

An algorithm to compute the t -value of a digital net and of its projections[☆]

Pierre Marion^a, Maxime Godin^a, Pierre L'Ecuyer^{b,*}

^a*CMAP, Ecole Polytechnique, 1 route de Saclay, 91128 Palaiseau, France*

^b*DIRO, Université de Montréal, Pavillon Aisenstadt, C.P. 6128, Succ. Centre-Ville,
Montréal, H3C 3J7, Canada*

Abstract

Digital nets are among the most successful methods to construct low-discrepancy point sets for quasi-Monte Carlo integration. Their quality is traditionally assessed by a measure called the t -value. A refinement computes the t -value of the projections over subsets of coordinates and takes a weighted average (or some other function) of these values. It is also of interest to compute the t -values of embedded nets obtained by taking subsets of the points. In this paper, we propose an efficient algorithm to compute such measures and we compare our approach with previously proposed methods both empirically and in terms of computational complexity.

Keywords: Quasi-Monte Carlo, digital net, t -value

2000 MSC: 11K45, 65C05, 82C80

1. Introduction

The *Monte Carlo method* (MC) simulates random variables from arbitrary distributions by uniform sampling of independent random points from the s -dimensional unit hypercube $[0, 1]^s$ for some integer s , followed by an appropriate transformation to achieve the target distribution [2]. *Quasi-Monte*

[☆]This research was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Institute for Data Valorization (IVADO).

*Corresponding author. lecuyer@iro.umontreal.ca

Email addresses: pierre.marion@polytechnique.edu (Pierre Marion),
maxime.godin@polytechnique.edu (Maxime Godin), lecuyer@iro.umontreal.ca
(Pierre L'Ecuyer)

Carlo (QMC) replaces these independent points by deterministic point sets that are carefully constructed to cover the hypercube much more evenly than i.i.d. points [4, 11, 12, 20, 21]. The most popular construction methods are lattice rules and digital nets. This paper concerns the latter.

An s -dimensional *digital net in base b* is a construction that defines $n = b^k$ points in $[0, 1)^s$ for some integer $k > 0$, via a linear mapping, as detailed in Section 2. A traditional measure of quality for these constructions is the so-called t -value, which measures in some sense the lack of uniformity of the points. Bounds on the integration error in terms of the t -value of the net have been derived, as well as bounds on the variance when the digital net is randomized in specific ways [4, 20, 22]. More refined error or variance bounds can be obtained if we rewrite the integrand as a sum of lower-dimensional components, e.g., via an ANOVA decomposition, and bound the error or variance of each term. Doing this using t -values involves bounding a weighted sum of the t -values of the $2^s - 1$ projections of the point set over subsets of coordinates. This motivates *figures of merit* defined by such sums, to evaluate the quality of digital nets and select good ones for QMC integration. To compute these figures, one must compute the t -value for each projection, which may take excessive computing time, given that calculating a single t -value already requires testing the rank of a combinatorial number of matrices.

It is also useful to construct and test sequences of *embedded digital nets* for which each net is a subset of the next one and such that each net in the sequence has good quality. Generally, if the cardinality of a net is b^k , the cardinality of the next one will be b^{k+1} . One then wants to compute the figure of merit for each net in the sequence, and perhaps take the worst-case as a global measure of quality of the sequence. This is useful when the required number of points is unknown in advance and one wishes to increase n until reaching the desired accuracy [4, 7].

The aim of this paper is to propose a new algorithm to compute the t -value of a given digital net, or the t -values of an arbitrary choice of its projections on subsets of coordinates, and do this for sequences of embedded nets as well. Our algorithm outperforms previously known methods when n is large or when we want the t -values of many projections. We provide an open-source software that implements our algorithm as well as three previously-proposed ones from [3, 23, 24], and we provide a detailed comparison in terms of both empirical performance and computational complexity. For now, the software is implemented only for basis $b = 2$, which is by far the most widely used basis in practice, for efficiency reasons.

In Section 2, we recall the notions of digital net, t -value, as well as error bounds and figures of merit based on t -values. The new algorithm is introduced in Section 3. A time complexity comparison is presented in Section 4. In Section 5, we present empirical performance comparisons.

2. Digital nets and t -values

2.1. Definitions

Let $b \geq 2$ be a prime number and let \mathbb{F}_b denote the finite field of cardinality b . We identify the elements of \mathbb{F}_b with $\{0, 1, \dots, b-1\}$ where 1 is the unit element and the operations correspond to the arithmetic modulo b . Let $s \geq 1$, $k \geq 1$ and $r \geq k$ be integers. Let C_1, \dots, C_s be s matrices of size $r \times k$ over \mathbb{F}_b . For $i = 0, \dots, b^k - 1$, let $\mathbf{i} = (a_0, \dots, a_{k-1})^T \in \mathbb{F}_b^k$ be the column vector that contains the k digits of the base- b expansion of $i = \sum_{\ell=0}^{k-1} a_\ell b^\ell$. For each coordinate $j = 1, \dots, s$, let

$$(y_{i,j,1}, \dots, y_{i,j,r})^T = C_j \cdot \mathbf{i} \quad \text{and} \quad x_{i,j} = \sum_{\ell=1}^r y_{i,j,\ell} b^{-\ell}.$$

and let $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,s})^T$. The point set $\mathcal{P} = \{\mathbf{x}_0, \dots, \mathbf{x}_{n-1}\}$ is called a *digital net in base b* . It contains $n = b^k$ points laying in the half-open s -dimensional unit hypercube $[0, 1)^s$. The matrices C_1, \dots, C_s are called the *generator matrices* of \mathcal{P} . Niederreiter [19, 20] introduced the notion of digital net with a more general definition, in which the operations are in a general finite commutative ring and an arbitrary bijection can be used between the elements of this ring and the ring \mathbb{Z}_b . Popular constructions that are special cases of the above definition were introduced earlier by Sobol' [26] (in base 2) and Faure [6] (in prime base $b \geq 2$).

