $\rho_{\mathcal{W}^*}(I) = \min\{n, |I|\}$, the rank function of the uniform matroid $U_{n,m}$.

This case also clarifies the content of Theorem 4.4. For $r = 1$ the theorem carries no information: in every spanning family of m distinct lines one can choose c of them spanning a space of dimension $\min\{n, c\}$, so $L_c(\mathcal{W}) = \min\{n, c\}$ and $\kappa(\mathcal{W}) = n$ for every configuration. What distinguishes the Reed–Solomon configuration is that every coalition of size c , not only the worst one, leaks exactly $\min\{n, c\}$; in geometric terms, the points form an arc. For $r \geq 2$, by contrast, the extremal problem is not trivial: if $n \geq 3r$, three members in general position already give $\dim(W_1 + W_2 + W_3) = 3r > g(3)$.

5. OPTIMAL PRIVACY AND RECONSTRUCTION THRESHOLDS

We translate the previous result back into privacy quantities.

Define the worst-case residual privacy dimension against coalitions of size c by

$$P_c(\mathcal{W}) = \min_{\substack{I \subseteq [m] \\ |I|=c}} \left(n - \dim \sum_{i \in I} W_i \right) = n - L_c(\mathcal{W}).$$

Corollary 5.1. Under the assumptions of Theorem 4.4, the largest possible worst-case residual privacy dimension is

$$P_1^{\text{opt}} = n - r,$$

and, for $2 \leq c \leq m$,

$$P_c^{\text{opt}} = \max\{0, n - 2r + h - c + 2\} = \max\{0, t - c\},$$

where $t = n - 2r + h + 2$. All these values are attained by $\mathcal{W}^*$.

Proof. Since $P_c(\mathcal{W}) = n - L_c(\mathcal{W})$, maximizing P_c is the same as minimizing L_c . The case $c = 1$ follows from $L_1(\mathcal{W}) = r$, and the case $c \geq 2$ from Theorem 4.4. $\square$

For uniform plaintexts, Corollary 2.4 shows that the corresponding conditional entropy is

$$P_c^{\text{opt}} \log_2 q.$$

Corollary 5.2. Under the assumptions of Theorem 4.4,

$$\max_{\mathcal{W}} \kappa(\mathcal{W}) = n - 2r + h + 2,$$

and the maximum is attained by $\mathcal{W}^*$.

Proof. Put $t = n - 2r + h + 2$. Since $m \geq t$, Corollary 3.7 gives $\kappa(\mathcal{W}) \leq t$ for every configuration. Conversely, in $\mathcal{W}^*$ every coalition of size $t - 1$ spans a proper subspace of V : its dimension is $r < n$ if $t = 2$, and $g(t - 1) = n - 1$ if $t \geq 3$, by Theorem 4.3. Hence $\kappa(\mathcal{W}^*) = t$. $\square$

Thus the tradeoff between diversity and collusion resistance is exact. Reducing the maximum allowed intersection h forces different institutions to possess more distinct information, but necessarily lowers the largest achievable reconstruction threshold.

Example 5.3. *Consider*

$$n = 10, \quad r = 4, \quad h = 1,$$

so that $t = 10 - 8 + 1 + 2 = 5$. For $q \geq 5$ and $5 \leq m \leq q$, Theorem 4.3 gives a configuration in which distinct institutions have authorization spaces intersecting in exactly one dimension, and every coalition of size c has the following leakage and residual privacy dimensions:

c	leakage dimension	residual privacy dimension
1	4	6
2	7	3
3	8	2
4	9	1
5	10	0

Thus no coalition of four institutions can determine the complete plaintext, while every coalition of five can. Moving from one to two institutions produces a jump of $r - h = 3$ leaked dimensions, which is the unavoidable price of requiring two institutions to have only a one-dimensional overlap. After the second institution, each additional institution contributes only one further independent dimension.

6. FINITE-GEOMETRIC INTERPRETATION

In this section we place the configurations of Section 4 in the framework of finite geometry. We first collect the notions we need. We then show that the optimal configuration is simultaneously a sunflower, an equidistant subspace code and a generalized arc, and that its orthogonal complements form a generalized dual arc.

All dimensions below are vector-space dimensions over $\mathbb{F}_q$. In finite geometry, subspaces of V are often viewed as subspaces of the projective space $\text{PG}(n - 1, q)$, whose projective dimension is the vector dimension minus one; when we quote definitions from the literature, we translate them into vector dimensions.

