Strong keys for tensor isomorphism cryptography

Anand Kumar Narayanan

SandboxAQ, Palo Alto, CA, USA
anand.kumar@sandboxaq.com

Abstract. Sampling a non degenerate (that is, invertible) square matrix over a finite field is easy, draw a random square matrix and discard if the determinant is zero. We address the problem in higher dimensions, and sample non degenerate *boundary format* tensors, which generalise *square* matrices. Testing degeneracy is conjectured to be hard in more than two dimensions [10, Conj. 13.1(i)], precluding the "draw a random tensor and discard if degenerate" recipe. The difficulty is in computing hyperdeterminants, higher dimensional analogues of determinants. Instead, we start with a structured random non degenerate tensor and scramble it by infusing more randomness while still preserving non degeneracy. We propose two kinds of scrambling. The first is multiplication in each dimension by random invertible matrices, which preserves dimension and format. Assuming pseudo randomness of this action, which also underlies tensor isomorphism based cryptography, our samples are computationally indistinguishable from uniform non degenerate tensors. The second scrambling employs tensor convolution (that generalises multiplication by matrices) and can increase dimension. Inspired by hyperdeterminant multiplicativity, we devise a recursive sampler that uses tensor convolution to reduce the problem from arbitrary to three dimensions.

Our sampling is a candidate solution for drawing public keys in tensor isomorphism based cryptography, since non degenerate tensors elude recent weak key attacks targeting public key tensors either containing geometric structures such as "triangles" [19] or being deficient in tensor rank [7]. To accommodate our sampling, tensor isomorphism based schemes need to be instantiated in boundary formats such as $(2k+1) \times (k+1) \times (k+1)$, away from the more familiar $k \times k \times k$ cubic formats. Our sampling (along with the recent tensor trapdoor one-way functions [15]) makes an enticing case to transition tensor isomorphism cryptography to boundary formats tensors, which are true analogues of square matrices.

Keywords: tensors · finite fields · post-quantum cryptography

1 Introduction

1.1 Cryptographic context and consequences

Tensor isomorphism based cryptography. A three dimensional tensor can be transformed into another by a triple of invertible square matrices, by multiplying in each dimension. The action is symmetric: if a triple of matrices takes

a tensor A to B, then the triple of inverses of those matrices takes B to A. Two tensors that can be transformed into each other are called isomorphic. The decision version of the tensor isomorphism problem (TI) is given two tensors over a finite field to tell if they are isomorphic. The promise search version asks for an isomorphism (a triple of invertible matrices) between two given isomorphic tensors over a finite field. The tensor isomorphism problem is complete for the complexity class **TI** which contains other cryptographically important problems such as cubic polynomial equivalence, matrix code equivalence, alternating trilinear form equivalence (ATFE) etc. and longstanding group theoretic problems such as p-group isomorphism [9]. Many of these problems are connected by tight nearly linear or quadratic time complexity reductions. Matrix code equivalence is in fact the same as **TI**, merely phrased differently. ATFE is **TI** restricted to alternating tensors with the same invertible matrix acting in all three dimensions, and remains **TI**-complete. Cubic polynomial equivalence is **TI** restricted to symmetric tensors. These problems may be phrased in arbitrary dimensions, but they remain hard even when restricted to three dimensions. With a difficult to find isomorphism problem at hand, it is hard not to design a zero knowledge identification scheme using the Goldreich-Micali-Wigderson construction, which yields signature schemes through the Fiat-Shamir transform. Following this motif, Patarin in his pioneering work on multivariate cryptography proposed a signature scheme based on the hardness of cubic polynomial equivalences [18]. Recently, NIST first round on-ramp signature schemes MEDS [3] and ALTEQ [1] are built on TI and ATFE respectively, with competitive efficiency and signature sizes. Beyond efficiency, part of the appeal of tensor isomorphism problems is that there is strong evidence of average case hardness. Further, attempts to solve it using extensions of Shor's algorithm on a quantum computer must confront the hidden subgroup problem over (products of) general linear groups, believed to be among the hardest [8]. But recently the following weak key vulnerability was identified, warranting caution in the key generation algorithms of tensor isomorphism based cryptosystems.

Weak key attacks on tensor isomorphism based cryptosystems. The public verification key in MEDS was a random $(k+1)\times(k+1)\times(k+1)$ tensor. But very recently, in reaction to a cryptanalytic attack by Narayanan, Qiao and Tang [16], MEDS adopted a $(k_1+1)\times(k_2+1)\times(k_3+1)$ format where k_1, k_2, k_3 are not all equal, but still fairly balanced. We will informally call it a nearly cubic format. For levels I,III and V, the respective choices are $26\times25\times25$, $35\times34\times34$ and $45\times44\times44$ [19, Table 7]. The private signing key is a triple of matrices. In ALTEQ, the public key is a random $(k+1)\times(k+1)\times(k+1)$ alternating tensor and the private key is an invertible matrix. Recently, Ran and Samardjiska identified tensors that have certain geometric structures called "triangles" for MEDS and 3-dimensional 2-singularity for ALTEQ as weak public keys [19]. Heuristically, a public key has a triangle (or a 3-dimensional 2-singularity) with probability roughly $1/q$, where q is the field size. They further devised Gröbner basis algorithms (augmented with equations encoding the triangle/ 3-dimensional 2-

singular structure) to compute such a triangle (or a 3-dimensional 2-singularity), which upon finding reveals a secret key. Conditioned on a weakness being present, their algorithm can detect and find the weakness faster than previously known (but still in exponential time, so there is no contradiction to the belief that testing singularity is hard). For MEDS, since the underlying prime 4093 is small, taking the conservative assumption that weak keys are likely, the triangle finding algorithm takes away 6 bits of security [19, Table 6]. For ALTEQ, the algorithm to find 3-dimensional 2-singular structures is much faster, exploiting the antisymmetry of the alternating tensors and the fact that the same matrix acts in all three dimensions. In fact, for ALTEQ level I parameters, the weakness (if it exists) can be found fast in practice. Luckily for ALTEQ, the field size is a large prime close to 2^{32}. Therefore, the probability of drawing a weak key is heuristically at most 2^{-32}. In the future, ALTEQ has the choice of doubling the bit length of the current prime, and safely ignoring the weak key issue as a rare occurrence that only happens with probability within the NIST accepted bound of 2^{-64}. Beyond signature schemes, certain commitment schemes were also under attack, exploiting the chance that the public key is tensor rank deficient [7]. We will focus on evading the attacks on MEDS as the testing ground. Since non degeneracy is a stronger guarantee than tensor rank, the attacks on the commitment scheme in [7] is also evaded by our sampling algorithm for MEDS. Sampling for ALTEQ is complicated by the antisymmetry. Symmetric tensors (polynomials) and antisymmetric tensors have hyperdeterminants that split into irreducible factors indexed by the symmetries [17]. A modified version of our construction taking this splitting into account might work for ALTEQ, but we defer that to future work.

Security boost by non degenerate sampling in boundary formats. Increasing the field size to dodge the weak key attacks on MEDS seems natural. Yet to be certain it works, the heuristic weak key probability estimate of $1/q$ needs to be proven. In validating the need for their weak key attack, [19] proved a lower bound on the weakness probability. But to be certain that increasing the field sizes is a provable remedy to the weak key attacks, we need an upper bound. Such a rigorous upper bound using hyperdeterminants and counting points on projective varieties over finite fields was proven in [11]. Consequently, increasing the field size to be exponentially large is a provably sound strategy against the weak key attacks. But increasing the field size q comes at the cost of efficiency and signature sizes, both of which grow roughly as $\log q$. For instance, if ALTEQ were to double the bit length of their field size, then the signature sizes would double too. We present a sampling algorithm for boundary formats that applied in the MEDS context samples exclusively from strong public keys, evading the weak key attacks. The sampling algorithm works for boundary formats in arbitrary dimensions, accommodating any future applications in cryptography and beyond. In the MEDS context, the existence of “triangles" coincides precisely with degeneracy and we merely need to specialize our algorithm to the three dimensional and move the nearly cubic MEDS format $(k_1+1)\times(k_2+1)\times(k_3+1)$ to

a boundary format of the form $(k_2+k_3+1)\times(k_2+1)\times(k_3+1)$. The public keys sampled are indistinguishable from uniformly drawn strong public keys assuming hardness of the tensor isomorphism problem. Therefore, the sampling sidesteps the weak key issue without the need of any new assumptions. We gain back the bits of security lost to the weak key attacks and further open up the possibility of instantiating MEDS and similar schemes over small field sizes. More generally, there is a recent trapdoor construction from tensors that works exclusively in boundary formats [15]. We anticipate future cryptographic constructions set in boundary formats, for which our sampling algorithms apply in full generality in arbitrary dimensions. A high level description of the non degenerate tensor sampling problem and our sampling algorithms follows, aided by illustrative small examples in three dimensions.

