Fully Anonymous Perfect Threshold Secret Sharing: Explicit One-Bit Construction and Rate Amplification

Chongxu Ren and Hongbo Yu

Department of Computer Science and Technology,
Tsinghua University, Beijing, China
rcx23@mails.tsinghua.edu.cn, yuhongbo@tsinghua.edu.cn

Abstract. Secret sharing distributes a secret among participants so that only sufficiently large groups can recover it. Fully anonymous secret sharing also allows recovery from share values alone, without participant identities. Below the reconstruction threshold, the shares must look exactly like independent uniform random values, revealing neither the secret nor their holders' identities. Con, Li, and Mazor established that such schemes with error-free recovery exist for every threshold, leaving an efficient general construction open.

We give efficient constructions for every threshold $2 \leq t \leq n$ among n participants. For an L-bit secret, each share has $L + O(t^2 \log^2 n)$ bits. Sharing runs in expected polynomial time, and reconstruction runs in deterministic polynomial time. Our approach combines an efficient one-bit construction with a general transformation that uses it to protect a short message needed for anonymous reconstruction of a long secret. The additional share length has an upper bound depending only on the number of participants and the threshold. For a fixed participant count and threshold, the rate approaches one as secrets grow longer, while anonymity remains exact and recovery never errs.

1 Introduction

Secret sharing makes recovery of a secret depend on cooperation. A dealer assigns shares to n participants so that authorized groups can recover the secret, while unauthorized groups learn nothing about it [1]. In a (t, n) threshold scheme, every t participants can reconstruct and every $t - 1$ participants have no information about the secret. Shamir's scheme [16] achieves this by evaluating a random polynomial at public points assigned to the participants. Reconstruction uses both the values and the points at which they were evaluated.

This dependence on labels is natural for ordinary secret sharing, but secrecy alone says nothing about the identities revealed by a share. Appending a participant's name or a common execution tag preserves secrecy while revealing who holds the share or which shares were generated together. Anonymous secret sharing seeks to hide this information from unauthorized groups and to reconstruct without participant labels. These requirements arise, for example,

in abuse-resistant location tracking, where anonymous sharing helps reconcile owner privacy with the detection of covert tracking devices [9].

We study the strongest information-theoretic anonymity requirement of Bishop et al. [2], together with perfect reconstruction. *Share anonymity* requires that, for each fixed secret, the shares held by any $t-1$ distinct participants are independent and uniform over the share alphabet. An unauthorized collection therefore has the same distribution regardless of the secret or the identities of its holders, and its values carry no common tag linking them to a sharing. *Perfect anonymous reconstruction* requires that every t distinct participants recover the secret with certainty from their share values alone, without using their identities or the order in which the shares are received. Section 2 gives the formal definition.

The two requirements create a tension. Reconstruction must succeed without the evaluation labels used by standard schemes, yet unauthorized shares must look exactly like independent random values. A natural modification of Shamir's scheme includes a fresh random evaluation point in each share. This removes the fixed link to a participant, but collisions can prevent reconstruction. Sampling distinct points avoids collisions, but two resulting shares have dependent evaluation points. For thresholds above two, such a pair is unauthorized, so this dependence violates exact share anonymity. The random-point approach of Eldridge et al. [9] allows a small reconstruction error. Con [4] obtains perfect guarantees by separating the privacy and reconstruction thresholds. The exact threshold target asks for uniformity of $t-1$ shares and certain recovery from t shares.

Con, Li, and Mazor [6] established that this exact target is attainable. They give an efficient one-bit construction for threshold three and prove that $O(t \log n)$-bit shares suffice at every threshold. Their general construction determines a suitable distribution by finite rational linear algebra, but does not give a polynomial-time general dealer. The remaining question is therefore algorithmic: how can one sample the required distribution efficiently and exactly for every threshold?

Share length raises a second question. If a one-bit scheme uses b bits per share, sharing an L-bit secret one bit at a time uses bL bits per share. The rate, defined as secret length divided by share length, remains $1/b$ regardless of L. To approach rate one, the cost of anonymous reconstruction must be shared across the secret rather than paid separately for each bit. This leads to our central question:

> *Can every threshold admit efficient exact anonymous sharing whose additive overhead has an upper bound independent of the secret length?*

1.1 Our results

We construct fully anonymous perfect (t, n) threshold sharing for L-bit secrets with $M \le L + O(t^2 \log^2 n)$ bits per share, expected-polynomial-time sharing, and deterministic polynomial-time reconstruction. The bound includes padding and holds at every threshold $2 \le t \le n$. No computational assumption is required. For fixed (n, t), the overhead has an upper bound independent of L, so the rate L/M approaches one. This answers the constant-rate question of Con, Li, and Mazor [6] for sufficiently long secrets. Theorems 1 and 2 give the two ingredients; Corollary 1 derives the rate guarantee. Table 1 compares share-length bounds and dealer guarantees under this same anonymity and reconstruction target. Write $m = \lceil \log_2 n \rceil$; in the compiler row, b is the one-bit share length.

Table 1. Exact fully anonymous perfect threshold sharing. Share lengths are in bits. Poly. time refers to expected polynomial dealer time after public preprocessing; the compiler inherits its primitive's time bounds. Our one-bit construction also has deterministic polynomial-time preprocessing.

Work	Secret	Threshold	Share bits	Poly. time	Ref.
CLM	1	$2 \le t \le n$	$O(t \log n)$	Not given	[6]
CLM	1	$t = 3$	$2m$	Yes	[6]
Ours	1	$3 \le t < n$	$m \min\{2t-4, n-2\}$	Yes	Theorem 2
Ours	L	$2 \le t \le n$	$m\lceil L/m \rceil + bm \min\{t-1, \lceil L/m \rceil\}$	Inherited	Theorem 1

A compiler with additive overhead. From any exact one-bit scheme with b-bit shares, our compiler produces an L-bit scheme with

$$M \le L + (m-1) + bm(t-1).$$

Sharing and reconstruction inherit the primitive's expected and deterministic polynomial-time guarantees. When $m \mid L$ and $L \ge m(t-1)$, the length is exactly $L + bm(t-1)$. Thus a one-bit primitive with $b = O(t \log n)$ suffices for the combined bound above. Section 3 proves this generic transformation before choosing its one-bit implementation.

An efficient exact one-bit primitive. For $3 \le t < n$, our one-bit scheme has

$$b = \min\{2t-4, n-2\}\lceil \log_2 n \rceil$$

bits per share. Threshold two uses m bits, and $t = n$ uses one bit. The exact dealer terminates almost surely and introduces no reconstruction or anonymity error. This supplies an efficient primitive at every threshold with $O(t \log n)$-bit shares, matching the asymptotic existence bound of Con, Li, and Mazor [6]. At threshold three, both constructions use $2m$ bits per share. Section 4 gives the construction and proof; Section 5.1 accounts for the exact sampler's preprocessing and complete bit costs.

1.2 Technical overview

Amortize anonymity through a common direction space. The compiler makes every authorized group derive the same space containing the secret, then uses one-bit sharing to identify the secret within that space. Encode the secret as a vector S over a field with at least n elements. At distinct public points α_i, the dealer forms the payloads

$$Y_i = A_0 + \sum_{j=1}^{t-2} \alpha_i^j A_j + \alpha_i^{t-1} S,$$

where $A_0, \ldots, A_{t-2}$ are independent uniform mask vectors. Any $t-1$ payloads are independent and uniform. Every t payloads, however, determine the same span of their pairwise differences,

$$U = \operatorname{span}(A_1, \ldots, A_{t-2}, S).$$

The reconstruction algorithm obtains U directly from the payload values; it never needs to match them to evaluation points. This identity holds even when the sampled mask vectors are linearly dependent or payloads coincide.

The space U has dimension at most $t-1$. In its canonical reduced basis, the coordinates of S at the pivot positions uniquely identify S. The dealer shares these coordinates using independent invocations of the one-bit scheme. Every authorized group computes the same basis from its payloads and recovers the secret from the auxiliary shares. The basis itself need not be transmitted. At most $t-1$ field coordinates require this extra protection, regardless of the length of S, which gives an additive overhead bound independent of the secret length. For unauthorized groups, the auxiliary shares remain uniform after conditioning on the payloads, preserving exact share anonymity. Section 3 gives the complete encoding, padding, algorithms, and proof.

Encode a bit in affine relations. The auxiliary primitive represents shares as vectors over a finite field. An affine relation is a linear combination that vanishes and whose coefficients sum to zero. The decoder asks whether every received vector is an affine combination of the others. This test uses only the values. For secret zero, evaluations of low-degree polynomials make the answer yes. For secret one, the dealer chooses a relation space supported on at most $t-1$ participants and realizes it using polynomial-evaluation vectors. Every authorized set contains a participant outside that support, whose vector cannot be an affine combination of the others, so the answer is no.

To hide this distinction, the sampled relation space must restrict on every unauthorized set to the relation law of independent uniform vectors. We obtain the exact masses by inverting containment probabilities and sample efficiently by grouping spaces according to support size and dimension. A uniformly random invertible affine transformation then makes the values uniform conditional on their relations. Matching both laws gives independent uniform shares, including

collisions and dependent configurations. Section 4.2 explains the relation test, Section 4.4 proves correctness and anonymity, and Section 4.5 derives the exact sampler.

1.3 Related work

Early work studied reconstruction without participant identities [17,14,3,12]. Deng, Guo, and Liu [8] developed further threshold-two constructions using combinatorial designs, including group-divisible designs and difference families. Daza and Domingo-Ferrer [7] introduce partial anonymity, which hides a holder's identity within a specified subset of participants, and give a general construction. The classical literature also establishes limits that apply to our stronger target. In particular, the ideality characterization of Phillips and Phillips [14] and the alphabet bounds of Kishimoto et al. [12] rule out finite rate one for $1 < t < n$.

Hiding information carried by unauthorized shares is a separate requirement [11,13]. Bishop et al. [2] formalized a hierarchy of share-anonymity notions, from hiding participant identities within one sharing to hiding linkage across independent sharings and, most strongly, uniform anonymity. Paskin-Cherniavsky and Olimid [13] give statistical approximations to anonymity. Eldridge et al. [9] use random evaluation points for abuse-resistant location tracking, allowing a small reconstruction error. Definition 1 specifies the exact uniform-anonymity and reconstruction requirements used here.

Reed–Solomon codes offer another approach to anonymous reconstruction. Con [4] relates the problem to codes correcting permutations, insertions, and deletions, obtaining perfect anonymous ramp schemes with separated privacy and reconstruction thresholds. For ordinary insertion and deletion errors, Con et al. [5] approach the half-Singleton tradeoff over linear-size alphabets. Zhang et al. [18] obtain polynomial-size alphabets while also permitting permutations, yielding fully anonymous ramp schemes through Con's reduction.

Sahni and Koga [15] give a one-bit anonymous ramp scheme based on Fermat's little theorem, with perfect reconstruction and single-dealer anonymity. These ramp constructions separate the privacy and reconstruction thresholds; our exact threshold target places them at $t - 1$ and t.

Farràs and Guiot [10] use anonymous sharing to construct traceable schemes for general access structures. They introduce local share projections whose reconstruction inputs hide the participating set, and also study access-structure hiding. Their Shamir-based threshold construction has perfect correctness and single-dealer anonymity, and their general-access construction has statistical correctness.