Definition 6.1. A constant-dimension code is a subset $\mathcal{C} \subseteq \text{Gr}_q(r, n)$. Its minimum distance is the minimum of the subspace distance

$$d_S(U, W) = 2r - 2 \dim(U \cap W)$$

over distinct $U, W \in \mathcal{C}$. The code is equidistant if $d_S(U, W)$ takes the same value for all distinct $U, W \in \mathcal{C}$, that is, if $\dim(U \cap W)$ is constant.

Definition 6.2. A family $\{W_1, \dots, W_m\} \subseteq \text{Gr}_q(r, n)$ with $m \geq 2$ is a sunflower with center H if $W_i \cap W_j = H$ for all $i \neq j$. A sunflower with center $\{0\}$, that is, a family of r -dimensional subspaces intersecting pairwise trivially, is a partial r -spread.

Every sunflower is an equidistant code. The converse fails in general, but Etzion and Raviv [10] showed that, for r -dimensional subspaces of a sufficiently large ambient space, the largest equidistant codes are always sunflowers.

The following notions, in the form given by Klein, Schillewaert and Storme [16], generalize arcs and dual arcs of projective spaces.

Definition 6.3. Let $d \geq 1$.

- (a) Let $\delta_1 < \delta_2 < \dots < \delta_{d+1} < n$. A generalized arc of order d with dimensions $\delta_1, \dots, \delta_{d+1}$ is a family of δ_1 -dimensional subspaces of V such that any j distinct members span a subspace of dimension δ_j , for $1 \leq j \leq d + 1$, and any $d + 2$ distinct members span V .
- (b) Let $n > \delta_1 > \delta_2 > \dots > \delta_{d+1} > 0$. A generalized dual arc of order d with dimensions $\delta_1, \dots, \delta_{d+1}$ is a family of δ_1 -dimensional subspaces of V such that any j distinct members intersect in a subspace of dimension δ_j , for $1 \leq j \leq d + 1$, and any $d + 2$ distinct members intersect in $\{0\}$.

For instance, a generalized arc of order $n - 2$ with dimensions $1, 2, \dots, n - 1$ is a classical arc: a set of points of $\text{PG}(n - 1, q)$, any n of which span the whole space. A generalized dual arc of order 1 with dimensions δ_1 and 1 is a family of δ_1 -dimensional subspaces, any two of which meet in a one-dimensional subspace (a projective point) and any three of which meet trivially; these are the dimensional dual arcs studied, for example, in [23]. Constructions of generalized dual arcs from Veronesean-type varieties can be found in [13, 16].

Arcs and dual arcs are exchanged by orthogonal duality, for any nondegenerate bilinear form.

Lemma 6.4. *Let $U, W, W_1, \dots, W_c \leq V$. Then:*

- (i) $\dim W^\perp = n - \dim W$;
- (ii) $(W_1 + \dots + W_c)^\perp = W_1^\perp \cap \dots \cap W_c^\perp$;
- (iii) $(U \cap W)^\perp = U^\perp + W^\perp$.

Consequently, $W \mapsto W^\perp$ maps $\text{Gr}_q(r, n)$ to $\text{Gr}_q(n - r, n)$ and preserves the subspace distance, and a family $\{W_1, \dots, W_m\}$ is a generalized arc of order d with dimensions $\delta_1, \dots, \delta_{d+1}$ if and only if $\{W_1^\perp, \dots, W_m^\perp\}$ is a generalized dual arc of order d with dimensions $n - \delta_1, \dots, n - \delta_{d+1}$.

Proof. Part (i) is proved as in Lemma 2.3, and (ii) follows directly from the definition of $W^\perp$. For (iii), the inclusions $U \cap W \leq U$ and $U \cap W \leq W$ give $U^\perp + W^\perp \leq (U \cap W)^\perp$. By (i) and (ii),

$$\dim(U^\perp + W^\perp) = (n - \dim U) + (n - \dim W) - (n - \dim(U + W)) = n - \dim(U \cap W),$$

which is $\dim(U \cap W)^\perp$, so equality holds. If $U, W \in \text{Gr}_q(r, n)$, then (i) and (ii) give $\dim(U^\perp \cap W^\perp) = n - 2r + \dim(U \cap W)$, hence

$$d_S(U^\perp, W^\perp) = 2(n - r) - 2(n - 2r + \dim(U \cap W)) = d_S(U, W).$$

Finally, by (i) and (ii), any j members $W_{i_1}, \dots, W_{i_j}$ satisfy $\dim(W_{i_1}^\perp \cap \dots \cap W_{i_j}^\perp) = n - \dim(W_{i_1} + \dots + W_{i_j})$, which translates each condition of Definition 6.3(a) into the corresponding condition of (b), and conversely. $\square$

For the cryptographic model, the dual family has a direct meaning. By Proposition 2.2 and Lemma 6.4(ii), a coalition I knows the plaintext exactly modulo

$$W_I^\perp = \bigcap_{i \in I} W_i^\perp.$$

Thus, while the spaces W_i describe what the institutions learn, their orthogonal complements describe what they do not learn, and the residual uncertainty of a coalition is the intersection of the complements of its members.