1.2 The non degenerate tensor sampling problem.

Sampling a non degenerate matrix over a finite field is easy: draw a random square matrix and test if the determinant is non zero. We next take time to informally state its higher dimensional generalisation, sampling non degenerate tensors of three or more dimensions, as solving it is our primary goal.

Degeneracy. An r dimensional tensor to us is an element A in the tensor product of (dual) vector spaces of dimensions $k_1+1, k_2+1, \ldots, k_r+1$ over a finite field $\mathbb{F}_q$. We also think of such a tensor as a multilinear form given by a $(k_1+1)\times(k_2+1)\times\ldots\times(k_r+1)$ format r-dimensional matrix. We will call k_j+1 as the length in the j-th dimension. The lengths k_j+1 have a plus one for convenience, as working in k_j dimensional projective space will make k_j appear often in formulae. Square matrices correspond to the $r=2$, $k_1=k_2$ with the same length in both dimensions. A square matrix being degenerate corresponds to its determinant vanishing. Since the determinant is polynomial time computable, it is easy to test for degeneracy. Further, the determinant is a monic irreducible polynomial in the entries of the square matrix. Therefore, the set of degenerate matrices forms a closed sub variety, in fact a hypersurface. A generic matrix lies outside this hypersurface. Therefore, drawing a generic matrix and testing if the determinant is non-zero samples from non degenerate square matrices. Consider the following algebraic definition of degeneracy of square matrices: A matrix A is degenerate if there is a pair of non zero vectors $(w^{\text{left}}, w^{\text{right}})$ such that $w^{\text{left}}A$ and Aw^{right} are both zero vectors. The definition may seem atypical, but is clarified by thinking of $(w^{\text{left}}, w^{\text{right}})$ as a pair of left and right kernel vectors. It is this motif that easily generalises in the following definition for three dimensional tensors. A three dimensional tensor A (trilinear form) is degenerate, if there exists a triple $(w^{(1)}, w^{(2)}, w^{(3)})$ of non zero kernel vectors such that evaluating the trilinear form at all but one (so, two) of the vectors results in the all zero dual vector. That is,

$$A\left(*, w^{(2)}, w^{(3)}\right) = 0,\ A\left(w^{(1)}, *, w^{(3)}\right) = 0,\ A\left(w^{(1)}, w^{(2)}, *\right) = 0,$$

are zero (dual) vectors in the first, second and third dimension respectively. Likewise, an r dimensional tensor is degenerate if there is an r-tuple of non zero vectors such that evaluating the r-linear form at all but one of the vectors gives the zero (dual) vector. See [5, Chap. 4](also § 2) for a formal algebraic definition and also an equivalent analytic definition in terms of singularity. Singular and degenerate are equivalent in our contexts.

Hyperdeterminants and tensor formats. Cayley discovered an analogue of the determinant for the $2 \times 2 \times 2$ format, depicted below with the vertices of the cube indexing the tensor entries [2].

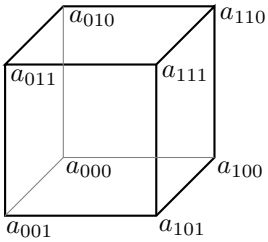

$$\begin{aligned}\text{Hyperdeterminant}= &a_{000}^2a_{111}^2 + a_{001}^2a_{110}^2 + a_{010}^2a_{101}^2 + a_{011}^2a_{100}^2\\ &-2(a_{000}a_{001}a_{110}a_{111} + a_{000}a_{010}a_{101}a_{111} + a_{000}a_{011}a_{100}a_{111}\\ &+a_{001}a_{010}a_{101}a_{110} + a_{001}a_{011}a_{110}a_{100} + a_{010}a_{011}a_{101}a_{100})\\ &+4(a_{000}a_{011}a_{101}a_{110} + a_{001}a_{010}a_{100}a_{111}).\end{aligned}$$

Nearly two centuries later, Gelfand, Kapranov and Zelevinsky generalised Cauchy's construction to arbitrary r dimensional formats satisfying the convexity condition

$$k_j \leq \sum_{\ell \neq j} k_\ell, \quad \forall 1 \leq j \leq r$$

under the name of hyperdeterminants. Analogous to the determinant, the hyperdeterminant is an integer polynomial in the entries of the tensor that vanishes precisely when the tensor is singular. However, the hyperdeterminant is conjectured to be hard (VNP hard to compute, NP hard to zero test) in three or more dimensions [10]. Therefore, it is unlikely for the strategy of generating a random tensor and testing for its hyperdeterminant to be non zero to work efficiently. Formats satisfying the convexity condition satisfies as a strict inequality for all dimensions j are called **interior formats**. Formats satisfying the convexity condition with equality for at least one dimensions j are called **boundary formats**. We will call formats not satisfying the convexity condition as **exterior formats**. Interior and boundary formats have an associated hyperdeterminant, to help reason about degeneracy. Degenerate tensors of an exterior format form a variety of co-dimension more than one, meaning there is no single polynomial such as a hyperdeterminant to characterise degeneracy. Across all formats, non degenerate tensors form a Zariski closed set of co-dimension at least one. Therefore, most tensors are non degenerate.

The Problem: Given an r-dimensional format $(k_1+1)\times(k_2+1)\times\ldots\times(k_r+1)$ and a finite field $\mathbb{F}_q$, sample a non degenerate tensor of that format over $\mathbb{F}_q$.

1.3 Our contribution: Sampling algorithms for boundary formats

We present two families of sampling algorithms for boundary and exterior formats. The first step in the first family of algorithms is to construct a structured random non degenerate tensor of the desired format, leveraging the theory of hyperdeterminants. The structures we exploit vary across interior, boundary and exterior formats and all the necessary mathematical ingredients were developed by Gelfand-Kapranov-Zelevinsky [6,5], Weyman-Zelevinsky [22], Kaji [12], Dionisi-Ottaviani [4], and others. We are merely identifying their relevance to cryptography and putting the ingredients together. Once such a non degenerate diagonal tensor is at hand, we scramble it by acting on each dimension with multiplication by a non-singular matrix, as the second step. This scrambling is precisely the same action intrinsic to the definition of the tensor isomorphism problem. Crucially, scrambling preserves the non-singularity of the tensors we so carefully constructed. Beyond non degeneracy (and other similar hard to compute tensor isomorphism invariants), the scrambling hides structure of the tensor we started with. The resulting tensor is non degenerate and indistinguishable from random tensors under the tensor isomorphism hardness assumption. Even finding one non degenerate tensor of the given format is interesting! But, we want to sample from a distribution with support close to the number $q^{(k_1+1)(k_2+1)\dots(k_r+1)}$ of tensors of that format. We set $\log_q(|\text{Support}|)$ as a coarse estimate of the sampler's "degrees of freedom" and try to maximise it, hoping to approach $\Theta((k_1+1)(k_2+1)\dots(k_r+1))$. This is a crude metric, since a true measure of pseudo randomness would judge the distribution of our samples across the orbits of the scrambling action. But this is difficult to gauge, given the difficulty of the tensor isomorphism problem and we leave it for future work.

We next sketch the main ideas of the first step for interior, boundary and exterior formats, with the aid of illustrative small examples.

Diagonal interior and boundary format tensors. For every interior or boundary format, Weyman and Zelevinsky defined a notion of diagonal tensors. The diagonal entries are easiest to describe for three dimensional boundary formats $(k_1+1)\times(k_2+1)\times(k_3+1)$, say with $k_1 = k_2 + k_3$. In this case, the diagonal entries are those whose first coordinate index is the sum of the rest of the coordinate indices. A $3\times 2\times 2$ boundary format diagonal tensor is pictured below, with only the diagonal entries labeled and visible.

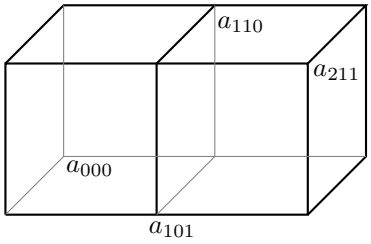

Hyperdeterminant $= \pm a_{000}^{\times} a_{101}^{\times} a_{110}^{\times} a_{211}^{\times}$,
where a "$\times$" denotes some positive exponent.