For the exact threshold target, Con, Li, and Mazor [6] establish existence at every threshold and supply the direct one-bit baseline in Table 1. Their general argument constructs a distribution avoiding all forbidden authorized tuples while keeping small marginals uniform. Our construction specifies an affine-relation law and samples it exactly by grouping relation spaces according to support size and dimension.

1.4 Organization

Section 2 defines the model and the one-bit interface. Section 3 gives the generic rate compiler. Section 4 supplies the efficient exact one-bit primitive, with its sampling derivation after the correctness argument. Section 5 analyzes costs, gives a sparse refinement, and relates the results to known alphabet limits. Section 6 discusses the remaining questions on share length.

2 Preliminaries

All logarithms are to base two. We write $[n] = \{1, \dots, n\}$. For a binary field $\mathbb{F}_{2^a}$, a fixed polynomial basis identifies each element with an a-bit string. We use span, ker, and rank for linear span, kernel, and rank. The direction space of points $y_1, \dots, y_t$ is $\mathrm{span}(y_2 - y_1, \dots, y_t - y_1)$; it is independent of the chosen basepoint. Section 3.1 defines the canonical basis used by the compiler, and Section 4.2 introduces the affine-relation notation needed for one-bit sharing.

For public parameters (n, t), a nonempty finite secret alphabet Σ, and a nonempty finite share alphabet Ω, a scheme is a pair of algorithms

$$(\mathsf{Share}_{n,t}, \mathsf{Recon}_{n,t}), \qquad \mathsf{Share}_{n,t} : \Sigma \to \Omega^n, \quad \mathsf{Recon}_{n,t} : \Omega^t \to \Sigma.$$

The sharing algorithm is randomized, uses participant labels internally, and terminates almost surely. The reconstruction algorithm is a total deterministic function on Ω^t, receives only share values, and must be invariant under permutations of its inputs, including tuples outside the sharing support. The dealer may use public preprocessing independent of the secret. We charge it separately from online sharing and reconstruction. Randomized running-time bounds are expectations over fair-bit randomness; exact samplers run to completion.

Definition 1 (Fully anonymous perfect threshold sharing). *Let $2 \le t \le n$. The pair $(\mathsf{Share}, \mathsf{Recon})$ is a fully anonymous perfect (t, n) threshold scheme if the following two properties hold.*

- **Share anonymity.** *For every secret $s \in \Sigma$, let $W \leftarrow \mathsf{Share}(s)$. For every ordered sequence of $t-1$ distinct participants $(i_1, \dots, i_{t-1})$, and every $(x_1, \dots, x_{t-1}) \in \Omega^{t-1}$,*

$$\Pr[(W_{i_1}, \dots, W_{i_{t-1}}) = (x_1, \dots, x_{t-1}) \mid s] = |\Omega|^{-(t-1)}.$$

- **Perfect anonymous reconstruction.** *For every ordered sequence of t distinct participants, every secret s, and every $w \in \mathrm{supp}(\mathsf{Share}(s))$,*

$$\mathsf{Recon}(w_{i_1}, \dots, w_{i_t}) = s.$$

For more than t honest shares, any t may be used. Marginalizing the first property gives the same product-uniform law for every number $a < t$ of distinct participants. In particular, it implies perfect secrecy and makes the observed law independent of the participant sequence. All reconstruction guarantees concern

shares from one honest execution. Distinct participant origins are an input guarantee: repeating one participant's share does not count toward the threshold, and different participants may hold equal share values. The decoder is defined on other inputs as well; robustness against altered shares and recovery from mixtures of different executions are separate properties.

Definition 2 (One-bit auxiliary primitive). *Fix (t, n). A one-bit auxiliary primitive is a scheme satisfying Definition 1 with $\Sigma = \{0,1\}$ and share alphabet $\Gamma = \{0,1\}^b$. Its algorithms are denoted by Share_1 and Recon_1. Independent invocations use independent randomness.*

Section 3 uses this interface without assuming a particular implementation. Section 4 supplies an efficient exact primitive at every threshold.

3 High-rate anonymous sharing

We turn any one-bit primitive with b-bit shares into a scheme for long secrets. Independent one-bit sharings would use bL bits per share for an L-bit secret. Our compiler gives each participant a payload of essentially L bits and uses the one-bit primitive only for a short auxiliary message. That message has at most $t-1$ field coordinates, so its sharing cost has an upper bound independent of L.

The payloads are evaluations of a vector polynomial whose leading coefficient is the secret. An authorized group receives these values without their evaluation labels. It can nevertheless compute a common subspace containing the secret; the auxiliary message identifies the secret within this space. We use only the interface of Definition 2. Its efficient implementation is given in Section 4.

3.1 Recovering the secret without labels

Fix $2 \le t \le n$ and $L \ge 1$. Put $m = \lceil \log_2 n \rceil$, choose a public field width $a \ge m$, and set $k = t-1$ and $d = \lceil L/a \rceil$. Public setup fixes the field $F = \mathbb{F}_{2^a}$, represented using a monic irreducible binary polynomial of degree a, and distinct evaluation points $\alpha_1, \ldots, \alpha_n \in F$. Zero is allowed among these points; the default width is $a = m$. Append $ad - L$ zero bits to the secret and use the fixed polynomial basis to encode it as $S \in F^d$.

What the payloads reveal. The dealer samples independent uniform masks $A_0, \ldots, A_{k-1} \in F^d$ and gives participant i the payload $Y_i = P(\alpha_i)$, where

$$P(X) = A_0 + \sum_{j=1}^{k-1} X^j A_j + X^k S.$$

The k random coefficients make the payloads of any k participants independent and uniform, as in polynomial secret sharing. With $k + 1 = t$ labeled evaluations, one could interpolate P and read off S. Anonymous reconstruction has no pairing

between values and evaluation points. The useful quantity it can recover is the span of the payload differences: for any t distinct participants,

$$\operatorname{span}(Y_{i_2} - Y_{i_1}, \ldots, Y_{i_t} - Y_{i_1}) = \operatorname{span}(A_1, \ldots, A_{k-1}, S) = U.$$

To see why, augment each payload with a leading coordinate 1. Evaluation at t distinct points is an invertible transformation of the augmented coefficient rows. Taking differences selects the part with first coordinate zero, removing the translation A_0 and leaving exactly U. Lemma 1 proves this identity. It holds even when the masks are linearly dependent or payload values coincide.

What the auxiliary message supplies. The space U contains S and has dimension at most $h = \min(k, d)$. Every authorized group obtains this same space from its payload values, so the dealer and decoder can choose the same basis without transmitting it. Let $\operatorname{rref}(U)$ be the unique matrix of nonzero rows in reduced row echelon form spanning U. Write its rows as $v_1, \ldots, v_r$ and their pivot positions as $p_1 < \cdots < p_r$. In particular, $v_j[p_\ell] = 1$ if $j = \ell$ and 0 otherwise. Every $S \in U$ therefore satisfies

$$S = \sum_{j=1}^{r} S[p_j] v_j. \tag{1}$$

Thus the r pivot coordinates identify S within U. Pad them with zeros to a vector $\Theta \in F^h$ and share its ah encoded bits using independent invocations of the one-bit primitive. Padding to h gives every share the same length even when $\dim U$ varies with the masks and secret. Each complete share consists of the payload and these auxiliary shares, with alphabets and length

$$\Sigma = \{0,1\}^L, \qquad \Omega = F^d \times \Gamma^{ah}, \qquad M = ad + bah.$$

3.2 Compiler and guarantee

Theorem 1 (Direction-space rate compiler). *Given a one-bit auxiliary primitive as in Definition 2, Algorithm 1 is a fully anonymous perfect (t, n) threshold scheme for every $2 \le t \le n$ and $L \ge 1$. Its secret and share lengths are, respectively,*

$$L \quad \textit{and} \quad M = a\lceil L/a \rceil + ba \min(t-1, \lceil L/a \rceil), \qquad a \ge \lceil \log_2 n \rceil. \tag{2}$$

After public field setup, sharing has expected polynomial cost whenever the auxiliary sharing does; reconstruction is deterministic polynomial time whenever auxiliary reconstruction is.

Algorithm 1 gives the complete transformation, using the public encodings fixed above. Section 3.3 proves the theorem, and Section 3.4 derives the rate bound.

Algorithm 1 (common-direction-space compiler).

Public parameters: $(n, t, L, a, k, d, h, F, \alpha_1, \ldots, \alpha_n)$ as above; $(\mathsf{Share}_1, \mathsf{Recon}_1)$ as in Definition 2.

Share$(s \in \{0,1\}^L)$:

1. Encode s with trailing zeros as $S \in F^d$. Sample $A_0, \ldots, A_{k-1} \xleftarrow{\$} F^d$ independently.
2. Set
$$Y_i = A_0 + \sum_{j=1}^{k-1} \alpha_i^j A_j + \alpha_i^k S \quad (i \in [n]),$$
$$U = \mathrm{span}(A_1, \ldots, A_{k-1}, S).$$
3. Compute the rows $v_1, \ldots, v_r$ and pivots $p_1, \ldots, p_r$ of $\mathrm{rref}(U)$. Set
$$\Theta = (S[p_1], \ldots, S[p_r], \underbrace{0, \ldots, 0}_{h-r}) \in F^h.$$
Encode Θ as $\theta \in \{0,1\}^{ah}$ in the fixed field basis.
4. For each $\ell \in [ah]$, independently compute $(Z_{1,\ell}, \ldots, Z_{n,\ell}) \leftarrow \mathsf{Share}_1(\theta_\ell)$. Return $(W_i)_{i \in [n]}$, where $W_i = (Y_i, Z_{i,1}, \ldots, Z_{i,ah}) \in \Omega$.

Recon$(w_1, \ldots, w_t \in \Omega)$:

1. Parse $w_i = (y_i, z_{i,1}, \ldots, z_{i,ah})$. Set $\theta_\ell = \mathsf{Recon}_1(z_{1,\ell}, \ldots, z_{t,\ell})$ for every $\ell \in [ah]$, and parse θ as $\Theta \in F^h$.
2. Set $U' = \mathrm{span}(y_2 - y_1, \ldots, y_t - y_1)$. Compute the nonzero rows $v_1, \ldots, v_r$ and pivots $p_1, \ldots, p_r$ of $\mathrm{rref}(U')$. Set
$$S' = \sum_{j=1}^{r} \Theta_j v_j.$$
3. If some $\Theta_j \neq 0$ for $r < j \leq h$, or the last $ad - L$ bits of S' are not all zero, return 0^L. Otherwise return the first L bits of S'.

3.3 Correctness and exact anonymity

We first prove that authorized payloads recover U and that unauthorized payloads are uniform. The theorem proof then combines these facts with the one-bit primitive. For anonymity, the auxiliary message depends on the same masks as the payloads, so the proof must establish uniformity of the complete shares jointly. We do this by conditioning on the masks.

Lemma 1 (Common direction space). *For every choice of $S, A_0, \ldots, A_{k-1}$ and every t distinct participants $i_1, \ldots, i_t$,*

$$\mathrm{span}(Y_{i_2} - Y_{i_1}, \ldots, Y_{i_t} - Y_{i_1}) = U.$$

Proof. Let $B \in F^{t\times(d+1)}$ have first row $(1, A_0)$, middle rows $(0, A_j)$ for $1 \le j < k$, and last row $(0, S)$. For the Vandermonde matrix $V_I = (\alpha_{i_\ell}^j)_{\ell\in[t],\, 0\le j\le k}$, the payload equations give

$$\begin{pmatrix} 1 & Y_{i_1} \\ \vdots & \vdots \\ 1 & Y_{i_t} \end{pmatrix} = V_I B.$$

The points are distinct, so V_I is invertible. Consequently the augmented payload rows span the same space as the augmented coefficient rows, without any rank assumption on the latter.