Let $\mathcal{W}^* = \{W_{\lambda_1}, \dots, W_{\lambda_m}\}$ be the configuration of Theorem 4.3. By Lemma 4.1, $\mathcal{W}^*$ is a sunflower with center H , where $\dim H = h$; in particular, it is an equidistant code in $\text{Gr}_q(r, n)$ with subspace distance $2(r - h)$.

Sunflowers reduce the extremal problem to the case $h = 0$.

Lemma 6.5. *Let $\{W_1, \dots, W_m\} \subseteq \text{Gr}_q(r, n)$ be a sunflower with center H of dimension h , put $k = r - h$, and let $\pi : V \rightarrow V/H$ be the canonical projection. Then $\{\pi(W_1), \dots, \pi(W_m)\}$ is a partial k -spread of $V/H \cong \mathbb{F}_q^{n-h}$ (i.e. a set of pairwise disjoint subspaces), and for every nonempty $I \subseteq [m]$,*

$$\dim\left(\sum_{i \in I} W_i\right) = h + \dim\left(\sum_{i \in I} \pi(W_i)\right).$$

Proof. Since $m \geq 2$, the center H is contained in every W_i . Hence $\dim \pi(W_i) = r - h = k$, and for $i \neq j$,

$$\pi(W_i) \cap \pi(W_j) = \pi((W_i + H) \cap (W_j + H)) = \pi(W_i \cap W_j) = \pi(H) = \{0\}.$$

If I is nonempty, then $H \leq \sum_{i \in I} W_i$, so

$$\dim\left(\sum_{i \in I} W_i\right) = h + \dim \pi\left(\sum_{i \in I} W_i\right) = h + \dim\left(\sum_{i \in I} \pi(W_i)\right).$$

□

For $\mathcal{W}^*$, the projection π identifies V/H with $A \oplus C \oplus E$ and $\pi(W_\lambda)$ with the graph U_λ . Thus $\mathcal{W}^*$ is obtained by adding the common subspace H to the partial k -spread $\{U_{\lambda_1}, \dots, U_{\lambda_m}\}$ of $\mathbb{F}_q^{n-h}$. Writing $g_{n,r,h}$ for the profile g of Remark 4.5, one checks that

$$g_{n,r,h}(c) = h + g_{n-h,k,0}(c) \quad (c \geq 1).$$

In cryptographic terms, every institution can evaluate all inner products with vectors of H , so this part of the plaintext is revealed as soon as a single institution colludes. The diversity parameter h only shifts the problem, and the actual extremal question lives in the quotient V/H . This reduction applies to sunflowers only; general configurations may have pairwise intersections of varying dimension, but by Theorem 4.4 they cannot have a smaller worst-case leakage.

Proposition 6.6. *Assume $t \geq 3$ and $t \leq m \leq q$, and let $\mathcal{W}^*$ be the configuration of Theorem 4.3. Then:*

(i) $\mathcal{W}^*$ is a generalized arc of order $t - 2$ with dimensions

$$r, \quad 2r - h, \quad 2r - h + 1, \quad \dots, \quad n - 1;$$

(ii) the family $\{W_{\lambda_1}^\perp, \dots, W_{\lambda_m}^\perp\}$ is a generalized dual arc of order $t - 2$ with dimensions

$$n - r, \quad t - 2, \quad t - 3, \quad \dots, \quad 1.$$

More precisely, for every c distinct indices $i_1, \dots, i_c$ with $2 \leq c \leq m$,

$$\dim(W_{\lambda_{i_1}}^\perp \cap \dots \cap W_{\lambda_{i_c}}^\perp) = \max\{0, t - c\}.$$