The description of diagonal entries for interior formats needs a little more notation and we defer it to later sections. Either way, for all interior and boundary formats, there is a simple rule to identify which tensor coordinates are diagonal. Weyman and Zelevinsky proved that hyperdeterminants of interior or boundary format have a monomial lying as a vertex of the Newton polytope, consisting purely of positive powers of all the diagonal entries [22][Theorem 7.1]! Further, a boundary format tensor whose diagonal entries are precisely the non-zero entries is non degenerate. This leads to a simple sampling algorithm. Pick a diagonal boundary format tensor with random non zero diagonal entries. The remedy for the weak key attacks on MEDS is thus simple, to merely sample the public key tensors this way and scramble. For the MEDS relevant $(2k+1)\times(k+1)\times(k+1)$ boundary formats, the degrees of freedom is $\Theta(k^2)$. For interior format diagonal tensors, we can write down a monomial in the hyperdeterminant involving all the diagonal entries. But there could be other monomials making it difficult to find an assignment of diagonal entries such that the hyperdeterminant provably does not vanish. The obvious question we leave open is if there is a polynomial time algorithm to sample non degenerate interior format tensors, even when restricted to cubic three dimensional formats. We work out some small examples and hint at the possibility of using diagonal tensors towards this goal.

Vandermonde-Weyman-Zelevinsky boundary format tensors. Boundary formats accomodate a curious alternative to diagonal tensors, albeit with much fewer degrees of freedom. Weyman and Zelevinsky constructed boundary format analogues of Vandermonde matrices, which we propose to use to generate non-singular tensors [22]. Classical Vandermonde matrices are structured square matrices completely described by a vector of entries, whose singularity is characterised by the distinctness of the entries of the vector. Weyman and Zelevinsky constructed structured boundary format tensors completely characterised by $r-1$ vectors, packaged as the columns of a defining matrix. Its degeneracy is characterised by the distinctness of the column entries of the description matrix. We will call such tensors Vandermonde-Weyman-Zelevinsky (VWZ) tensors. By simply drawing a description matrix whose columns each have distinct elements, we construct a non-singular boundary format tensor. This construction works in arbitrary dimension. As an example, a 2×3 matrix defines a $3\times 2\times 2$ boundary format VWZ tensor, as in figure 1.

Binet-Cauchy and high dimensional boundary format tensors. As the second family of samplers, we propose a recursive algorithm for constructing high dimensional non degenerate boundary format tensors from fewer dimensional ones by exploiting the multiplicativity of hyperdeterminants. Dionisi and Ottaviani's [4] high dimensional analogue of the the Binet-Cauchy theorem is key to preserving non degeneracy during the recursion. This may be seen as a reduction of the non-singular tensor sampling problem (for boundary formats) from arbitrary to three dimensions. The smaller dimensional problems may be solved by sampling diagonal tensors, Vandermonde-Weyman-Zelevnisky tensors

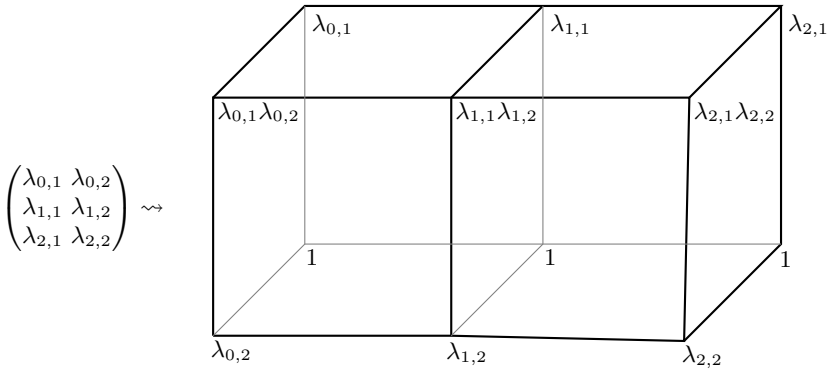

Fig. 1. A $3 \times 2 \times 2$ Vandermonde-Weyman-Zelevinsky tensor.

or by other methods devised in the future. Further, if the smaller dimensional instance is of small enough format, it can be solved exhaustively or using the algorithms in [11]. Curiously, the scrambling we do in the recursive framework at the interior vertices of the recursion tree can scramble the tensors across the orbits of the tensor isomorphism action. Therefore, this is a robust framework promising pseudorandom boundary format non degenerate tensors. Tensor isomorphism based cryptography is based on the action of random tuples of invertible matrices multiplying in each dimension. In our recursive sampling, we identify a more general scrambling action by convolution of non degenerate tensors. It is a fascinating open problem to construct cryptographic primitives using convolution of non degenerate tensors. A challenge or an opportunity in using convolution of non degenerate tensors is that the tensor dimension increases.

Exterior format tensors as boundary format slices. For exterior formats, we start with a non degenerate boundary format tensor in one higher dimension. This boundary format is chosen to envelope the target exterior format, such that the slices of the boundary format in the longest dimension are of the target exterior format. Non degeneracy of the enveloping boundary format ensures non degeneracy of every slice in the longest dimension.

1.4 Nuances in positive characteristic.

Evidently, the theory of hyperdeterminants developed by Gelfand, Kapranov and Zelevinsky [6][5, Chap. 14] play a central role in our constructions, alongside further developments by Weyman and Zelevinsky [22]. But these foundational works built the theory over the complex numbers. We need the theory to hold in positive characteristic. We could look to generic model theoretic tools such as Lefschetz principle to lift the theory from complex numbers to positive characteristic, say to an algebraic closure of our finite field. But such a generic

method would exclude a finite set of primes from being the characteristic, without explicit knowledge of which primes are exclude. This is unsatisfactory for cryptography, where we need to know if the theorems work for our chosen characteristic. Further, the very definition of hyperdeterminants uses geometric tools (such as tangency and projective duality) that need great care while translating to positive characteristic [13,14]. Thankfully, Kaji proved that the hyperdeterminant theory does indeed translate to every prime characteristic [12]. We take this for granted and invoke theorems from hyperdeterminant theory originally stated in characteristic zero without further clarification. One exception is the degeneracy characterisation of Vandermonde-Weyman-Zelevinsky tensors [22, Prop. 7.3]. We observe that only one direction of this characterisation is needed by us. For this direction, we present a self contained elementary exposition of Weyman and Zelevinsky's proof in theorem 1, making it apparent that is works in all characteristics.

2 Sampling non degenerate tensors

Let $\mathbb{F}_q$ denote the finite field with q elements. For positive numbers $k_1, k_2, \ldots, k_r$, an r-dimensional tensor over $\mathbb{F}_q$ of format $(k_1+1) \times (k_2+2) \ldots \times (k_r+1)$ is an element

$$A \in \left(\mathbb{F}_q^{k_1+1}\right)^* \otimes \left(\mathbb{F}_q^{k_2+1}\right)^* \otimes \ldots \otimes \left(\mathbb{F}_q^{k_r+1}\right)^*$$

in the tensor product of dual vector spaces. We will use j exclusively to index dimensions $\{1, 2, \ldots, r\}$. Fix a coordinate system $x^{(j)} = (x_0^{(j)}, x_1^{(j)}, \ldots, x_{k_j}^{(j)})$ for the j^{th}-vector space $\mathbb{F}_q^{k_j+1}$, or equivalently an ordered basis for the dual $\left(\mathbb{F}_q^{k_j+1}\right)^*$. Then, identify A with the r-dimensional matrix

$$A = \left(a_{i_1,i_2,\ldots,i_r}, 0 \le i_1 \le k_1, 0 \le i_2 \le k_2, \ldots, 0 \le i_r \le k_r\right).$$

Associated with A is the multilinear form

$$\begin{aligned} f_A : \mathbb{F}_q^{k_1+1} \times \mathbb{F}_q^{k_2+1} \times \ldots \times \mathbb{F}_q^{k_r+1} &\longrightarrow \mathbb{F}_q \\ \left(w^{(1)}, w^{(2)}, \ldots, w^{(r)}\right) &\longmapsto \sum_{\substack{0 \le i_1 \le k_1 \\ \cdots \\ 0 \le i_r \le k_r}} a_{i_1,i_2,\ldots,i_r} w_{i_1}^{(1)} w_{i_2}^{(2)} \ldots w_{i_r}^{(r)}. \end{aligned}$$

There is a trichotomy of tensor formats depending on the convexity constraint

$$\forall j \in \{1, 2, \ldots, r\}, \qquad k_j \le \sum_{\ell \ne j} k_\ell. \tag{2.1}$$

Formats that satisfy equation 2.1, with the further assurance that there is at least one j satisfying the equation with equality are called as boundary formats. Formats that satisfy equation 2.1 that are not boundary are called as interior. We call all other formats exterior.