Consider the vectors in this common row space whose first coordinate is zero. On the coefficient side, these are precisely $\{0\} \times U$. On the payload side, they are linear combinations of $(1, Y_{i_\ell})$ whose coefficients sum to zero. Each such combination is a linear combination of $(0, Y_{i_\ell} - Y_{i_1})$ for $2 \le \ell \le t$, and conversely every difference has first coordinate zero. Hence these differences span U. □

Lemma 2 (Payload uniformity). *Fix a secret s and an ordered sequence I of $u < t$ distinct participants. Then $(Y_i)_{i\in I}$ is uniform over $(F^d)^u$.*

Proof. Fix one of the d vector coordinates. Its k mask entries are independent uniform field elements, and the corresponding payload entries are their linear image plus the fixed secret contribution. The linear map has rows $(1, \alpha_i, \ldots, \alpha_i^{k-1})$ for $i \in I$. Its first u columns form an invertible Vandermonde matrix, so the map is surjective onto F^u. A surjective linear map sends a uniform input to a uniform output. Adding the fixed secret contribution preserves that uniformity. The mask entries are also independent across vector coordinates, giving the claimed joint law. □

Proof (of Theorem 1). Reconstruction. On honest shares, the auxiliary decoders recover the dealer's Θ. Lemma 1 gives $U' = U$, so the decoder obtains the same reduced basis and pivot positions as the dealer. Equation (1) then gives $S' = S$. The unused coordinates of Θ and the secret's trailing padding are zero, so the validation succeeds and the output is s.

The decoder is also total and permutation-invariant on all of Ω^t. Every auxiliary decoder has these properties. The difference span U' is independent of the order and chosen basepoint of the received payloads, and its reduced basis is unique. Its rank is at most h, since it is generated by $t - 1 = k$ vectors in F^d, so the h entries of Θ always suffice. Thus S', the validation conditions, and the output depend only on the input multiset. Every step is defined and terminates, with output 0^L on a validation failure.

Anonymity. Fix a secret s and an ordered sequence I of $u < t$ distinct participants. Write $Y_I = (Y_i)_{i\in I}$ and $Z_I = (Z_{i,\ell})_{i\in I,\, \ell\in[ah]}$ for their payloads and auxiliary shares. For any prescribed values $y \in (F^d)^u$ and $z \in \Gamma^{u\times ah}$, we compute their joint probability. Condition on all masks $A = (A_0, \ldots, A_{k-1})$. This fixes both the payloads and the auxiliary message θ. Each auxiliary invocation uses fresh independent randomness and has a uniform marginal on I for either

message bit. Across all ah invocations, this gives

$$\Pr[Z_I = z \mid A, s] = 2^{-bahu}.$$

In particular, this conditional probability is the same for every mask choice, even though the message depends on the masks. Summing over the choices for which $Y_I = y$ and applying Lemma 2 gives

$$\Pr[(Y_I, Z_I) = (y, z) \mid s] = \Pr[Y_I = y \mid s]2^{-bahu} = 2^{-adu}2^{-bahu} = 2^{-Mu}.$$

This is the joint product-uniform law required by Definition 1.

Length and running time. Each share contains one d-coordinate payload and ah auxiliary b-bit shares, hence $M = ad + bah$ bits. Substituting $d = \lceil L/a \rceil$ and $h = \min(t-1, d)$ gives (2). The compiler makes $ah \le a(t-1)$ auxiliary invocations and uses polynomially many field operations. It therefore inherits the primitive's expected polynomial sharing time and deterministic polynomial reconstruction time. The complete setup and online bit costs are given in Section 5.2. □

3.4 Rate and secret length

The payload occupies $a\lceil L/a \rceil$ bits, which is L plus the field padding. The auxiliary part occupies at most $ba(t-1)$ bits. At the default width $a = m$, fixing (n, t) and the one-bit primitive therefore gives an additive overhead bound independent of L.

Corollary 1 (Rate approaching one). *For a fixed one-bit primitive with b-bit shares, Algorithm 1 with $a = m = \lceil \log_2 n \rceil$ satisfies*

$$M \le L + (a-1) + ba(t-1).$$

If $a \mid L$ and $L \ge a(t-1)$, then

$$M = L + ba(t-1), \qquad \rho_L = \frac{L}{L + ba(t-1)}.$$

For every integer $c \ge 1$, taking $L = cba(t-1)$ gives rate $c/(c+1)$. With $b = O(t \log n)$, constant rate is therefore attained with $L = O(t^2 \log^2 n)$, and the rate tends to one as L grows.

Proof. The padding adds at most $a-1$ bits and $h \le t-1$, giving the first bound. Under alignment and the stated length condition, $ad = L$ and $h = t-1$ in (2). Since $b, c \ge 1$, the chosen L satisfies that condition and gives the displayed rate.

Rate one is the information-theoretic upper bound. For a uniform secret S, a set I of $t-1$ participants, and $j \notin I$, privacy and correctness give

$$\log|\Sigma| = H(S \mid W_I) = I(S; W_j \mid W_I) \le H(W_j) \le \log|\Omega|.$$

It follows that $\rho_L \le 1$. □

Efficient instantiation. For $3 \le t < n$, substitute the one-bit primitive of Theorem 2, with $b = m \min\{2t-4, n-2\}$. At the default width $a = m$, the resulting share length is

$$M = m\lceil L/m \rceil + m^2 \min\{2t-4, n-2\} \min\{t-1, \lceil L/m \rceil\}.$$

The endpoint primitives in Section 4 also have $b = O(t \log n)$, so Corollary 1 gives $L + O(t^2 \log^2 n)$ bits per share at every threshold, including padding. Public preprocessing is deterministic polynomial time; sharing takes expected polynomial time and reconstruction takes deterministic polynomial time. Section 5.2 derives the complete setup and online costs.

For short secrets, the public parameters may select the shorter of independent one-bit sharing, with length bL, and Algorithm 1 at $a = m$. This choice depends only on (n, t, L) and preserves the exact guarantees. For $t = n$, vector XOR sharing already gives L-bit shares: sample $n-1$ independent uniform L-bit strings and choose the last to make their XOR equal to the secret. Reconstruction is XOR, and every $n-1$ shares are independent uniform. The trivial $t = 1$ endpoint uses replication.

For $1 < t < n$ and a nontrivial finite secret alphabet, the anonymous-sharing bounds discussed in Section 5.4 rule out attaining rate one. The limiting rate in Corollary 1 is a supremum.

4 Efficient exact one-bit sharing

We now construct the one-bit primitive used in Section 3. The shares are vectors, and the decoder recovers the bit from the affine relations among them. Every t participants must exhibit the relation pattern for the chosen bit, while every $t-1$ participants must receive independent uniform vectors. To meet both requirements, the secret-one dealer first samples a relation space, then realizes it as points and applies a random affine map. On each unauthorized set, the sampled space determines the relation pattern, and the affine map makes the values uniform among tuples with that pattern. Section 4.2 explains the reconstruction test and why matching relation-space distributions suffices for exact anonymity.

4.1 Parameters and guarantee

Fix $3 \le t < n$ and set

$$m = \lceil \log_2 n \rceil, \qquad q = 2^m, \qquad k = t-1, \qquad D = D_0 = \min\{2k-2, n-2\}.$$

The secret alphabet is $\Sigma = \{0, 1\}$ and the share alphabet is $\Omega = \mathbb{F}_q^D$. Here D is the dimension of a one-bit share; the compiler's payload dimension d depends on the length of the secret being compiled. The choice D_0 meets two requirements: it permits the prescribed relations to be realized on every relevant small set, and it makes their exact probability masses positive. Lemmas 4 and 3 establish these requirements, respectively.

Choose distinct public points $\alpha_1, \ldots, \alpha_n \in \mathbb{F}_q$. A polynomial-basis encoding represents a field element by m bits; concatenating coordinates identifies $\mathbb{F}_q^D$ with $\{0,1\}^{Dm}$. Public preprocessing fixes the field representation and the integer type weights used by the sampler in Section 4.5. It is performed once for fixed (n,t).

Theorem 2 (Explicit general-threshold construction). *For every $3 \le t < n$, there is a fully anonymous perfect one-bit (t,n) threshold scheme with share alphabet $\mathbb{F}_q^D$, where*

$$q = 2^{\lceil \log_2 n \rceil}, \qquad D = \min\{2t-4, n-2\}.$$

Each share has exactly Dm bits, where $m = \lceil \log_2 n \rceil$. Public preprocessing is deterministic polynomial time in (n,t). After preprocessing, sharing takes expected polynomial time and reconstruction takes deterministic polynomial time in (n,t). Sharing terminates almost surely and has no correctness or anonymity error. The complete operation and bit costs are given in Section 5.1.

The endpoints have simpler constructions. For $t = 2$, choose distinct $\alpha_i \in \mathbb{F}_q$, draw $z \leftarrow \mathbb{F}_q$ and $a \leftarrow \mathbb{F}_q \setminus \{0\}$ uniformly and independently, and set $W_i = z$ for secret zero and $W_i = a\alpha_i + z$ for secret one. The decoder returns zero exactly when its two inputs are equal. Every individual share is uniform, and the share length is m bits. For $t = n$, draw independent uniform bits $W_1, \ldots, W_{n-1}$ and set $W_n = s \oplus W_1 \oplus \cdots \oplus W_{n-1}$; reconstruct by XOR. These endpoint schemes are fully anonymous and perfect, with one-bit shares at $t = n$. Replication covers $t = 1$. The rest of this section treats $3 \le t < n$.

4.2 Affine relations and the reconstruction test

A relation records a dependence among values without naming their holders. For an ordered tuple $x = (x_i)_{i \in A}$ with $x_i \in \mathbb{F}_q^D$, define

$$K(x) = \left\{ \lambda \in \mathbb{F}_q^A : \sum_{i \in A} \lambda_i = 0, \quad \sum_{i \in A} \lambda_i x_i = 0 \right\}.$$

The condition on the coefficient sum makes these *affine* relations: adding the same vector to every x_i preserves them. If C_x has columns $(1, x_i)^{\mathsf{T}}$, then $K(x) = \ker C_x$. There exists a relation $\lambda \in K(x)$ with $\lambda_i \neq 0$ if and only if the ith augmented column lies in the span of the remaining columns. Equivalently, x_i is an affine combination of the remaining points. The tuple is affinely independent if $K(x) = \{0\}$.

We identify $\mathbb{F}_q^A$ with the vectors in $\mathbb{F}_q^n$ that vanish outside A, by extending each coefficient vector with zeros. Define

$$H_A = \left\{ \lambda \in \mathbb{F}_q^n : \operatorname{supp}(\lambda) \subseteq A, \ \sum_{i \in A} \lambda_i = 0 \right\}.$$

Then $K(x)$ is a subspace of H_A. For a relation space $R \leq H_{[n]}$, write $\operatorname{supp} R$ for the union of the supports of its vectors. Its restriction to A is $R \cap H_A$: this keeps the relations supported in A, rather than projecting their coefficients. This distinction will let us compare a global choice of R with the relations visible to an unauthorized set.