Proof. By Theorem 4.3, any j distinct members of $\mathcal{W}^*$ span a subspace of dimension $\delta_1 = r$ if $j = 1$ and $\delta_j = 2r - h + j - 2$ if $2 \leq j \leq t - 1$, and any t members span V . These dimensions satisfy $\delta_1 < \delta_2$, because $h < r$, and $\delta_{t-1} = 2r - h + t - 3 = n - 1 < n$. This proves (i) with $d = t - 2 \geq 1$. Part (ii) and the displayed formula follow from (i), Theorem 4.3 and Lemma 6.4, since $n - \delta_j = t - j$ for $2 \leq j \leq t - 1$. □

Being a generalized arc is a statement about *every* coalition, whereas Theorem 4.4 concerns the worst one. By Proposition 6.6, the optimal configuration has both properties: every coalition of size c leaks exactly $g(c)$, as observed in Remark 4.5. Two special cases illustrate the picture. For $r = 1$ and $h = 0$, the members of $\mathcal{W}^*$ are the points spanned by $(1, \lambda, \dots, \lambda^{n-1})$ in suitable coordinates, that is, m points of the normal rational curve, and Proposition 6.6(i) recovers the classical fact that they form an arc; dually, their orthogonal complements are hyperplanes in general position. For $t = 3$, that is, $n = 2r - h + 1$, the orthogonal complements are $(n - r)$ -dimensional subspaces any two of which meet in a one-dimensional subspace and any three of which meet trivially, so they form a dimensional dual arc in the sense recalled above.

7. CONCLUSION AND OPEN PROBLEMS

We have proposed a finite-geometric model for the cumulative leakage of functional keys in inner-product functional encryption. An institution is represented by the span W_i of its authorized key vectors, and a coalition by the sum W_I of these spaces. A coalition determines the plaintext exactly modulo $W_I^\perp$ and, for a uniform plaintext, learns exactly $\dim W_I \cdot \log_2 q$ bits; the resulting leakage function is a representable, and entropic, polymatroid. For authorization spaces of dimension r with pairwise intersections of dimension at most h , we proved that some coalition of size $c \geq 2$ always leaks at least $g(c) = \min\{n, 2r - h + c - 2\}$ dimensions, and we gave a Vandermonde-type configuration in which every coalition of size c leaks

exactly $g(c)$. The optimal reconstruction threshold is therefore $n - 2r + h + 2$: demanding more diverse permissions (smaller h) necessarily helps small coalitions. Geometrically, the optimal configuration is a sunflower, an equidistant subspace code and a generalized arc, and its orthogonal complements form a generalized dual arc. These results do not replace cryptographic security: a secure IPFE scheme guarantees that coalitions learn nothing beyond the outputs of their authorized functions, and our results quantify the information already contained in those outputs. We conclude with some open problems.

Size of optimal configurations. Our construction has at most q members. Determine the largest m for which there is an $(n, r, h; m)_q$ authorization configuration in which every coalition of size c leaks exactly $g(c)$; for $t \geq 3$, this is the largest generalized arc as in Proposition 6.6(i), and the diversity constraint is then automatic. For $r = 1$ and $h = 0$, this is the maximum size of an arc in $\text{PG}(n - 1, q)$, which by the MDS conjecture should equal $q + 1$ for $2 \leq n \leq q$, except when q is even and $n \in \{3, q - 1\}$, where it equals $q + 2$; the conjecture was proved by Ball for q prime [3]. For sunflower configurations, Lemma 6.5 gives the upper bound $m \leq (q^{n-h} - 1)/(q^k - 1)$. A first question is whether, as for the normal rational curve, the construction of Section 4 can be extended by a member “at infinity”.

Nonuniform and constrained authorizations. Different institutions may need authorization spaces of different dimensions r_i , in which case the problem is governed by a nonuniform subspace arrangement rather than a constant-dimension code. More realistically, one may prescribe mandatory functionality spaces $R_i \leq W_i$, encoding the computations that institution i must be able to perform, and ask for extensions W_i minimizing the leakage function. This is a constrained subspace-arrangement optimization problem, closely related to the theory of span programs [14].

Dynamic authorization. Revocation, key updates, or changes of responsibilities correspond to replacing, enlarging or shrinking individual authorization spaces. Understanding how the leakage profile changes under such operations could complement cryptographic approaches based on fine-grained or updatable functional encryption [2, 7].

Structured plaintexts. The exact entropy statements of Corollary 2.4 assume a uniform plaintext. By Remark 2.6, $\dim W_I \cdot \log_2 q$ remains an upper bound on the leakage for every source, but the actual loss depends on the distribution. Combining the geometric framework with entropy or min-entropy models of structured data is a natural next step.