Degeneracy and hyperdeterminants. Call the tensor A degenerate if and only if there is an r-tuple of non zero vectors $\left(w^{(1)}, w^{(2)}, \ldots, w^{(r)}\right) \in \mathbb{F}_q^{k_1+1} \times \mathbb{F}_q^{k_2+1} \times \ldots \times \mathbb{F}_q^{k_r+1}$ such that in every dimension $j \in \{1, 2, \ldots, r\}$,

$$\sum_{0 \leq i_j \leq k_j} \left(\sum_{\substack{0 \leq i_1 \leq k_1 \\ \ldots \\ 0 \leq i_r \leq k_r}} a_{i_0, i_1, \ldots, i_r} w_{i_0}^{(0)} w_{i_1}^{(1)} \ldots w_{i_{j-1}}^{(j-1)} w_{i_{j+1}}^{(j+1)} \ldots w_{i_r}^{(r)} \right) x_{i_j}^{(j)} = 0 \left(\in \left(\mathbb{F}_q^{k_j+1}\right)^* \right). \tag{2.2}$$

The inner summation is over all dimensions except j. That is,

$$\sum_{\substack{0 \leq i_0 \leq k_0 \\ \ldots \\ 0 \leq i_r \leq k_r}} a_{i_0, i_1, \ldots, i_r} w_{i_0}^{(0)} w_{i_1}^{(1)} \ldots w_{i_{j-1}}^{(j-1)} w_{i_{j+1}}^{(j+1)} \ldots w_{i_r}^{(r)} = 0, \ \forall 0 \leq j \leq r, \ 0 \leq i_j \leq k_j, \tag{2.3}$$

where again the summation is over all dimensions except j. This notion of degeneracy is identical to that in [5], except that there it is stated in terms of an r-tuple of projective vectors instead of non zero vectors. It is also identical to the existence of "triangles" in [19]. Consider a format $(k_1+1) \times (k_2+2) \ldots \times (k_r+1)$ that is either interior or boundary. The hyperdeterminant

$$Det \in \mathbb{F}_q[a_{i_1, i_2, \ldots, i_r}, 0 \leq i_1 \leq k_1, 0 \leq i_2 \leq k_2, \ldots, 0 \leq i_r \leq k_r.]$$

is an element in the coordinate ring of the tensor that vanishes precisely when A is degenerate. In particular, it is a monic irreducible polynomial in the entries of the tensor whose degree can be exponential in the lengths of the dimensions, even for three dimensions. Each boundary or interior format has an associated hyperdeterminant polynomial, but we suppress this from the notation for hyperdeterminants, as it will be clear from the context. We refer the reader to [5] for a comprehensive treatment on hyperdeterminants. The most critical fact we use is that a boundary or interior format A is degenerate precisely when $Det(A) = 0$.

Tensor isomorphism. An r-tuple

$$(X_1, X_2, \ldots, X_r) \in GL_{k_1+1}(\mathbb{F}_q) \times GL_{k_2+1}(\mathbb{F}_q) \times \ldots \times GL_{k_r+1}(\mathbb{F}_q)$$

of invertible matrices acts on $(k_1 + 1) \times (k_2 + 2) \ldots \times (k_r + 1)$ format tensors A by multiplication in the respective dimensions. We denote the result of the action by $(X_1, X_2, \ldots, X_r) \circ A$. To clarify, the multilinear form $f_{(X_1, X_2, \ldots, X_r) \circ A}$ associated with $(X_1, X_2, \ldots, X_r) \circ A$ is

$$\left(w^{(1)}, w^{(2)}, \ldots, w^{(r)}\right) \longmapsto \sum_{\substack{0 \leq i_1 \leq k_1 \\ \ldots \\ 0 \leq i_r \leq k_r}} a_{i_1, i_2, \ldots, i_r} X_1 w_{i_1}^{(1)} X_2 w_{i_2}^{(2)} \ldots X_r w_{i_r}^{(r)}.$$

We call two tensors A, B isomorphic if there exists

$$(X_1, X_2, \ldots, X_r) \in GL_{k_1+1}(\mathbb{F}_q) \times GL_{k_2+1}(\mathbb{F}_q) \times \ldots \times GL_{k_r+1}(\mathbb{F}_q)$$

such that

$$B = (X_1, X_2, \ldots, X_r) \circ A.$$

If $B = (X_1, X_2, \ldots, X_r) \circ A$, then $A = \left(X_1^{-1}, X_2^{-1}, \ldots, X_r^{-1}\right) \circ B$. Therefore, tensor isomorphism is symmetric and an equivalence relation. The tensor isomorphism problem is to decide if two given tensors are isomorphic, which is believed to be hard. Grochow and Qiao built an intricate web of hard problems that reduce to tensor isomorphism, including some longstanding hard problems that lay at the foundation of multivariate cryptography. Complexity theoretically, the tensor isomorphism problem is $NP \cap co\text{-}AM$, and believed to be hard on average in theory and practice [9]. The best known run time of $p^{O(n^{11/6})}$ is through Sun's p-group isomorphism algorithm [20] (in conjunction with a reduction in [9]). The promise search version, asks for an isomorphism (an r-tuple of invertible matrices) between two given two isomorphic tensors. Spurred on by this hardness, several post-quantum digital signature schemes including MEDS [3] and ALTEQ [1,21] have recently been proposed and part of NIST's first round of on-ramp post-quantum signatures, all reliant on tensor isomorphism hardness assumptions, or hardness assumptions that reduce to tensor isomorphism.

2.1 Diagonal non degenerate tensors

We next describe the sampling algorithm based on the diagonal tensor theory of Weyman and Zelevinsky [22]. For a dimension r, define the semigroup

$$\Phi_r := \left\{ (\ell_1, \ell_2, \ldots, \ell_r) \in \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0} \times \ldots \times \mathbb{Z}_{\geq 0} \mid \ \forall\ 0 \leq j \leq r,\ \ell_j \leq \sum_{j' \neq j} \ell_{j'} \right\}$$

of all non negative index r-tuples satisfying the convexity constraint. Consider r-dimensional formats $(k_1+1) \times (k_2+2) \ldots \times (k_r+1)$ that are either interior or boundary. That is, $(k_1, k_2, \ldots, k_r) \in \Phi_r$. The set of all diagonal indices for the format $(k_1+1) \times (k_2+2) \ldots \times (k_r+1)$ is defined as

$$\mathcal{I}^{Diagonal}_{(k_1+1)\times(k_2+2)\times\ldots\times(k_r+1)} := \{(\ell_1, \ell_2, \ldots, \ell_r) \in \Phi_r \mid (k_1 - \ell_1, k_2 - \ell_2, \ldots, k_r - \ell_r) \in \Phi_r\}.$$

For boundary formats with $k_1 = k_2 + k_3 + \ldots + k_r$, this simplifies to

$$\mathcal{I}^{Diagonal}_{(k_1+1)\times(k_2+2)\times\ldots\times(k_r+1)} = \{(\ell_1, \ell_2, \ldots, \ell_r) \mid \ell_1 = \ell_2 + \ell_3 + \ldots + \ell_r\}.$$

As an example, the diagonal $3 \times 3 \times 3$ interior format tensor looks as in figure 2, with only the diagonal entries visible and labeled. The hyperdeterminant of the $3 \times 3 \times 3$ format was determined in [22, 7.11] using the computer algebra package MACAULAY as

$$\begin{aligned} Det(A) =& (a_{000}a_{222})^8 (a_{110}a_{101}a_{011}a_{112}a_{121}a_{211})^2 [a_{000}^2 a_{111}^4 a_{222}^2 \\ &+ 8a_{000}a_{111}^2 a_{222}(a_{000}a_{112}a_{121}a_{211} + a_{222}a_{110}a_{101}a_{011}) + 16(a_{000}a_{112}a_{121}a_{211})^2 \\ &+ 16(a_{222}a_{110}a_{101}a_{011})^2 - 32a_{000}a_{110}a_{101}a_{011}a_{112}a_{121}a_{211}a_{222}]. \end{aligned} \tag{2.4}$$

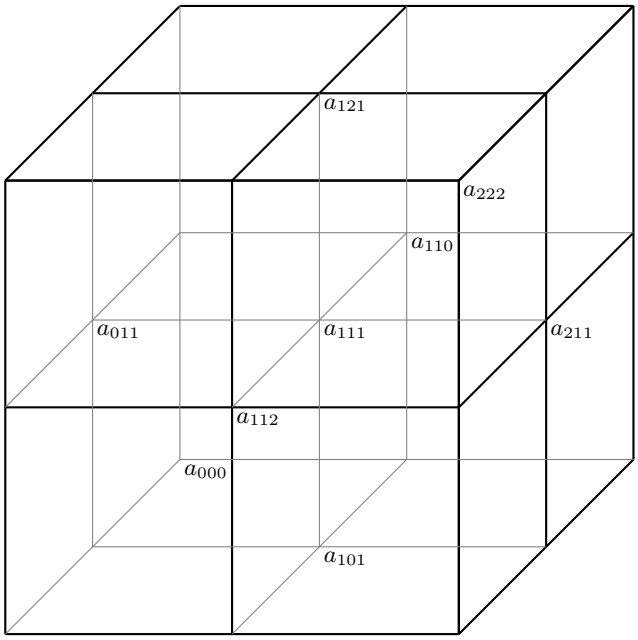

Fig. 2. A $3 \times 3 \times 3$ diagonal tensor.