An invertible affine map has the form $x \mapsto Mx + z$, with $M \in \mathrm{GL}_D(\mathbb{F}_q)$ and $z \in \mathbb{F}_q^D$. These maps form $\mathrm{AGL}_D(\mathbb{F}_q)$ and preserve affine relations.

The two relation patterns. The decoder outputs zero if every augmented column lies in the span of the others, and one if some column does not. Reordering the values permutes the columns and preserves this answer. Affine dependence alone is not enough for output zero: each coordinate must occur with a nonzero coefficient in some relation, and the relation may differ from one coordinate to another. We arrange the two secret cases to produce these different answers.

For zero, the dealer evaluates D independent polynomials of degree at most $k-1$. Every k shares are independent uniform vectors. On $k+1 = t$ shares, polynomial interpolation supplies a relation with a nonzero coefficient at every coordinate. Thus every augmented column lies in the span of the others.

For one, the dealer chooses a relation space R supported on at most k participants. The geometric construction in Section 4.3 gives points whose relation space on each set A of at most $k+1$ participants is exactly $R \cap H_A$. An authorized set has $k+1$ members, so at least one member lies outside $\operatorname{supp} R$. Every relation on this set has zero coefficient at that member. Its augmented column therefore lies outside the span of the others, and the decoder outputs one. The construction needs to control relations only on sets of at most $k+1$, the sizes used for anonymity and reconstruction; Lemma 4 gives this local guarantee.

The remaining task is to hide this difference from unauthorized sets. We choose R so that $R \cap H_A$ has the same distribution as the relations among independent uniform vectors for every unauthorized set A. Tuples with a fixed relation space are exactly the images of one another under invertible affine maps (Lemma 5). A uniform such map makes our tuple uniform among tuples with that relation space. A tuple of independent uniform vectors, conditioned on its relation space, has the same conditional distribution. Matching the relation-space probabilities therefore matches the full distribution of the observed shares. Section 4.3 specifies this law, and Section 4.5 implements it exactly.

4.3 Sharing and reconstruction algorithms

Sharing zero Polynomial evaluations supply the relation pattern for zero. Their degree is one lower than in a labeled (t, n) Shamir scheme because only $t - 1 = k$ shares must be uniform.

Secret-zero sharing. Public parameters are as above.
$\mathsf{Share}(0) \to (W_1, \ldots, W_n) \in (\mathbb{F}_q^D)^n$:
1. Sample $f_1, \ldots, f_D$ independently and uniformly from $\mathbb{F}_q[X]_{\le k-1}$.
2. For each $i \in [n]$, return $W_i = (f_1(\alpha_i), \ldots, f_D(\alpha_i))$.

Localized relation spaces To ensure that every authorized set has a member outside the sampled support, we draw R from

$$\mathcal{P}_k = \{R \le H_{[n]} : |\operatorname{supp} R| \le k\}.$$

We now choose its distribution to match the relations visible below threshold. Fix an unauthorized set A and a subspace $T \le H_A$. For independent uniform vectors on A, each of the D coordinate rows satisfies the $\dim T$ independent equations specified by T with probability $q^{-\dim T}$. Thus the probability that all relations in T hold is $q^{-D \dim T}$. Also $T \subseteq R \cap H_A$ if and only if $T \subseteq R$. Writing $Q = q^D = |\Omega|$, we therefore require, for every $T \in \mathcal{P}_k$,

$$\Pr[T \subseteq R] = Q^{-\dim T}.$$

A containment event sums the masses of all sampled spaces containing T. Consequently the unknown masses must satisfy $\sum_{R \supseteq T} \pi(R) = Q^{-\dim T}$. Möbius inversion solves these equations by subtracting the contributions of larger spaces. For $R \in \mathcal{P}_k$, it gives

$$\pi(R) = \sum_{\substack{S \in \mathcal{P}_k \\ R \subseteq S}} (-1)^{\dim S - \dim R} q^{\binom{\dim S - \dim R}{2}} Q^{-\dim S}. \tag{3}$$

The alternating factor is the Möbius coefficient of a finite subspace lattice. Inversion gives the prescribed containment sums, but positivity also has to be proved. The leading term is $Q^{-\dim R}$; the dimension choice makes the total contribution from proper superspaces smaller in absolute value. Lemma 3 states the resulting probability law. Section 4.5 derives the inversion formula, proves positivity, and samples the law without enumerating $\mathcal{P}_k$.

Lemma 3 (Möbius law). *Suppose $D \ge k$ and*

$$\frac{2\binom{n}{k}}{q^{D-k+2} - 1} < 1. \tag{4}$$

The values in (3) *form a strictly positive probability distribution on $\mathcal{P}_k$ and for every subspace $T \le H_{[n]}$ with $|\operatorname{supp} T| \le k$,*

$$\Pr[T \subseteq R] = Q^{-\dim T}. \tag{5}$$

In particular, these conclusions hold for $D = D_0$.

Sharing one The secret-one dealer turns the sampled space R into points whose local relations are exactly $R \cap H_A$. There are three steps: encode the relations as a matrix kernel, realize the columns as affine combinations of points, and randomize the resulting configuration.

For nonzero R with support $B = \operatorname{supp} R$, a matrix with row space $R^\perp$ on the coordinates B has kernel R. Normalizing a set of independent columns gives the coordinate matrix C used below. Its columns sum to one, so each column describes an affine combination of anchor points. A linear relation among these columns is then an affine relation among the resulting participant points.

The choice of anchors must prevent additional relations on sets of at most $k+1$ participants. We take the anchors and the points outside B from the moment curve

$$\nu_D(x) = (x, x^2, \ldots, x^D) \in \mathbb{F}_q^D.$$

Any $D+1$ distinct curve points are affinely independent, by a Vandermonde determinant. With our choice of D, every relevant local relation expands into a combination of at most $D+1$ distinct curve points. Their independence forces the relation to come from R; Lemma 4 gives the count. Finally, a uniform invertible affine map makes the resulting values uniform conditional on their relation space.

Secret-one sharing.

Public parameters are as above.

$\mathsf{Share}(1) \to (W_1, \ldots, W_n) \in (\mathbb{F}_q^D)^n$:

1. Draw $R \leftarrow \mathsf{LocalRelation}(n, k, q, D)$ using Section 4.5. If $R = 0$, set $v_i = \nu_D(\alpha_i)$ for every i and continue at step 4.
2. Set $B = \operatorname{supp} R$, $\ell = |B|$, $r = \dim R$, and $\delta = \ell - r$. Compute a $\delta \times \ell$ matrix Λ with row space $R^\perp$ on B and first row $(1, \ldots, 1)$. Select δ independent columns, indexed by J, and set
$$C = \Lambda_J^{-1} \Lambda \in \mathbb{F}_q^{\delta \times \ell}.$$
Index the columns of C by the sorted set B.
3. Assign distinct field elements $\beta_1, \ldots, \beta_\delta$ to the anchors and γ_i, $i \notin B$, to the outsiders, using the first $n - r$ public field elements in a fixed order. Set
$$v_i = \begin{cases} \sum_{j=1}^{\delta} C_{j,i} \nu_D(\beta_j), & i \in B, \\ \nu_D(\gamma_i), & i \notin B. \end{cases}$$
4. Sample $D \times D$ matrices uniformly over $\mathbb{F}_q$ until an invertible matrix M is obtained. Independently sample z uniformly from $\mathbb{F}_q^D$. For every $i \in [n]$, return
$$W_i = M v_i + z. \tag{6}$$

The normalization in step 2 preserves the required kernel and makes the coordinates affine. Since $R \le H_B$, the all-ones row belongs to $R^\perp$ and can be completed to the rows of Λ. Left multiplication by Λ_J^{-1} preserves its kernel, so $\ker C = R$ on the coordinates B. The first row of Λ_J is all ones; taking the

first row of $\Lambda_J C = \Lambda$ gives $\mathbf{1}_\delta^\mathsf{T} C = \mathbf{1}_\ell^\mathsf{T}$. Thus the coefficients assigned to each participant sum to one, as required for the affine combinations in step 3.

There are δ anchors and $n - \ell$ outsider points, using $\delta + n - \ell = n - r \le q$ distinct field parameters. Each field draw uses m fair bits. Each participant receives the final vector W_i.

Anonymous reconstruction The reconstruction rule reads the relation pattern directly from the values. No sampled support, coordinate matrix, or participant index is needed.

Reconstruction. $\mathsf{Recon}(x_1, \ldots, x_t) \to \{0, 1\}$, for $x_i \in \mathbb{F}_q^D$:

1. Form $C = \begin{pmatrix} 1 & \cdots & 1 \\ x_1 & \cdots & x_t \end{pmatrix}$.
2. Return zero if every column of C lies in the span of its remaining columns; otherwise return one.

The rule is a total deterministic rank or nullspace test. Permuting the inputs permutes the columns and leaves its answer unchanged.

4.4 Correctness and exact anonymity

Two facts connect the dealer's relation space to the observed shares. The geometric realization preserves exactly the prescribed local relations (Lemma 4), and affine randomization is uniform among all tuples with those relations (Lemma 5). Together with the law of Lemma 3, these facts give the required unauthorized distribution. Positivity of that law and its efficient exact sampler are proved in Section 4.5.

Lemma 4 (Localized relations). *For nonzero R, the vectors in Section 4.3 satisfy, for every $A \subseteq [n]$ with $|A| \le k + 1$,*

$$K((v_i)_{i \in A}) = R \cap H_A \tag{7}$$

whenever

$$D \ge \max\{k, \min\{k + \delta - 1, n - r - 1\}\}, \qquad \delta = \ell - r. \tag{8}$$

For $R = 0$, the same identity holds whenever $D \ge k$. In particular, $D = D_0$ realizes every $R \in \mathcal{P}_k$.

Proof. Write $\eta_D(x) = (1, x, \ldots, x^D)^\mathsf{T}$. Any $a \le D+1$ columns $\eta_D(x_1), \ldots, \eta_D(x_a)$ at distinct parameters are independent: their first a rows have Vandermonde determinant $\prod_{i<j}(x_j - x_i) \ne 0$.

For nonzero R, the columns of C sum to one, so for $i \in B$ the augmented participant column is $\sum_j C_{j,i} \eta_D(\beta_j)$. Suppose A meets B, and let $\lambda \in K((v_i)_{i \in A})$,

extended by zero outside A. Substituting the anchor expressions into this relation gives

$$\sum_{j=1}^{\delta}(C\lambda_B)_j\eta_D(\beta_j)+\sum_{i\in A\setminus B}\lambda_i\eta_D(\gamma_i)=0,$$

where λ_B is extended by zero on $B\setminus A$. At most k outsiders occur because $|A|\le k+1$ and $A\cap B\neq\varnothing$. Thus the number of distinct underlying curve columns is at most $k+\delta$, and also at most the total $n-r$. Condition (8) makes their number at most $D+1$, so they are independent. Every outsider coefficient vanishes and $C\lambda_B=0$, so $\lambda\in R\cap H_A$. Conversely, every vector in $R\cap H_A$ gives a relation because $\ker C=R$.