ACKNOWLEDGMENTS

The author was partially supported by the Italian National Group for Algebraic and Geometric Structures and their Applications (GNSAGA – INdAM).

REFERENCES

- [1] M. Abdalla, F. Bourse, A. De Caro and D. Pointcheval, *Simple functional encryption schemes for inner products*, Public-Key Cryptography – PKC 2015, Lecture Notes in Computer Science **9020** (2015), 733–751.
- [2] M. Abdalla, D. Catalano, R. Gay and B. Ursu, *Inner-product functional encryption with fine-grained access control*, Advances in Cryptology – ASIACRYPT 2020, Lecture Notes in Computer Science **12493** (2020), 467–497.
- [3] S. Ball, *On sets of vectors of a finite vector space in which every subset of basis size is a basis*, Journal of the European Mathematical Society **14** (2012), 733–748.
- [4] A. Beigel, *Secret-sharing schemes: a survey*, Coding and Cryptology – IWCC 2011, Lecture Notes in Computer Science **6639** (2011), 11–46.
- [5] G. R. Blakley and C. Meadows, *Security of ramp schemes*, Advances in Cryptology – CRYPTO ’84, Lecture Notes in Computer Science **196** (1985), 242–268.
- [6] E. F. Brickell, *Some ideal secret sharing schemes*, Journal of Combinatorial Mathematics and Combinatorial Computing **6** (1989), 105–113.
- [7] V. Cini, S. Ramacher, D. Slamanig, C. Striecks and E. Tairi, *(Inner-product) functional encryption with updatable ciphertexts*, Journal of Cryptology **37** (2024), Article 8.
- [8] L. Csirmaz, *The size of a share must be large*, Journal of Cryptology **10** (1997), 223–231.

- [9] Y. Dodis, R. Ostrovsky, L. Reyzin and A. Smith, *Fuzzy extractors: how to generate strong keys from biometrics and other noisy data*, SIAM Journal on Computing **38** (2008), 97–139.
- [10] T. Etzion and N. Raviv, *Equidistant codes in the Grassmannian*, Discrete Applied Mathematics **186** (2015), 87–97.
- [11] S. Fujishige, *Polymatroidal dependence structure of a set of random variables*, Information and Control **39** (1978), 55–72.
- [12] A. W. Ingleton, *Representation of matroids*, in: Combinatorial Mathematics and its Applications (D. J. A. Welsh, ed.), Academic Press, London, 1971, 149–167.
- [13] W. M. Kantor and E. E. Shult, *Veroneseans, power subspaces and independence*, Advances in Geometry **13** (2013), 511–531.
- [14] M. Karchmer and A. Wigderson, *On span programs*, Proceedings of the 8th Annual Structure in Complexity Theory Conference (1993), 102–111.
- [15] R. Kinser, *New inequalities for subspace arrangements*, Journal of Combinatorial Theory, Series A **118** (2011), 152–161.
- [16] A. Klein, J. Schillewaert and L. Storme, *Generalised Veroneseans*, preprint, arXiv:1309.2008.
- [17] R. Koetter and F. R. Kschischang, *Coding for errors and erasures in random network coding*, IEEE Transactions on Information Theory **54** (2008), 3579–3591.
- [18] R. La Scala and A. Meneghetti, *Inner product functional encryption based on the UOV scheme*, Journal of Mathematical Cryptology **19** (2025), no. 1, Article 20240026. <https://doi.org/10.1515/jmc-2024-0026>
- [19] C. Mascia, M. Sala and I. Villa, *A survey on functional encryption*, Advances in Mathematics of Communications **17** (2023), 1251–1289.
- [20] R. J. McEliece and D. V. Sarwate, *On sharing secrets and Reed–Solomon codes*, Communications of the ACM **24** (1981), 583–584.
- [21] J. Oxley, *Matroid Theory*, 2nd ed., Oxford Graduate Texts in Mathematics **21**, Oxford University Press, Oxford, 2011.
- [22] A. Shamir, *How to share a secret*, Communications of the ACM **22** (1979), 612–613.
- [23] S. Yoshiara, *Ambient spaces of dimensional dual arcs*, Journal of Algebraic Combinatorics **19** (2004), 5–23.

DIPARTIMENTO DI MATEMATICA E FISICA, UNIVERSITÀ DEGLI STUDI DELLA CAMPANIA “LUIGI VANVITELLI”, CASERTA, ITALY

Email address: ferdinando.zullo@unicampania.it