Apparent from the $3 \times 3 \times 3$ example, the hyperdeterminant of the interior format diagonal tensors is not a monomial. Therefore, merely picking a diagonal tensor with non zero diagonal entries does not imply non degeneracy for interior formats. For $3 \times 3 \times 3$ formats, it is easy to sample non degenerate diagonal tensors: set all but one of the diagonal entries to be random non zero elements, and set the remaining diagonal entry such that equation 2.4 is non zero. For this strategy to generalise, one needs to compute the hyperdeterminant of interior format diagonal tensors, which remains a curious open problem. By [22, Theorem 7.1], the hyperdeterminant (for both interior and boundary formats) has a monomial

$$\pm 1 \prod_{(i_1,i_2,\ldots,i_r)\in\mathcal{I}^{Diagonal}_{(k_1+1)\times(k_2+2)\times\ldots\times(k_r+1)}} a^{>0}_{(i_1,i_2,\ldots,i_r)}$$

lying as a vertex of the Newton polytope, consisting purely of positive powers of all the diagonal entries with coefficient ± 1. Further, the exponents can be determined [22, Remark 7.2(b)]. For the $3 \times 3 \times 3$ example, this corresponds to the monomial $a_{000}^{10} a_{111}^{4} a_{222}^{10} (a_{110} a_{101} a_{011} a_{112} a_{121} a_{211})^2$. Knowing there is such a monomial lying as a vertex of the Newton polytope may be a first step in sampling non degenerate interior format tensors. But we leave this as an open problem. However, hyperdeterminants of boundary format diagonal tensors only consist of this monomial, as observed by Weyman and Zelevinsky [22, Remark 7.3(c)], which we exploit in the following sampling algorithm.

Sampling non degenerate diagonal boundary format tensors.
Input: A finite field $\mathbb{F}_q$ and a boundary tensor format $(k_1+1)\times(k_2+2)\ldots\times(k_r+1)$ with $k_1 = k_2 + k_3 + \ldots + k_r$.

1. Construct a tensor A by setting all the off diagonal entries

$$\left(a_{(i_1,i_2,\ldots,i_r)}, (i_1, i_2, \ldots, i_r) \notin \mathcal{I}^{Diagonal}_{(k_1+1)\times(k_2+2)\ldots\times(k_r+1)}\right)$$

   to zero and drawing the diagonal elements

$$\left(a_{(i_1,i_2,\ldots,i_r)}, (i_1, i_2, \ldots, i_r) \in \mathcal{I}^{Diagonal}_{(k_1+1)\times(k_2+2)\times\ldots\times(k_r+1)}\right)$$

   uniformly and independently at random from $\mathbb{F}_q \setminus \{0\}$.

2. Draw a uniformly random r-tuple

$$(X_1, X_2, \ldots, X_r) \in GL_{k_1+1}(\mathbb{F}_q) \times GL_{k_2+1}(\mathbb{F}_q) \times \ldots \times GL_{k_r+1}(\mathbb{F}_q)$$

   and output

$$(X_1, X_2, \ldots, X_r) \circ A.$$

Lemma 1. *The output $(X_1, X_2, \ldots, X_r) \circ A$ is non degenerate.*

Proof. Weyman and Zeleinsky [22, Remark 7.3(c)] remark that the hyperdeterminant of the boundary format diagonal tensors under consideration is the monomial

$$\pm 1 \prod_{(i_1,i_2,\ldots,i_r)\in\mathcal{I}^{Diagonal}_{(k_1+1)\times(k_2+2)\times\ldots\times(k_r+1)}} a^{>0}_{(i_1,i_2,\ldots,i_r)}$$

consisting purely of positive powers of all the diagonal entries, defined in [22, Theorem 7.1]. To see why, they refer to the remarkable fact proven by Gelfand-Kapranov-Zelevinsky ([6, Theorem 4.3] or [5, Theorem 3.3]) that boundary format tensors have a hyperdeterminant that is the identical to the determinant of a certain (exponentially sized in k_1) square matrix (see also [4], for an alternate proof). When the relation is specialised to diagonal boundary format tensors, the associated exponentially large square matrix has as its diagonal entries, precisely the diagonal entries of the tensor (with possible repetition). Therefore, it is clear that hyperdeterminant of diagonal boundary format tensors is a monomial. Therefore, it must be the monomial whose existence is proven in theorem [22, Theorem 7.1]. In summary, by construction $Det(A) \neq 0$, meaning A is non degenerate. The hyperdeterminant is not an invariant under the $GL_{k_1+1}(\mathbb{F}_q) \times GL_{k_2+1}(\mathbb{F}_q) \times \ldots \times GL_{k_r+1}(\mathbb{F}_q)$ action. It is only an invariant under the $SL_{k_1+1}(\mathbb{F}_q) \times SL_{k_2+1}(\mathbb{F}_q) \times \ldots \times SL_{k_r+1}(\mathbb{F}_q)$ action. But the hyperdeterminant is a relative invariant under the $GL_{k_1+1}(\mathbb{F}_q) \times GL_{k_2+1}(\mathbb{F}_q) \times \ldots \times GL_{k_r+1}(\mathbb{F}_q)$ [5, Chap 14, Prop. 1.4], which implies that

$$Det(A) = 0 \Leftrightarrow Det\left((X_1, X_2, \ldots, X_r) \circ A\right) = 0.$$

Therefore, the output $(X_1, X_2, \ldots, X_r) \circ A$ is non degenerate. □

For the 3-dimensional $(2k+1) \times (k+1) \times (k+1)$ boundary format, we get $\Theta(k^2)$ degrees of freedom, which is fewer in comparison to the number $\Theta(k^3)$ of tensor entries. More generally, for balanced r-dimensional boundary formats $((r-1)k+1) \times (k+1) \times \ldots \times (k+1)$, we get a commendable $\Theta(k^{r-1})$ degrees of freedom compared to the $\Theta(k^r)$ entries. Here, we assume $q > 2$, for otherwise there is only one non zero element to place as the diagonal entries.

2.2 Vandermonde-Weyman-Zelevinsky non degenerate tensors

While constructing a Vandermonde matrix, one starts with a vector and ends up with a square matrix whose determinant vanishes precisely when the vector entries are distinct. Weyman and Zelevinsky constructed higher dimensional analogues of Vandermonde matrices for boundary format tensors. Consider a $(k_1+1) \times (k_2+2) \ldots \times (k_r+1)$ boundary format with $k_1 = k_2 + k_3 + \ldots + k_r$. Start with a $(k_1+1) \times (r-1)$ matrix

$$\Lambda = (\lambda_{i_1,j})_{0 \leq i_1 \leq k_1, 2 \leq j \leq r},$$

and define the Vandermonde-Weyman-Zelevinsky tensor A^Λ with entries

$$\left(a^\Lambda_{i_1,i_2,\ldots,i_r} := \lambda^{i_2}_{i_1,2}\lambda^{i_3}_{i_1,3}\ldots\lambda^{i_r}_{i_1,r}\right)_{0 \leq i_1 \leq k_1, 0 \leq i_2 \leq k_2, \ldots, 0 \leq i_r \leq k_r}.$$

Theorem 1. *[Weyman-Zelevinsky [22, Prop. 7.3]] If* $\forall j \in \{2,3,\ldots,r\}$, $(\lambda_{i_1,j}, 0 \leq i_1 \leq k_1)$ *are distinct (that is, if each column of the starting matrix* Λ *consists of distinct elements), then* A^Λ *is non degenerate.*

Proof. Let $\forall j \in \{1,2,\ldots,r\}$, $(\lambda_{i,j}, 0 \leq i \leq k_1)$ be distinct. Assume that A^Λ is degenerate. For A^Λ to be degenerate, there is an r-tuple of non zero vectors $\left(w^{(2)}, w^{(3)}, \ldots, w^{(r)}\right) \in \mathbb{F}_q^{k_2+1} \times \mathbb{F}_q^{k_3+1} \times \ldots \times \mathbb{F}_q^{k_r+1}$ such that

$$\sum_{\substack{0 \leq i_2 \leq k_2 \\ \ldots \\ 0 \leq i_r \leq k_r}} a^\Lambda_{i_1,i_2,\ldots,i_r} w^{(2)}_{i_2} w^{(3)}_{i_3} \ldots w^{(r)}_{i_r} = 0, \ \forall\ 0 \leq i_1 \leq k_1. \tag{2.5}$$

indexed by the first dimension vanish. To see why, the constraints in equation 2.5 form a subset of the constraints in the defining equation 2.2 of degeneracy. Substituting $a^\Lambda_{i_1,i_2,\ldots,i_r} = \lambda^{i_2}_{i_1,2}\lambda^{i_3}_{i_1,3}\ldots\lambda^{i_r}_{i_1,r}$, we get

$$\sum_{\substack{0 \leq i_2 \leq k_2 \\ \ldots \\ 0 \leq i_r \leq k_r}} \lambda^{i_2}_{i_1,1}\lambda^{i_3}_{i_1,2}\ldots\lambda^{i_r}_{i_1,r} w^{(2)}_{i_2} w^{(3)}_{i_3} \ldots w^{(r)}_{i_r} = 0, \ \forall\ 0 \leq i_1 \leq k_1,$$

which decouples into products as

$$\left(\sum_{i_2=0}^{k_2} w^{(2)}_{i_2}\lambda^{i_2}_{i_1,2}\right)\left(\sum_{i_3=0}^{k_3} w^{(3)}_{i_3}\lambda^{i_3}_{i_1,3}\right)\ldots\left(\sum_{i_r=0}^{k_r} w^{(r)}_{i_r}\lambda^{i_r}_{i_1,r}\right) = 0, \ \forall\ 0 \leq i_1 \leq k_1.$$