If A misses B, its points are at most $k+1$ distinct curve points and have no relation for $D\ge k$. Also $R\cap H_A=0$, giving the same identity. When $R=0$, all participant points lie on the curve, so this argument applies to every A. Finally, for nonzero R we have $\delta\le k-1$ and $r\ge 1$, so $\min\{k+\delta-1,n-r-1\}\le\min\{2k-2,n-2\}=D_0$. Also $D_0\ge k$ because $2\le k\le n-2$. This proves the uniform claim. □

The preceding lemma fixes the local relations. The next shows that affine randomization makes the local values uniform conditional on those relations.

Lemma 5 (Affine orbits). *Two ordered tuples in $(\mathbb{F}_q^D)^a$ have the same affine-relation space if and only if one is obtained from the other by an element of $\mathrm{AGL}_D(\mathbb{F}_q)$. Consequently, uniform affine randomization is uniform on the set of tuples with a fixed affine-relation space.*

Proof. If x and y have the same relation space, define on the affine span of x,

$$\phi\left(\sum_i c_ix_i\right)=\sum_i c_iy_i,\qquad\sum_i c_i=1.$$

Two affine representations differ by an element of the common relation space, so ϕ is well-defined. Interchanging x and y gives a well-defined inverse. Hence its linear part is an isomorphism between the translation spaces. Extending bases of these spaces extends ϕ to an element of $\mathrm{AGL}_D(\mathbb{F}_q)$ mapping every x_i to y_i. Conversely, an invertible affine map preserves exactly the affine relations.

Thus the tuples with a given relation space form one orbit. For a fixed tuple, the number of affine maps taking it to any prescribed point in that orbit is the size of its stabilizer. These numbers are equal, so a uniform affine map induces the uniform distribution on the orbit. □

Proof (of Theorem 2).

Secret zero. The $k\times k$ evaluation matrix at any k distinct public points is Vandermonde and invertible. Hence any k evaluations of a uniformly sampled polynomial of degree at most $k-1$ are independent uniform field elements. Applying this to the D independent polynomials proves uniformity of every

unauthorized tuple. On any $k+1$-tuple of distinct evaluation points, the Lagrange relation

$$\lambda_j = \left(\prod_{i \neq j} (\alpha_j - \alpha_i) \right)^{-1}$$

is nonzero in every coordinate and annihilates every polynomial of degree at most $k-1$, including the constant polynomial. To see the identity directly, write any polynomial of degree at most $k-1$ in its Lagrange interpolation form on $k+1$ points and compare the coefficient of X^k; that coefficient on the left is zero and on the right is $\sum_j \lambda_j f(\alpha_j)$. Here the α_j denote the selected evaluation points. Hence every augmented column lies in the span of the others and reconstruction outputs zero.

Reconstruction for secret one. Let A be any $k+1$ participants. Since $|\operatorname{supp} R| \leq k$, choose $i \in A \setminus \operatorname{supp} R$. By (7), every relation in $K((v_j)_{j \in A})$ has zero ith coordinate. Affine transformations preserve relations, so the same holds for the shares. The ith column is not in the span of the others, and reconstruction outputs one on every support point.

Anonymity for secret one. Fix A with $|A| \leq k$. If X_i are independent uniform points in $\mathbb{F}_q^D$, then for every $T \leq H_A$ of dimension r,

$$\Pr[T \subseteq K((X_i)_{i \in A})] = q^{-Dr} = Q^{-r}. \tag{9}$$

Indeed, each of the D coordinate rows must satisfy r independent linear equations. Lemma 3 gives the same containment moments for $R \cap H_A$. Since the containment transform on the finite subspace lattice is invertible, the two relation spaces have the same distribution.

It remains to compare the values conditional on their relation spaces. For each fixed dealer choice of R, Lemma 4 gives relation space $R \cap H_A$ on A. Uniform M, z then make the tuple uniform on that affine orbit by Lemma 5. This conditional distribution depends on R only through $R \cap H_A$, so it is unchanged when we condition only on the local relation space.

The independent uniform tuple X_A is likewise uniform among all tuples with its given relation space. The two tuples therefore have both the same relation-space law and the same conditional law of values. Their full laws agree, giving independent uniform shares on A for secret one.

Length and running time. The two secret cases have exact uniform unauthorized tuples, and reconstruction is correct on every point in each support. Its rank test is total and invariant under input permutations, as required by Definition 1. The share alphabet has size q^D, giving Dm bits. The exact sampler is expected polynomial time by Lemma 6; the remaining operations and their costs are analyzed in Section 5.1. □

4.5 Exact sampling of the relation space

The law in (3) assigns a probability to each space in $\mathcal{P}_k$. An efficient sampler uses the fact that this probability depends only on the type $(\ell, r) = (|\operatorname{supp} R|, \dim R)$,

as proved in Lemma 6. There are only $O(k^2)$ nonzero types. The dealer chooses a type with its aggregate probability, chooses a support of size ℓ uniformly, and then chooses a uniform r-subspace with exactly that support. The zero space is a separate outcome.

Let $G(a, j) = \begin{bmatrix} a \\ j \end{bmatrix}_q$ be the Gaussian binomial coefficient, with $G(a, j) = 0$ unless $0 \le j \le a$, including when $a < 0$. It counts the j-dimensional subspaces of $\mathbb{F}_q^a$; Section 4.5 gives its product formula and recurrence. For a fixed ℓ-set B, the number of r-subspaces of H_B with support exactly B is

$$F(\ell, r) = \sum_{h=0}^{\ell-1} (-1)^h \binom{\ell}{h} G(\ell - h - 1, r), \tag{10}$$

for $1 \le r < \ell \le k$. These are the full-support subspaces used by the sampler. To represent their probabilities by integers, put

$$P_0 = 1, \qquad P_j = \prod_{i=0}^{j-1} (Q - q^i), \qquad Z = Q^{k-1}.$$

Every space in $\mathcal{P}_k$ has dimension at most $k-1$, so the denominators in (3) divide Z. Since q is a power of two, Z is also a power of two. Thus integer type weights summing to Z allow an exact draw using $\log_2 Z$ fair bits. We distinguish the weight of one space from the weight of its whole type. For any fixed space of type (ℓ, r), its probability is $p_{\ell,r} = w_{\ell,r}/Z$, where

$$w_{\ell,r} = \sum_{c=\ell}^{k} (-1)^{k-c} \binom{n-\ell}{c-\ell} \binom{n-c-1}{k-c} \cdot Q^{k-c} P_{c-r-1}. \tag{11}$$

There are $\binom{n}{\ell} F(\ell, r)$ spaces of this type, so its aggregate integer weight is

$$\omega_{\ell,r} = \binom{n}{\ell} F(\ell, r) w_{\ell,r}.$$

The zero space receives the remaining weight:

$$\omega_0 = Z - \sum_{\ell=2}^{k} \sum_{r=1}^{\ell-1} \omega_{\ell,r}, \qquad p_0 = \omega_0 / Z. \tag{12}$$

Preprocessing stores the aggregate weights and their cumulative sums in a fixed order, omitting types with $F(\ell, r) = 0$. Lemmas 3 and 6 prove that the retained weights are positive integers summing to Z. The sampler therefore chooses the type by a uniform integer draw and implements the two conditional choices as follows.

LocalRelation(n, k, q, D)

Precomputed input: cumulative type weights at $Q = q^D$, with total $Z = Q^{k-1}$.

1. Draw $U \in \{0, \dots, Z-1\}$ using $\log_2 Z$ fair bits. Locate the unique cumulative-ω interval containing U by binary search. If its type is zero, return $R = 0$.
2. For its type (ℓ, r), initialize the array $(1, \dots, n)$. For $j = 1, \dots, \ell$, draw an integer uniformly from $\{j, \dots, n\}$ and swap that position with position j. Let B be the first ℓ entries, sorted. An integer in a range of size a is sampled with $\lceil \log_2 a \rceil$ bits, rejecting values at least a; a singleton range needs no bits.
3. Draw an $r \times (\ell - 1)$ matrix Y uniformly over $\mathbb{F}_q$. Append to each row its negative row sum, obtaining an $r \times \ell$ matrix A. If $\operatorname{rank} A < r$ or a column of A is zero, repeat this step.
4. Embed the rows of A in the coordinates B, with zeros elsewhere, and return their row span, stored by the support list and this local basis.

The last representation avoids constructing an n-column matrix for the small relation space. Sorting the support changes only its coordinate order.

Subspace counts and inversion We first derive the subspace identities needed for the law and its type weights. The number of j-subspaces of $\mathbb{F}_q^a$ is

$$G(a, j) = \prod_{i=0}^{j-1} \frac{q^a - q^i}{q^j - q^i}.$$

The numerator counts ordered independent j-tuples and the denominator counts bases of a fixed j-space. For $0 \le j \le a$, fixing a hyperplane gives the recurrence

$$G(a, j) = G(a-1, j) + q^{a-j} G(a-1, j-1), \qquad G(a, 0) = 1.$$

Indeed, a subspace outside the hyperplane meets it in dimension $j - 1$, and has q^{a-j} extensions for each such intersection. Induction using this recurrence proves the finite polynomial identity

$$\sum_{j=0}^{a} G(a, j)(-1)^j q^{j(j-1)/2} z^j = \prod_{i=0}^{a-1} (1 - q^i z).$$

At $z = 1$ this is zero for $a > 0$ and one for $a = 0$. Thus on a full subspace-lattice interval $[U, V]$, the coefficients $\mu(U, V) = (-1)^j q^{j(j-1)/2}$, $j = \dim V - \dim U$, invert the containment sums: $\sum_{U \le S \le V} \mu(U, S) = \mathbf{1}_{U=V}$. All intervals in $\mathcal{P}_k$ are full intervals, because every subspace of a space supported on at most k coordinates has that property. This proves the inversion formula used in (3).

The same product formula gives an upper bound for the positivity argument. For $q \ge 4$,

$$G(a, j) \le q^{j(a-j)} \prod_{i=1}^{j} (1 - q^{-i})^{-1} \le \frac{q^{j(a-j)}}{1 - \sum_{i \ge 1} q^{-i}} < 2q^{j(a-j)}.$$

Here the elementary inequality $\prod(1 - x_i) \ge 1 - \sum x_i$ applies to finite products; their positive lower bound persists in the limit. □

Positivity of the localized law

Proof (of Lemma 3). Containment probabilities. By the preceding inversion identity, summing (3) over $R \supseteq T$ and interchanging the finite sums cancels every term indexed by $S \neq T$. Hence

$$\sum_{R \supseteq T} \pi(R) = Q^{-\dim T}.$$

Taking $T = 0$ gives normalization.

Positive masses. For R of dimension r, its own term in (3) is Q^{-r}. We show that the absolute contribution of all proper superspaces is smaller than this term. Let $N_j(R)$ be the number of superspaces in $\mathcal{P}_k$ of dimension $r + j$. For $j \geq 1$, every such superspace lies in H_C for some k-set C, so

$$N_j(R) \leq \binom{n}{k} G(k-1-r, j) \leq 2\binom{n}{k} q^{j(k-1-r-j)}.$$

For the first inequality, each superspace has a support contained in a k-set C, and the spaces containing R inside H_C correspond to j-subspaces of H_C/R. Counting a space more than once only increases the bound. Consequently,

$$\begin{aligned}\left|\frac{\pi(R)}{Q^{-r}} - 1\right| &\leq 2\binom{n}{k} \sum_{j \geq 1} q^{-j(D-k+r+2)-j(j-1)/2} \\ &\leq \frac{2\binom{n}{k}}{q^{D-k+2}-1} < 1.\end{aligned}$$

The second inequality drops the nonnegative terms $jr + j(j-1)/2$ and sums a geometric series. Hence every mass is positive under (4).