For $2 \leq j \leq r$, consider the following polynomials

$$P_j(\Lambda_j) := \sum_{i_j=0}^{k_j} w_{i_j}^{(j)} \Lambda_j^{i_j} \in \mathbb{F}_q[\Lambda_j]$$

in commuting indeterminates Λ_j with the coordinates of $w^{(j)}$ as the coefficients. Every $P_j(\Lambda_j)$ is a non zero polynomial, since each one has coefficients encoding coordinates of non zero vectors. The constraints in equation 2.5 are equivalent to the system of polynomial equations

$$P_2(\lambda_{i_1,2})P_3(\lambda_{i_1,3}) \dots P_r(\lambda_{i_1,r}) = 0, \ \forall 0 \leq i_1 \leq k_1. \tag{2.6}$$

For $2 \leq j \leq r$, let $I_j := \{i_1 \mid P_j(\lambda_{i_1,j}) = 0\}$ be the set of sub indices of the roots of $P_j(\Lambda_j)$. To satisfy equation 2.6, it is necessary that

$$I_2 \cup I_3 \cup \ldots \cup I_r = \{1, 2, \ldots, k_1\}.$$

Therefore,

$$|I_2 \cup I_3 \cup \ldots \cup I_r| = k_1 + 1 = k_2 + k_3 + \ldots + k_r + 1,$$

implying there is at least one dimension j such that $|I_j| > k_j$. But then, the non-zero polynomial $P_j(\Lambda_j)$ has more roots than its degree k_j, a contradiction in any field, irrespective of the characteristic. Therefore our assumption is wrong and A^{Λ} is indeed non-singular. □

Sampling non degenerate Vandermonde-Weyman-Zelevinsky tensors.
Input: A finite field $\mathbb{F}_q$ and a boundary tensor format $(k_1+1)\times(k_2+2)\ldots\times(k_r+1)$ with $k_1 = k_2 + k_3 + \ldots + k_r$ and $q \geq k_1 + 1$.

1. Draw a $(k_1 + 1) \times (r - 1)$ matrix

   $$\Lambda = (\lambda_{i_1,j})_{0 \leq i_1 \leq k_1, 2 \leq j \leq r} \in \mathbb{F}_q^{(k_1+1)\times(r-1)}$$

   uniformly with the restriction that each column has distinct entries. Let A^{Λ} be the Vandermonde-Weyman-Zelevinsky tensor A^{Λ} associated with Λ. That is,

   $$\left(a^{\Lambda}_{i_1,i_2,\ldots,i_r} := \lambda_{i_1,1}^{i_2}\lambda_{i_1,3}^{i_3}\ldots\lambda_{i_1,r}^{i_r}\right)_{0\leq i_1\leq k_1, 0\leq i_2\leq k_2,\ldots,0\leq i_r\leq k_r}.$$

2. Draw a uniformly random r-tuple

   $$(X_1, X_2, \ldots, X_r) \in GL_{k_1+1}(\mathbb{F}_q) \times GL_{k_2+1}(\mathbb{F}_q) \times \ldots \times GL_{k_r+1}(\mathbb{F}_q)$$

   and output

   $$(X_1, X_2, \ldots, X_r) \circ A^{\Lambda}.$$

By theorem 1, A^{Λ} is non degenerate. As before, since the hyperdeterminant is a relative invariant of the $GL_{k_1+1}(\mathbb{F}_q)\times GL_{k_2+1}(\mathbb{F}_q)\times\ldots\times GL_{k_r+1}(\mathbb{F}_q)$ action, the output $(X_1, X_2, \ldots, X_r) \circ A^{\Lambda}$ is non degenerate. The degrees of freedom of the Vandermonde-Weyman-Zelevinsky sampler for boundary formats whose largest dimension has length $k_1 + 1$ is $\Theta(rk_1)$ is small, in comparison to the number of tensor entries. The later can be as big as $\Theta((k_1/r)^r)$.

2.3 Sampling exterior format non degenerate tensors

The difficulty with sampling non degenerate exterior format tensors is that they do not have an associated hyperdeterminant polynomial characterising degeneracy. Our insight is to first look up to appropriate boundary formats in one higher dimension and project back down to a slice in an appropriate dimension.

Sampling non degenerate exterior format tensors.
Input: A finite field $\mathbb{F}_q$ and an exterior tensor format $(k_1+1)\times(k_2+2)\ldots\times(k_r+1)$.

1. Without loss of generality, assume that the first dimension is the longest, that is, $k_1 \geq k_j, \forall 2 \leq j \leq r$. Set $k_{r+1} := k_1 - \sum_{j=2}^{r} k_j$ to ensure that

$$(k_1+1) \times (k_2+1) \times \ldots \times (k_r+1) \times (k_{r+1}+1)$$

   is a boundary format in one higher dimension.
2. Sample a non degenerate tensor A of boundary format $(k_1+1) \times (k_2+1) \times \ldots \times (k_r+1) \times (k_{r+1}+1)$, either by the diagonal sampler, or the Vandermonde-Weyman-Zelevinsky sampler, or by any other means.
3. Draw a uniform $i_{r+1} \in \{0, 1, \ldots, k_{r+1}\}$ and set A^{slice} as the i_{r+1}-th slice of A in the last dimension. That is,

$$a^{slice}_{i_1,i_2,\ldots,i_r} := a_{i_1,i_2,\ldots,i_r,i_{r+1}}, \forall 0 \leq i_1 \leq k_1, 0 \leq i_2 \leq k_2, \ldots, 0 \leq i_r \leq k_r.$$

4. Draw a uniformly random r-tuple

$$(X_1, X_2, \ldots, X_r) \in GL_{k_1+1}(\mathbb{F}_q) \times GL_{k_2+1}(\mathbb{F}_q) \times \ldots \times GL_{k_r+1}(\mathbb{F}_q)$$

   and output

$$(X_1, X_2, \ldots, X_r) \circ A^{slice}.$$

Non degeneracy of A follows from [5, Cor. 3.11]. As before, since the hyperdeterminant is a relative invariant of the $GL_{k_1+1}(\mathbb{F}_q)\times GL_{k_2+1}(\mathbb{F}_q)\times\ldots\times GL_{k_r+1}(\mathbb{F}_q)$ action, the output $(X_1, X_2, \ldots, X_r) \circ A^{slice}$ is non degenerate.

3 Recursive sampling using tensor convolution

Gelfand, Kapranov and Zelevinsky considered a high dimensional analogue of matrix multiplication, which they called convolution [5]. It is a binary operation that takes an r dimensional tensor and an s dimensional tensor to result in an $r+s-2$ dimensional tensor. As the number of columns of the left matrix and number of rows of the right matrix has to agree in usual multiplication, the two dimensions involved in tensor convolution have to be of the same length.

Convolution of tensors. Let A and B be tensors of formats $(k_1+1)\times(k_2+1)\times\ldots\times(k_r+1)$ and $(\ell_1+1)\times(\ell_2+1)\times\ldots\times(\ell_s+1)$ respectively, with two distinguished dimensions (j_1,j_2) such that $k_{j_1}=\ell_{j_2}$. Their convolution $A\star_{(j_1,j_2)}B$ in the (j_1,j_2)-th dimension is defined as the $r+s-2$-dimensional tensor with entries

$$\sum_{i_{j_1}=0}^{k_{j_1}}\sum_{i'_{j_2}=0}^{\ell_{j_2}} a_{i_1,i_2,\ldots,i_{j_1},\ldots,i_r}b_{i'_1,i'_2,\ldots,i'_{j_2},\ldots,i'_s},$$

$$0\le i_1\le k_1, 0\le i_2\le k_2,\ldots,0\le i_{j_1-1}\le k_{j_1-1}, 0\le i_{j_1+1}\le k_{j_1+1},\ldots,0\le i_r\le k_r,$$
$$0\le i'_1\le \ell_1, 0\le i'_2\le \ell_2,\ldots,0\le i'_{j_2-1}\le \ell_{j_2-1}, 0\le i'_{j_2+1}\le \ell_{j_2+1},\ldots,0\le i'_s\le \ell_s.$$

We can think of it as an inner product coupling the j_1-th dimension of A with the j_2-th dimension of B, which are of the same length. To simplify the tedious notation, here on, we will (i) restrict to boundary formats with the implicit assumption that the longest dimension is the first, and (ii) fix $j_2=1$. Therefore, we are only considering convolutions involving the longest dimension on the right side tensor. Let A be a $(k_1+1)\times(k_2+1)\times\ldots\times(k_r+1)$ boundary format tensor with $k_1=k_2+k_3+\ldots+k_r$. Let B be $(\ell_1+1)\times(\ell_2+1)\times\ldots\times(\ell_s+1)$ boundary format tensor with $\ell_1=\ell_2+\ell_3+\ldots+\ell_s$. The convolution

$$A\star_j B$$

with respect to a dimension j (implicitly, on the left) such that $k_j=\ell_1$ is the

$$(k_1+1)\times(k_2+1)\times\ldots\times(k_{j-1}+1)\times(k_{j+1}+1)\ldots\times(k_r+1)\times (\ell_2+1)\times(\ell_3+1)\times\ldots\times(\ell_s+1)$$

format $r+s-2$-dimensional tensor with entries

$$\sum_{i_j=0}^{k_j}\sum_{i'_1=0}^{\ell_1} a_{i_1,i_2,\ldots,i_j,\ldots,i_r}b_{i'_1,i'_2,\ldots,i'_s},$$

$$0\le i_1\le k_1, 0\le i_2\le k_2,\ldots,0\le i_{j_1-1}\le k_{j_1-1}, 0\le i_{j_1+1}\le k_{j_1+1},\ldots,0\le i_r\le k_r,$$
$$0\le i'_2\le \ell_2, 0\le i'_3\le \ell_3,\ldots,0\le i'_s\le \ell_s.$$

Note that the resulting convolution is again of boundary format with the first dimension being the longest,

$$k_1=k_2+k_3+\ldots+k_{j-1}+k_{j+1}+\ldots+k_r+\ell_2+\ell_3+\ldots+\ell_s,$$

consistent with our implicit notation.

Multiplicativity of hyperdeterminants of boundary format. The determinant of a product of square matrices is the product of the determinants. Binet-Cauchy generalises this multiplicativity of determinants when the matrices multiplied are not necessarily square. Dionisi and Ottaviani proved a high

dimensional analogue of Binet-Cauchy for hyperdeterminants, which when specialised to boundary format tensors gives the following multiplicative property of hyperdeterminants.

Theorem 2. *[Dionisi-Ottaviani [4]] Let A and B be non degenerate boundary format tensors of formats $(k_1+1)\times(k_2+1)\times\ldots\times(k_r+1)$ and $(\ell_1+1)\times(\ell_2+1)\times\ldots\times(\ell_s+1)$ respectively. For every dimension j such that $k_j=\ell_1$,*

$$Det(A\star_j B) = Det(A)^{\binom{\ell_1}{\ell_2,\ell_3,\ldots,\ell_s}} Det(B)^{\binom{k_1+1}{k_2,k_3,\ldots,k_{j-1},k_{j+1},\ldots,k_r}},$$

where the exponents on the right are multinomial coefficients. In particular, if $Det(A)$ and $Det(B)$ are non zero, then so is $Det(A\star_j B)$.

This suggests a recipe to build non degenerate boundary format tensors of high dimension from non degenerate boundary format tensors of fewer dimensions.

Sampling across isomorphism orbits. Let us begin to explore this idea with sampling in 4-dimensions. Let A and B be non degenerate boundary format tensors of formats $(3k+1)\times(2k+1)\times(k+1)$ and $(2k+1)\times(k+1)\times(k+1)$ respectively. Then $A\star_2 B$ is a non degenerate $(3k+1)\times\times(k+1)\times(k+1)\times(k+1)$ boundary format tensor. We can then scramble $A\star_2 B$ to try to hide the convolution structure, if it were visible at all.

We can do better! Draw $(X_1,X_2,X_3)\in GL_{2k+1}(\mathbb{F}_q)\times GL_{k+1}(\mathbb{F}_q)\times GL_{k+1}(\mathbb{F}_q)$ uniformly at random and scramble B before convolving. The resulting convolutions are likely

$$A\star_2 B \ncong A\star_2((X_1,X_2,X_3)\circ B),$$

where $\ncong$ denotes that they are not isomorphic (unless by extraordinary chance, such as picking a triple of identity matrices to scramble.). This is because (X_1,X_2,X_3) couples the third and fourth dimensions of $A\star_2((X_1,X_2,X_3)\circ B)$[1]. Therefore, there is no reason for $A\star_2 B$ and $A\star_2((X_1,X_2,X_3)\circ B$ to be in the same $GL_{3k+1}(\mathbb{F}_q)\times GL_{k+1}(\mathbb{F}_q)\times GL_{k+1}(\mathbb{F}_q)\times GL_{k+1}(\mathbb{F}_q)$-orbit. Yet, by theorem 2, $A\star_2 B$ and $A\star_2((X_1,X_2,X_3)\circ B$ are both non degenerate. Therefore convolution goes beyond scrambling (which is merely a sequence of convolutions with 2-dimensional non degenerate boundary format tensors, that is, invertible square matrices). Scrambling woven into the fabric of a recursive convolution algorithm therefore promises pseudorandom non degenerate tensors with equidistribution across isomorphism orbits.

We next outline the framework for a recursive algorithm, leaving precise instantiations to the future. The reason is that there is great flexibility in how the problem is divided, which can and should be tailored to the needs of the application.

[1] The coupling is in the third and fourth dimensions because by convention, during convolution, we append the dimensions coming from tensor on the right to the end.

Recursive sampling for boundary formats.
Input: A finite field $\mathbb{F}_q$ and a boundary format $(k_1+1)\times(k_2+2)\ldots\times(k_r+1)$ such that $k_1 = k_2 + k_3 + \ldots + k_r$.

1. If $r = 3$, output a non degenerate tensor of the input format using a sampling technique that works for 3-dimensional boundary formats, such as using diagonal tensors or Vandermonde-Weyman-Zelevnisky tensors.
2. Else, partition the dimensions 2 through r into two non-empty subsets

$$\{2, 3, \ldots, r\} = J_{\text{left}} \bigsqcup J_{\text{right}},$$

such that J_{left} has at least one element and J_{right} has at least two elements. Such a partition always exists since since $r > 3$.
3. Order the sets J_{left} and J_{right} arbitrarily[2]. Set $\ell_1 := \sum_{j\in J_{\text{right}}} k_j$. Recursively solve the problem (over the same field) for boundary formats

$$(k_1+1)\times(\ell_1+1)\times\prod_{j\in J_{\text{left}}}(k_j+1) \quad\text{and}\quad (\ell_1+1)\times\prod_{j\in J_{\text{right}}}(k_j+1),$$

receiving non degenerate boundary format tensors A_{left} and A_{right} as the respective outputs of the recursive calls.

4. Compute the convolution

$$A_{\text{left}} \star_2 A_{\text{right}}.$$

5. Draw a uniformly random r-tuple

$$(X_1, X_2, \ldots, X_r) \in GL_{k_1+1}(\mathbb{F}_q)\times GL_{k_2+1}(\mathbb{F}_q)\times\ldots\times GL_{k_r+1}(\mathbb{F}_q)$$

and output

$$(X_1, X_2, \ldots, X_r)\circ(A_{\text{left}} \star_2 A_{\text{right}}).$$

A high level depiction of the recursive sampler framework is in figure 3.

As before, non degeneracy is preserved by the matrix multiplications in the last step (of each recursive call) due to the relative invariance of the hyperdeterminant. Non degeneracy is preserved by convolutions by theorem 2. Therefore, non degeneracy is preserved by compositions of matrix multiplications and convolutions, ensuring the output is non degenerate.

Implicit in the construction is proof that we can hit every desired boundary format in every desired dimension using this recursive technique. One can balance the split in step 2 by demanding that J_{left} and J_{right} are roughly the same size. This ensures that there are at most $r\log r$ recursive calls, which we have

[2] The ordering is not important here, but is there merely to conform to notation since formats are ordered. After the convolution, by the notational convention, all the dimensions indexing J_{right} and J_{left} will be appended to the end of the format.