The stated dimension. For $D = D_0$, put $u = \min\{k, n-k\} \geq 2$. Then $D_0 - k + 2 = u$ and $\binom{n}{k} = \binom{n}{u}$. If $u = 2$,

$$2\binom{n}{2} = n(n-1) < n^2 - 1 \leq q^2 - 1.$$

If $u \geq 3$, then $2\binom{n}{u} \leq 2n^u/u! \leq n^u/3 < n^u - 1 \leq q^u - 1$. Thus (4) holds at D_0 as well. □

Exact sampling and preprocessing

Lemma 6 (Exact type sampler). *Equations* (10), (11), *and* (12) *give the aggregate masses of the types and the rejection sampler in Section 4.5 returns a uniform full-support r-subspace conditional on its type. Its expected number of trials is at most an absolute constant. For $D = O(k)$, all type weights can be computed with $O(k^3)$ integer operations on $O(k^2 m)$-bit operands and $O(k^2)$ integer storage.*

Proof. Type masses. Equation (10) is inclusion-exclusion over the coordinates omitted by a subspace of H_B. The omitted $h = \ell$ term is zero because $r \geq 1$. To prove (11), put

$$A_j = P_j/Q^j = \prod_{i=0}^{j-1}(1 - q^i/Q).$$

The polynomial identity in Section 4.5 gives

$$\sum_{j=0}^{a}(-1)^j q^{\binom{j}{2}} G(a,j)Q^{-j} = A_a.$$

Fix a full-support $R \leq H_B$ of dimension r, where $|B| = \ell$. To evaluate (3), group its superspaces by their exact support. On a fixed containing support of size $\ell + u$, inclusion-exclude the newly used coordinates. Quotienting by R leaves ambient dimension $\ell + u - h - 1 - r$ when h new coordinates are omitted. Summing the weighted subspaces in that quotient therefore gives

$$\pi(R) = Q^{-r}\sum_{u=0}^{k-\ell}\binom{n-\ell}{u}\sum_{h=0}^{u}(-1)^h\binom{u}{h}A_{\ell+u-h-1-r}.$$

Set $j = u - h$. The coefficient of $A_{\ell+j-1-r}$ is

$$\binom{n-\ell}{j}\sum_{v=0}^{k-\ell-j}(-1)^v\binom{n-\ell-j}{v}$$
$$= (-1)^{k-\ell-j}\binom{n-\ell}{j}\binom{n-\ell-j-1}{k-\ell-j}.$$

The last equality follows by telescoping Pascal's identity; its truncation index is smaller than $n - \ell - j$. Multiply by Z, set $c = \ell + j$, and use $ZQ^{-r}A_{c-r-1} = Q^{k-c}P_{c-r-1}$ to obtain (11). This proves that the per-space probability depends only on (ℓ, r). Multiplying by $\binom{n}{\ell}F(\ell,r)$ and normalizing gives the aggregate weights and (12).

Uniform sampling within a type. For a fixed support B, appending a negative row sum maps a uniform row in $\mathbb{F}_q^{\ell-1}$ bijectively to a uniform vector in H_B. Conditioning on full rank and nonzero columns therefore gives a uniform ordered basis among all full-support r-subspaces of H_B. Each such subspace has the same number $|\operatorname{GL}_r(\mathbb{F}_q)|$ of ordered bases, so its conditional probability is $1/F(\ell,r)$.

The partial shuffle chooses each ℓ-element support with probability $1/\binom{n}{\ell}$. Combining the type, support, and subspace choices gives each nonzero R of type (ℓ, r) probability

$$\frac{\omega_{\ell,r}}{Z}\,\frac{1}{\binom{n}{\ell}}\,\frac{1}{F(\ell,r)} = \frac{w_{\ell,r}}{Z} = \pi(R).$$

The zero outcome has probability $\omega_0/Z = \pi(0)$, completing the distributional argument.

Acceptance probability. Assume first $r \geq 2$. Rank failure is at most

$$\sum_{i=0}^{r-1} q^{i-(\ell-1)} \leq \frac{1}{q-1},$$

and all r vectors have zero in a fixed coordinate with probability q^{-r}. Since $\ell \leq q-1$, the union bound gives total failure at most $1/(q-1) + (q-1)/q^r \leq 1/3 + 1/4 < 1$.

For $r = 1$, the number of vectors in H_B with every coordinate nonzero is

$$N_\ell = \frac{(q-1)^\ell + (q-1)(-1)^\ell}{q}.$$

These counts satisfy $N_1 = 0$ and $N_\ell = (q-1)^{\ell-1} - N_{\ell-1}$: choose the first $\ell - 1$ nonzero entries and require their sum to be nonzero before setting the last entry. Solving this recurrence gives the formula. If ℓ is even, $N_\ell/q^{\ell-1} \geq (1-1/q)^{q-1} \geq 1/3$. If ℓ is odd, then $\ell \geq 3$ and

$$N_\ell/q^{\ell-1} \geq 1/3 - (q-1)/q^3 > 1/4.$$

The bound $(1-1/q)^{q-1} > 1/3$ follows by reciprocation from $(1+1/(q-1))^{q-1} < \sum_{j\geq 0} 1/j! < 3$. Thus the local-basis loop has acceptance at least $1/4$ and expected at most four trials. The fixed-width integer-range loop accepts with probability greater than $1/2$.

Preprocessing. The displayed weight formulas depend only on (n, k, q, D) and are computed once for these parameters. Section 5.1 derives their cubic integer-operation count, operand sizes, and storage bound. □

5 Efficiency, refinements, and limits

The constructions use two kinds of computation: public preprocessing for fixed parameters and online computation for each sharing or reconstruction. We account for both, then show how sparse affine coordinates can reduce the one-bit share dimension. Finally, known alphabet bounds place the achieved share lengths in context. The share-length comparison appears in Table 1.

5.1 Complexity of the one-bit construction

We count bit operations using schoolbook arithmetic and fair-bit randomness, as in Section 2. Put $m = \lceil \log_2 n \rceil$ and $k = t - 1$. For $3 \leq t < n$, after public preprocessing, the one-bit construction satisfies

$$\mathbb{E}[T_{\mathsf{Share}}] = O((nt^2 + t^3)m^2), \qquad T_{\mathsf{Recon}} = O(t^3 m^2).$$

The field representation costs $O(nm^3)$ bit operations to construct, and the probability table costs $O(t^7 m^2)$ bit operations and occupies $O(k^4 m)$ bits. Both are reused across sharing executions. The equality and XOR endpoint schemes of Section 4.1 require no probability table. The following accounting also applies to the sparse realization in Section 5.3; its dimension search adds a separate preprocessing cost.

Field representation. Given a monic irreducible polynomial g of degree m, use $\mathbb{F}_q = \mathbb{F}_2[X]/(g)$ with m-bit coefficient strings. Addition costs $O(m)$ bit operations. Multiplication with reduction and inversion by the polynomial extended Euclidean algorithm each cost $O(m^2)$ bit operations. For inversion, the degrees decrease, and the classical division and coefficient-update costs sum to a quadratic bound.

A deterministic setup tests at most $2^m = q < 2n$ monic degree-m candidates. For each candidate compute $X^{2^j} \bmod g$ for $j = 1, \ldots, m$, require $X^{2^m} = X \bmod g$, and test $\gcd(g, X^{2^j} - X) = 1$ for $1 \le j \le \lfloor m/2 \rfloor$. The equality forces squarefreeness and factor degrees dividing m; a reducible polynomial has a factor of degree at most $m/2$ and fails a gcd test. Thus the conditions characterize irreducibility. An irreducible candidate exists, and each test costs $O(m^3)$, giving $O(nm^3)$ bit operations in total, including public point enumeration. For $m = 1$, use $g = X$.

Cubic probability preprocessing. The probability table stores cumulative weights for the relation types in Section 4.5. It is built once for fixed (n, t, D), using integers of $\ell_{\text{int}} = O(k^2 m)$ bits. Gaussian coefficients use the recurrence in Section 4.5; binomial coefficients are tabulated only for the needed offsets of size at most k. For example, $\binom{a}{j+1} = \binom{a}{j}(a-j)/(j+1)$ computes them by exact integer divisions without a table of n rows. Powers of q and Q are binary shifts. The product sequence P_j and full-support counts $F(\ell, r)$ take $O(k^3)$ integer operations. There are

$$\sum_{\ell=2}^{k} \sum_{r=1}^{\ell-1} (k - \ell + 1) = \binom{k+1}{3}$$

summands in the per-space numerators. Aggregate the type numerators and compute the zero-space numerator by normalization. Every coefficient and intermediate absolute sum has $O(k^2 m)$ bits: Gaussian exponents are $O(k^2)$, support choices contribute at most n^k, and the common denominator $Z = Q^{k-1}$ has $D(k-1)m = O(k^2 m)$ bits. Thus preprocessing costs $O(k^3 \ell_{\text{int}}^2) = O(t^7 m^2)$ bit operations. Storing $O(k^2)$ cumulative weights takes $O(k^4 m)$ bits; temporary tables may be discarded.

The optional dimension search of Corollary 2 tests at most $O(k)$ candidate dimensions. Each test has the same cubic arithmetic cost, so the complete search costs $O(k^4)$ integer operations or $O(k^8 m^2)$ bit operations. Its online sharing and reconstruction bounds are unchanged. Thus the cubic preprocessing bound is per dimension, and the additional factor accounts for testing all candidate dimensions.

Online integer sampling. The type draw uses $\log_2 Z$ fair bits, followed by binary search among $O(k^2)$ cumulative weights. Support selection can use a partial Fisher–Yates shuffle, with $O(n)$ initialization and at most k exact range draws. Sorting the selected indices costs $O(k \log(k+1))$ comparisons of m-bit integers. Including exact range rejection, these steps have expected bit cost

$$O\big(\ell_{\text{int}} \log(k+1) + nm + km^2 + k \log(k+1)m\big).$$

Since $k < n$ and $\log(k+1) = O(m)$, this is absorbed by the online bounds above. The accounting therefore includes both the number of integer operations and the lengths of their operands.

Online field arithmetic. Secret zero uses Dk field draws and Horner evaluation of D polynomials at n points, costing $O(nk^2)$ field operations. For secret one, local basis sampling and row reduction cost expected $O(k^3)$. The affine-basis substitution requires one matrix inversion and $O(k^3)$ local operations; computing the curve points and assigning outsiders costs $O(nk)$ operations. The sparse realization, when used, also fits $O(k^3 + nk^2)$: its kernel basis, normalizing polynomial, and change of coordinates have dimensions $O(k)$. A uniformly sampled invertible $D \times D$ matrix costs expected $O(k^3)$ operations by rank rejection, and applying it to every point costs $O(nk^2)$. For $q \geq 4$, the acceptance probability is at least $1 - \sum_{j \geq 1} q^{-j} \geq 2/3$.

Multiplying these field-operation counts by $O(m^2)$ gives the stated sharing bound. Reconstruction eliminates the $(D+1) \times t$ augmented matrix and checks the supports of a nullspace basis, giving $O(t^3 m^2)$ bit operations. Sampling field elements uses $O(k^2 m)$ expected fresh random bits in addition to the type and support draws, and writing all shares takes nDm bits. These costs are included in the online bounds.

5.2 Compiler costs

The compiler adds polynomial evaluation and linear algebra to ah calls to the one-bit primitive. Recall the parameters from Section 3.1: $k = t - 1$, $d = \lceil L/a \rceil$, $h = \min\{k, d\}$, and share length $M = ad + bah$, where b is the auxiliary one-bit share length. Let T_{share} and T_{recon} denote the bit costs of one auxiliary invocation after public preprocessing. Including input and output, the compiler satisfies

$$\mathbb{E}[T'_{\text{share}}] = ah\,\mathbb{E}[T_{\text{share}}] + O(nkd\,a^2 + nM),$$
$$T'_{\text{recon}} = ah\,T_{\text{recon}} + O(tdh\,a^2 + tM).$$

The first term in each bound accounts for protecting or recovering the ah auxiliary bits. For the remaining work, Horner evaluation of the vector polynomial takes $O(nkd)$ operations in F. Row reduction of the $k \times d$ matrix with rows $A_1, \ldots, A_{k-1}, S$ takes $O(kdh)$ operations, and reconstruction takes $O(tdh)$. Each field operation costs $O(a^2)$ bit operations. Since $kdh \leq nkd$, the dealer's row reduction is absorbed by its polynomial-evaluation bound.

The linear terms nM and tM cover writing and reading shares, field encodings, validation, and sampling the kda mask bits. The mask and row-reduction workspace occupies $O(kda)$ bits, in addition to the outputs and auxiliary workspace. At the default width $a = m$, the field degree is independent of L.

Instantiating the auxiliary primitive with Theorem 2 gives, for $3 \leq t < n$,

$$\mathbb{E}[T'_{\text{share}}] = O\big(ah(nt^2 + t^3)m^2 + nkd\,a^2 + nM\big),$$
$$T'_{\text{recon}} = O\big(ah\,t^3 m^2 + tdh\,a^2 + tM\big).$$

All auxiliary invocations reuse one type table, so its $O(t^7m^2)$ preprocessing cost is paid once. Using the dimension search in Corollary 2 changes this one-time cost as stated there. Field setup is also performed once for each required degree. At degree a, the enumeration method of Section 5.1 costs $O(2^a a^3 + na)$ bit operations for the payload field and its public points. With the default $a = m$, the payload and auxiliary schemes share the same field representation: field setup costs $O(nm^3)$ bit operations, and the field description and public points occupy $O(nm)$ bits.

5.3 A refinement by sparse affine coordinates

Return to the one-bit parameters of Section 4.1. Write $\ell = |\operatorname{supp} R|$, $r = \dim R$, and $\delta = \ell - r$ for a nonzero R. The dimension D_0 accounts for every curve anchor that a local relation might involve. In the realization of Section 4.3, one participant may use all δ anchors. If each participant instead uses at most $r + 1$, a small set of participants can involve fewer anchors, and a lower-dimensional curve can still keep those anchors independent. The lemma below constructs coordinates with this property.

Two further steps turn these coordinates into a sharing scheme. First, we must make them affine while preserving their sparsity. Second, a smaller D must still give positive masses in (3), evaluated at $Q = q^D$. We address the affine normalization after the lemma and give the dimension selection rule in Corollary 2.

Lemma 7 (A kernel basis with small column load). *Let $H \in \mathbb{F}_q^{r\times\ell}$ have full row rank and no zero column. There is a matrix $C \in \mathbb{F}_q^{(\ell-r)\times\ell}$ whose rows form a basis of $\ker H$ and whose every column has at most $r + 1$ nonzero entries. It can be computed with $O(\ell r^2)$ field operations.*

Proof. Basis construction. Process the columns of H in their index order. Maintain an independent set I of active columns spanning all columns processed so far. Add a new independent column to I. If column j is dependent, write $H_j = \sum_{i\in I} c_i H_i$, emit the vector

$$e_j - \sum_{i\in I} c_i e_i \in \ker H,$$

and replace the oldest active index with $c_i \neq 0$ by j. Such an index exists because $H_j \neq 0$. The replacement preserves the active span and independence. Exactly $\ell - r$ vectors are emitted. Their largest nonzero coordinates are their distinct insertion indices j, so they are independent and form a basis of $\ker H$.

Column sparsity. Fix a column index i. Its own insertion emits at most one vector containing i. When i is active, any later emitted vector containing it either evicts i or evicts an active index older than i. There are at most $r - 1$ such older active indices at insertion, and subsequent insertions cannot create an older one. Thus at most r later vectors contain i. After eviction it never appears again. The column load is at most $r + 1$.

Running time. Maintain an invertible $r \times r$ matrix U whose first $|I|$ columns are the active columns, with an independent completion, together with U^{-1}. Initially U is the identity. Coordinates of a new column are computed by $U^{-1}H_j$ in $O(r^2)$ operations. For an independent insertion, replace a completion column whose coordinate is nonzero; for a dependent insertion, replace the active column selected above. Both updates replace one column of U by H_j, with a nonzero replacement coordinate v_p, where $v = U^{-1}H_j$. The new inverse is obtained by dividing row p of U^{-1} by v_p and subtracting v_h/v_p times the old row p from each other row h. Reordering columns to keep the active columns first only reorders inverse rows. Each update costs $O(r^2)$, proving the claim. □

To apply the lemma, write $B = \operatorname{supp} R$ and restrict the coordinates of R to B. All kernels and orthogonal complements in this paragraph are taken in $\mathbb{F}_q^B$. A row-basis matrix H for R has no zero column, since every index in B occurs in a relation. Lemma 7 gives a $\delta \times \ell$ matrix C with

$$\delta = \ell - r, \qquad \operatorname{rowspan} C = R^{\perp}, \qquad \ker C = R.$$

Since $R \le H_B$ and C has full row rank, there is a unique $a \in \mathbb{F}_q^{\delta}$ satisfying $a^{\mathsf{T}}C = \mathbf{1}_{\ell}^{\mathsf{T}}$. The vector a is nonzero, but some entries may vanish. Scaling row j of C by a_j would make the column sums one, but it would erase that row whenever $a_j = 0$.

Instead, we choose augmented anchors U_j whose first coordinates are a_j. The supported column $V_i = \sum_j C_{j,i}U_j$ then has first coordinate $(a^{\mathsf{T}}C)_i = 1$, so it represents an affine point without changing C. To obtain these anchors from the curve, choose a polynomial g vanishing exactly at the parameters β_j for which $a_j = 0$. A change of coordinates makes $g(x)$ the first coordinate of the curve column at x. Scaling the remaining anchors sets their first coordinates to the required nonzero a_j; outsider columns are scaled to have first coordinate one. Invertible coordinate changes and nonzero scalings preserve the independence needed for the local-relation argument.

Sparse realization.

Input: nonzero $R \in \mathcal{P}_k$, with $B = \operatorname{supp} R$, $\ell = |B|$, $r = \dim R$, $\delta = \ell - r$, and a dimension $D \ge k$ satisfying (13) below. Output: $v_1, \dots, v_n \in \mathbb{F}_q^D$.

1. Compute C by Lemma 7, and solve $a^{\mathsf{T}}C = \mathbf{1}_{\ell}^{\mathsf{T}}$.
2. Assign distinct parameters $\beta_1, \dots, \beta_\delta$ and γ_i, $i \notin B$, from $\mathbb{F}_q$ in a fixed order. Set

$$g(X) = \prod_{j : a_j = 0} (X - \beta_j), \qquad \eta_D(x) = (1, x, \dots, x^D)^{\mathsf{T}}.$$

3. Let $T \in \mathrm{GL}_{D+1}(\mathbb{F}_q)$ have first row equal to the coefficient vector of g padded to length $D + 1$. For its remaining rows take all standard basis rows except one at a nonzero coefficient of g, in a fixed order.

Sparse realization (continued).

4. For $1 \le j \le \delta$, set

$$\sigma_j = \begin{cases} a_j / g(\beta_j), & a_j \neq 0, \\ 1, & a_j = 0, \end{cases} \qquad U_j = \sigma_j T \eta_D(\beta_j).$$

5. Set

$$V_i = \begin{cases} \sum_{j=1}^{\delta} C_{j,i} U_j, & i \in B, \\ T \eta_D(\gamma_i) / g(\gamma_i), & i \notin B. \end{cases}$$

Return v_i by deleting the first coordinate of V_i.

The next lemma verifies the normalization and gives the dimension needed to preserve local relations. Its two alternatives count either all curve columns used by the realization or only those involved in one local set.

Lemma 8 (Sparse local relations). *The sparse realization satisfies* (7) *for every* $|A| \le k+1$ *provided*

$$D \ge \max\left\{k, \min\left\{n - r - 1, k + \delta - \left\lceil \frac{\delta}{r+1} \right\rceil\right\}\right\}. \tag{13}$$

If $D \le D_0$, *it uses* $O(nk^2 + k^3)$ *field operations.*

Proof. Affine normalization. Because $a \neq 0$, the degree of g is at most $\delta - 1 \le D$. The prescribed coefficient row is nonzero, and adjoining all standard basis rows except one at a nonzero coefficient makes T invertible. Distinctness of the parameters gives $g(\beta_j) \neq 0$ whenever $a_j \neq 0$, and $g(\gamma_i) \neq 0$ for every outsider. Hence every σ_j is nonzero. The first coordinate of U_j is a_j in both cases, including $a_j = 0$. Thus the first coordinate of each V_i on B is $(a^{\mathsf{T}} C)_i = 1$, and that of each outsider column is one as well. These are the augmented columns of the returned points.

Local relations. Let A contain a_0 participants in B, where $|A| \le k+1$. If $a_0 = 0$, its at most $k+1$ columns are independent for $D \ge k$, since they are nonzero scalings of distinct transformed curve columns. If $a_0 \ge 1$, its supported columns use at most $\min\{\delta, (r+1)a_0\}$ anchors. Expanding an augmented relation therefore uses at most

$$k + 1 - a_0 + \min\{\delta, (r+1)a_0\} \le k + 1 + \delta - \left\lceil \frac{\delta}{r+1} \right\rceil$$

underlying curve columns, and also at most $n - r$ in total. To justify the inequality, write $s = r+1$, $z = \lceil \delta / s \rceil$, and $\delta = s(z-1) + v$ with $1 \le v \le s$. If $a_0 \ge z$, use $\min\{\delta, sa_0\} - a_0 \le \delta - z$. Otherwise use $\min\{\delta, sa_0\} - a_0 \le (s-1)(z-1) \le \delta - z$.

Condition (13) makes the smaller count at most $D+1$, so the underlying curve columns are independent. Expanding a relation with coefficient vector $\lambda \in H_A$ and applying T^{-1} makes every outsider coefficient zero and gives $\sigma_j (C\lambda_B)_j = 0$

at every anchor. Since $\sigma_j \neq 0$, $C\lambda_B = 0$. Since $\ker C = R$ on B and all outsider coefficients vanish, extending by zeros gives $\lambda \in R \cap H_A$. Conversely every vector in $R \cap H_A$ vanishes under C and supplies the required relation. This proves (7).