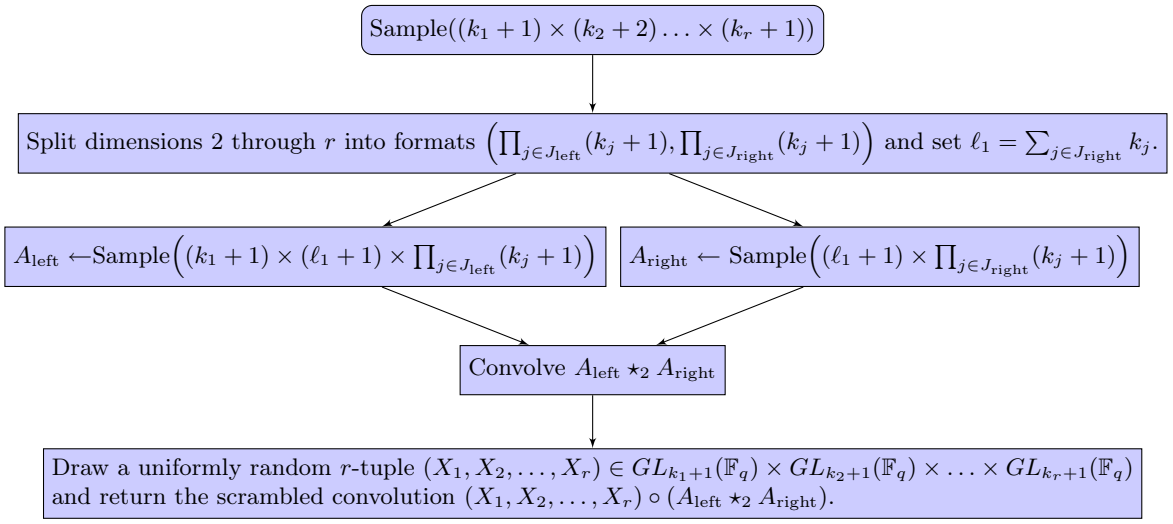

Fig. 3. Recursive sampler, with the field and the base 3-dimensional case implicit.

not tried to optimise.

One shortcoming of this convolution based sampling algorithm is that at each convolution, the degrees of freedom add up (as opposed to multiplying up). Therefore, in terms of degrees of freedom, convolution is comparable to sampling using Vandermonde-Weyman-Zelevinsky tensors and is inferior to diagonal samplers. But scrambling in the interior nodes of the recursion tree help increase the degrees of freedom. Further, merely comparing degrees of freedom may not be fair measure to the recursive sampler, since it scrambles across isomorphism orbits.

Recursive sampling for exterior formats. The recursive sampling technique using convolutions is of no use for interior formats, since the multiplicativity of hyperdeterminants can fail for interior formats [4]. But for exterior formats, we can can invoke the lifting strategy from § 2.3. In particular, given a target exterior format, lift the problem to a boundary format in one higher dimension as detailed in § 2.3, solve it using the recursive sampler and take a random slice to descend back to the target exterior format.

References

1. Bläser, M., Duong, D.H., Narayanan, A.K., Plantard, T., Qiao, Y., Sipasseuth, A., Tang, G.: The alteq signature scheme: Algorithm specifications and supporting documentation (2023), https://pqcalteq.github.io/ALTEQ_spec_2023.09.18.pdf

2. Cayley, A.: On the theory of elimination. Dublin Math. J. p. 116–120 (1848)
3. Chou, T., Niederhagen, R., Persichetti, E., Randrianarisoa, T.H., Reijnders, K., Samardjiska, S., Trimoska, M.: Take your meds: Digital signatures from matrix code equivalence. In: El Mrabet, N., De Feo, L., Duquesne, S. (eds.) Progress in Cryptology - AFRICACRYPT 2023. pp. 28–52. Springer Nature Switzerland, Cham (2023)
4. Dionisi, C., Ottaviani, G.: The binet–cauchy theorem for the hyperdeterminant of boundary format multi-dimensional matrices. Journal of Algebra **259**(1), 591–644 (2003). https://doi.org/https://doi.org/10.1016/S0021-8693(02)00537-9, https://www.sciencedirect.com/science/article/pii/S0021869302005379
5. Gelfand, I., Kapranov, M., Zelevinsky, A.: Discriminants, Resultants, and Multidimensional Determinants. Modern Birkhäuser Classics, Birkhäuser Boston (2009), https://books.google.es/books?id=ZxeQBAAAQBAJ
6. Gelfand, I., Kapranov, M., Zelevinsky, A.: Hyperdeterminants. Advances in Mathematics **96**(2), 226–263 (1992). https://doi.org/https://doi.org/10.1016/0001-8708(92)90056-Q, https://www.sciencedirect.com/science/article/pii/000187089290056Q
7. Gilchrist, V., Marco, L., Petit, C., Tang, G.: Solving the tensor isomorphism problem for special orbits with low rank points: Cryptanalysis and repair of an asiacrypt 2023 commitment scheme. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology – CRYPTO 2024. pp. 141–173. Springer Nature Switzerland, Cham (2024)
8. Grigni, M., Schulman, L., Vazirani, M., Vazirani, U.: Quantum mechanical algorithms for the nonabelian hidden subgroup problem. In: Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing. p. 68–74. STOC '01, Association for Computing Machinery, New York, NY, USA (2001). https://doi.org/10.1145/380752.380769, https://doi.org/10.1145/380752.380769
9. Grochow, J., Qiao, Y.: On the complexity of isomorphism problems for tensors, groups, and polynomials i: Tensor isomorphism-completeness. SIAM Journal on Computing **52**(2), 568–617 (2023). https://doi.org/10.1137/21M1441110, https://doi.org/10.1137/21M1441110
10. Hillar, C.J., Lim, L.H.: Most tensor problems are np-hard. J. ACM **60**(6) (nov 2013). https://doi.org/10.1145/2512329, https://doi.org/10.1145/2512329
11. Joux, A., Narayanan, A.K.: A High Dimensional Cramer's Rule Connecting Homogeneous Multilinear Equations to Hyperdeterminants. In: Meka, R. (ed.) 16th Innovations in Theoretical Computer Science Conference (ITCS 2025). Leibniz International Proceedings in Informatics (LIPIcs), vol. 325, pp. 62:1–62:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2025). https://doi.org/10.4230/LIPIcs.ITCS.2025.62, https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2025.62
12. Kaji, H.: On the duals of segre varieties. Geometriae Dedicata **99**(1), 221–229 (Jun 2003). https://doi.org/10.1023/A:1024968503486, https://doi.org/10.1023/A:1024968503486
13. Kleiman, S.: Tangency and Duality. Københavns Universitet. Matematisk Institut (1985), https://books.google.es/books?id=M9MlrgEACAAJ
14. Kleiman, S., Piene, R.: On the inseparability of the gauss map. American Journal of Mathematics. **123**, 107–129 (1991)
15. Narayanan, A.K.: Trapdoor one-way functions from tensors. Cryptology ePrint Archive, Paper 2025/624 (2025), https://eprint.iacr.org/2025/624
16. Narayanan, A.K., Qiao, Y., Tang, G.: Algorithms for matrix code and alternating trilinear form equivalences via new isomorphism invariants. In: Advances

in Cryptology – EUROCRYPT 2024: 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26–30, 2024, Proceedings, Part III. p. 160–187. Springer-Verlag, Berlin, Heidelberg (2024). https://doi.org/10.1007/978-3-031-58734-4_6, https://doi.org/10.1007/978-3-031-58734-4_6

17. Oeding, L.: Hyperdeterminants of polynomials. Advances in Mathematics **231**(3), 1308–1326 (2012). https://doi.org/https://doi.org/10.1016/j.aim.2012.06.023, https://www.sciencedirect.com/science/article/pii/S0001870812002447
18. Patarin, J.: Hidden fields equations (hfe) and isomorphisms of polynomials (ip): Two new families of asymmetric algorithms. In: Maurer, U. (ed.) Advances in Cryptology — EUROCRYPT '96. pp. 33–48. Springer Berlin Heidelberg, Berlin, Heidelberg (1996)
19. Ran, L., Samardjiska, S.: Rare structures in tensor graphs - bermuda triangles for cryptosystems based on the tensor isomorphism problem. Cryptology ePrint Archive, Paper 2024/1396 (2024), https://eprint.iacr.org/2024/1396
20. Sun, X.: Faster isomorphism for p-groups of class 2 and exponent p. In: Proceedings of the 55th Annual ACM Symposium on Theory of Computing. p. 433–440. STOC 2023, Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/10.1145/3564246.3585250, https://doi.org/10.1145/3564246.3585250
21. Tang, G., Duong, D.H., Joux, A., Plantard, T., Qiao, Y., Susilo, W.: Practical post-quantum signature schemes from isomorphism problems of trilinear forms. Eurocrypt 2022 (2022), https://eprint.iacr.org/2022/267
22. Weyman, J., Zelevinsky, A.: Singularities of hyperdeterminants. Annales de l'Institut Fourier **46**(3), 591–644 (1996). https://doi.org/10.5802/aif.1526, http://www.numdam.org/articles/10.5802/aif.1526/