Running time. Assume $D \leq D_0 \leq 2k-2$. The kernel basis, solving for a, and forming the supported combinations take $O(k^3)$ field operations. Evaluating the curve and applying the $O(k) \times O(k)$ matrix T at $O(n)$ parameters takes $O(nk^2)$. This gives the stated cost with ordinary matrix arithmetic. □

The dealer must choose one dimension that works for every sampled relation space. Maximizing the geometric requirement over all possible types gives

$$g(n,k) = \max\left\{k, \max_{1\leq r\leq k-1} \min\left\{n-r-1, 2k-r-\left\lceil\frac{k-r}{r+1}\right\rceil\right\}\right\}. \quad (14)$$

For fixed r, the expression $\delta - \lceil \delta/(r+1) \rceil$ is nondecreasing in δ. Since $\delta = \ell - r \leq k - r$, Lemma 8 applies to every nonzero R whenever $D \geq g(n,k)$. Ordinary curve points handle $R = 0$ in the same dimension.

This geometric threshold is the starting point for dimension selection. At each candidate D, we also recompute the exact probability weights. The sufficient positivity estimate (4) supplies an upper endpoint where the search must succeed; testing the weights directly can allow an earlier stopping point.

Corollary 2 (Dimension selection for the sparse realization). *Let e be the smallest integer such that $q^e > 2\binom{n}{k} + 1$, and put*

$$D_{\text{safe}} = \max\{g(n,k), k-2+e\}.$$

For each integer $D = g(n,k), \ldots, D_{\text{safe}}$, recompute the weights in (11)–(12) using $Q = q^D$. Select the first D for which $w_{\ell,r} > 0$ for every type with $F(\ell,r) > 0$ and $\omega_0 > 0$. This rule terminates and yields a fully anonymous perfect one-bit scheme with Dm-bit shares. It has the same online asymptotic costs as Theorem 2. The entire dimension selection uses $O(k^4)$ integer operations on $O(k^2 m)$-bit operands, or $O(k^8 m^2)$ bit operations with schoolbook arithmetic, and $O(k^2)$ integer storage. Moreover, $g(n,k) \leq D_{\text{safe}} \leq D_0$.

Proof. Termination. The definition of D_{safe} supplies both the geometric requirement and (4), so Lemma 3 guarantees success there. To see the upper bound, each inner term in (14) is at most both $n-2$ and $2k-2$, and $k \leq D_0$; hence $g(n,k) \leq D_0$. The positivity proof at D_0 gives $q^{D_0-k+2} > 2\binom{n}{k} + 1$, so $e \leq D_0 - k + 2$ and $D_{\text{safe}} \leq D_0$.

Sharing guarantees. At any earlier successful dimension, the exact weight tests themselves prove positivity for every realized type and the zero space. The counting identities and normalization prove that they are precisely the Möbius law at that Q. Its containment moments follow by inversion, independently of the sufficient estimate (4). Use the sparse realization in place of steps 2–3 of secret-one sharing and keep the zero-sharing and reconstruction algorithms. The local relation identity, affine-orbit argument, and correctness proof of Section 4.4 apply with this realization in place of Lemma 4.

Complexity. There are at most $D_0 - k + 1 = O(k)$ candidate dimensions. Each trial takes $O(k^3)$ integer operations by Lemma 6; its operand bound is uniform because $D \le D_0 = O(k)$. Reuse the dimension-independent binomial and Gaussian tables and discard unsuccessful weights, giving the stated time and storage costs. Lemma 8 and affine randomization give the same online field-operation bound. □

The selected dimension also reduces the compiler's auxiliary cost: use $b = Dm$ in Section 5.2. The payload length is unaffected, and the same compiler algorithms and cost formulas apply.

5.4 Known alphabet bounds

Anonymous reconstruction imposes a lower bound on the share alphabet even when secrecy is required without product uniformity. Kishimoto, Okada, Kurosawa, and Ogata [12] prove that every finite anonymous perfect threshold scheme satisfies

$$|\Omega| \ge (|\Sigma| - 1)(n - t + 1) + 1 \quad (1 < t < n), \tag{15}$$

with strict inequality for $3 \le t < n$ and $|\Sigma| \ge 2$. For $|\Sigma| \ge 2$, the bound forces $|\Omega| > |\Sigma|$, ruling out rate one at $1 < t < n$. This impossibility is also known from Phillips and Phillips [14].

To apply the bound to Definition 1, first observe that product uniformity implies perfect secrecy. The decoder also extends from t shares to any larger honest multiset: order its values by a fixed public ordering of the alphabet and select the first t occurrences. These occurrences belong to distinct participants even when some values coincide, so reconstruction returns the same secret for every such selection.

The dealer's probabilities fit the counting argument as well. For rational probabilities, represent each supported distribution rule by the appropriate number of equally likely indexed copies. The argument of [12] counts symbol multiplicities and supports, both of which are preserved by this representation. For an arbitrary finite law, fix its support. Normalization and uniform marginals impose rational linear equations on its positive probabilities. Rational solutions are dense in their affine solution space, so a sufficiently close rational solution preserves positivity on the same support. It also preserves correctness, which depends only on which share tuples occur.

In particular, for L-bit secrets and M-bit shares,

$$M \ge \left\lceil \log_2\bigl(1 + (2^L - 1)(n - t + 1)\bigr) \right\rceil .$$

For $3 \le t < n$ the expression inside the logarithm may be increased by one. For one-bit secrets these statements give $|\Omega| \ge n - t + 2$, and $|\Omega| \ge n - t + 3$ in the strict range.

To compare additive overheads, put $c = n - t + 1$. Dropping integer rounding gives

$$M - L \ge \log_2\bigl(c - (c - 1)2^{-L}\bigr) .$$

For fixed n, t, the right-hand side tends to $\log_2(n-t+1)$ as L grows, whereas our general construction has an overhead of $O(t^2 \log^2 n)$. Thus rate approaching one leaves a gap in the dependence on the number of participants and the threshold.

6 Conclusion

Efficient exact anonymous threshold sharing is possible with $L + O(t^2 \log^2 n)$ bits per share. The direction-space compiler reduces the extra cost of anonymity to a short coordinate description, and the affine-relation construction protects that description efficiently and exactly at every threshold. For fixed (n, t), the overhead has an upper bound independent of secret length, so the rate approaches one. This supports threshold recovery from share values alone without looking up participant identities.

The direction-space description has at most $t-1$ field coordinates, each protected by the one-bit construction. Reducing the cost of protecting that description, or reducing the one-bit dimension while keeping a positive exact relation law, may further shorten shares at general thresholds. For $1 < t < n$, the known alphabet bound gives an unrounded overhead lower bound tending to $\log_2(n-t+1)$ as the secret grows, while our general construction uses $O(t^2 \log^2 n)$ extra bits. Closing this gap would determine the cost of exact anonymity more precisely.

References

1. Beimel, A.: Secret-sharing schemes: A survey. In: Coding and Cryptology – IWCC 2011. Lecture Notes in Computer Science, vol. 6639, pp. 11–46. Springer (2011). https://doi.org/10.1007/978-3-642-20901-7_2
2. Bishop, A., Green, M., Ishai, Y., Jain, A., Lou, P.: Fully anonymous secret sharing. In: Advances in Cryptology – CRYPTO 2025. Lecture Notes in Computer Science, vol. 16003, pp. 356–389. Springer, Cham (2025). https://doi.org/10.1007/978-3-032-01884-7_12
3. Blundo, C., Stinson, D.R.: Anonymous secret sharing schemes. Discrete Applied Mathematics **77**(1), 13–28 (1997). https://doi.org/10.1016/S0166-218X(97)89208-6
4. Con, R.: Anonymous Shamir's secret-sharing via Reed–Solomon codes against permutations, insertions, and deletions. IEEE Transactions on Information Theory **71**(12), 9534–9547 (2025). https://doi.org/10.1109/TIT.2025.3622889
5. Con, R., Guo, Z., Li, R., Zhang, Z.: Random Reed–Solomon codes achieve the Half-Singleton bound for insertions and deletions over linear-sized alphabets. In: 52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025). Leibniz International Proceedings in Informatics, vol. 334, pp. 60:1–60:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2025). https://doi.org/10.4230/LIPIcs.ICALP.2025.60
6. Con, R., Li, R., Mazor, N.: Fully anonymous perfect secret-sharing. IACR Cryptology ePrint Archive, Report 2026/2145 (2026), https://eprint.iacr.org/2026/2145

7. Daza, V., Domingo-Ferrer, J.: On partial anonymity in secret sharing. In: Public Key Infrastructure – EuroPKI 2007. Lecture Notes in Computer Science, vol. 4582, pp. 193–202. Springer (2007). https://doi.org/10.1007/978-3-540-73408-6_14
8. Deng, Y.p., Guo, L.f., Liu, M.l.: Constructions for anonymous secret sharing schemes using combinatorial designs. Acta Mathematicae Applicatae Sinica, English Series **23**(1), 67–78 (2007). https://doi.org/10.1007/s10255-006-0350-3
9. Eldridge, H., Beck, G., Green, M., Heninger, N., Jain, A.: Abuse-resistant location tracking: Balancing privacy and safety in the offline finding ecosystem. In: 33rd USENIX Security Symposium (USENIX Security 24). pp. 5431–5448. USENIX Association, Philadelphia, PA, USA (2024), https://www.usenix.org/conference/usenixsecurity24/presentation/eldridge
10. Farràs, O., Guiot, M.: Traceable secret sharing schemes for general access structures. In: Advances in Cryptology – EUROCRYPT 2026. Lecture Notes in Computer Science, vol. 16543, pp. 3–32. Springer, Cham (2026). https://doi.org/10.1007/978-3-032-25324-8_1
11. Guillermo, M., Martin, K.M., O'Keefe, C.M.: Providing anonymity in unconditionally secure secret sharing schemes. Designs, Codes and Cryptography **28**(3), 227–245 (2003). https://doi.org/10.1023/A:1024198519111
12. Kishimoto, W., Okada, K., Kurosawa, K., Ogata, W.: On the bound for anonymous secret sharing schemes. Discrete Applied Mathematics **121**(1–3), 193–202 (2002). https://doi.org/10.1016/S0166-218X(01)00236-0
13. Paskin-Cherniavsky, A., Olimid, R.F.: On cryptographic anonymity and unpredictability in secret sharing. Information Processing Letters **161**, 105965 (2020). https://doi.org/10.1016/j.ipl.2020.105965
14. Phillips, S.J., Phillips, N.C.: Strongly ideal secret sharing schemes. Journal of Cryptology **5**(3), 185–191 (1992). https://doi.org/10.1007/BF02451114
15. Sahni, P.S., Koga, H.: Proposal of a ramp scheme with anonymous reconstruction for a one-bit secret using Fermat's little theorem. IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences (2026). https://doi.org/10.1587/transfun.2026CIL0003, advance online publication, 5 August 2026
16. Shamir, A.: How to share a secret. Communications of the ACM **22**(11), 612–613 (1979). https://doi.org/10.1145/359168.359176
17. Stinson, D.R., Vanstone, S.A.: A combinatorial approach to threshold schemes. In: Advances in Cryptology – CRYPTO '87. Lecture Notes in Computer Science, vol. 293, pp. 330–339. Springer (1988). https://doi.org/10.1007/3-540-48184-2_28
18. Zhang, Y., Sun, Y., Zhang, X., Ge, G.: Random Reed–Solomon codes correcting permutations, insertions, and deletions over polynomial-size alphabets. arXiv:2606.22344v1 (2026), https://arxiv.org/abs/2606.22344v1