In our definition, each coordinate has r digits, but only the first k digits will matter for the properties examined in this paper, so we can assume that $r = k$ and ignore the extra rows of the generator matrices. We also always assume that the first k rows of C_j are linearly independent. Then, the digital net is *fully projection-regular*, which means that for each j the one-dimensional projection of \mathcal{P} over the j th coordinate is the set $\{0, 1/n, \dots, (n-1)/n\}$. These numbers are enumerated in a different order for the different coordinates, i.e., the matrices C_j implement permutations of this set, and the choice of these permutations is crucial for the uniformity of the projections of \mathcal{P} over subspaces in two or more dimensions.

This uniformity of \mathcal{P} can be assessed as follows [4, 14, 20]. We can select s non-negative integers q_1, \dots, q_s such that $q := q_1 + \dots + q_s \leq k$, and divide axis j in b^{q_j} equal parts, for $j = 1, \dots, s$. This defines a partition of $[0, 1]^s$ into b^q identical rectangular boxes which are shifted copies of $[0, b^{-q_1}] \times \dots \times [0, b^{-q_s}]$. If each of those boxes contains exactly the same number of points from \mathcal{P} , which has to be b^{k-q} points in each box, we say that \mathcal{P} is (q_1, \dots, q_s) -*equidistributed in base b* .

The digital net \mathcal{P} is said to be a (t, k, s) -*net in base b* if it is (q_1, \dots, q_s) -equidistributed for any (q_1, \dots, q_s) such that $q = q_1 + \dots + q_s \leq k - t$. The t -*value* of a digital net is the smallest t for which it is a (t, k, s) -net. We denote it by $t(\mathcal{S})$ where $\mathcal{S} = (C_1, \dots, C_s)$ to emphasize that it is a function of the vector of generator matrices. We want the t -value to be as small as possible, ideally 0, although it is known that $t = 0$ is impossible unless $b \geq s - 1$ [20, Corollary 4.21], and much tighter lower bounds on the best possible t for any given values of b, k , and s can also be computed via linear programming [18].

If $\mathbf{u} = \{j_1, \dots, j_d\} \subseteq \{1, \dots, s\}$ is a nonempty subset of coordinates, the *projection* $\mathcal{P}_{\mathbf{u}}$ of \mathcal{P} over the coordinates in \mathbf{u} is also a digital net, in $d = |\mathbf{u}|$ dimensions, with vector of generator matrices $\mathcal{S}_{\mathbf{u}} = (C_{j_1}, \dots, C_{j_d})$. We are also interested in computing the t -values $t(\mathcal{S}_{\mathbf{u}})$ for these projected point sets $\mathcal{P}_{\mathbf{u}}$. The assumption of full projection-regularity implies that $t(\mathcal{S}_{\mathbf{u}}) = 0$ whenever $\mathbf{u} = \{j\}$, i.e., when the projection is over a single coordinate.

A sequence of embedded digital nets of cardinalities $b^{m_0}, b^{m_0+1}, \dots, b^k$, for two arbitrary integers $1 \leq m_0 \leq k$, can be obtained as follows. Let $\mathcal{S} = (C_1, \dots, C_s)$ denote the vector of generator matrices of the largest net, with $m = k$. For $j = 1, \dots, s$ and $m = m_0, \dots, k$, let $C_{j,m}$ be the $k \times m$ matrix formed by the first k rows and m columns of C_j . Then $\mathcal{S}_m = (C_{1,m}, \dots, C_{s,m})$ is the vector of generator matrices of a digital net which is the restriction of the largest net to its first b^m points. Each $C_{j,m}$ is assumed to be non-singular so that this digital net is also fully projection-regular. We denote by $\mathcal{S}_{\mathbf{u},m}$ the vector of generator matrices for the projection \mathbf{u} , restricted to their first k rows and m columns. In this context, we are interested in computing the t -values $t(\mathcal{S}_{\mathbf{u},m})$ for all $m \in \{m_0, \dots, k\}$ and a large set of projections \mathbf{u} .

2.2. Figures of merit based on t -values of projections

For a given vector of generator matrices \mathcal{S} , we are interested in computing figures of merit of the following general form:

$$D_{\tilde{t},\gamma}^{(\mathbf{q})}(\mathcal{S}) = \left[\sum_{\emptyset \neq \mathbf{u} \subseteq \{1, \dots, s\}} (\gamma_{\mathbf{u}} \tilde{t}(\mathcal{S}_{\mathbf{u}}))^{\mathbf{q}} \right]^{1/\mathbf{q}} \quad \text{if } \mathbf{q} < \infty, \quad (1)$$

$$D_{\tilde{t},\gamma}^{(\infty)}(\mathcal{S}) = \max_{\emptyset \neq \mathbf{u} \subseteq \{1, \dots, s\}} \gamma_{\mathbf{u}} \tilde{t}(\mathcal{S}_{\mathbf{u}}) \quad \text{if } \mathbf{q} = \infty, \quad (2)$$

where $\mathbf{q} \in [1, \infty]$, $\gamma = \{\gamma_{\mathbf{u}}, \emptyset \neq \mathbf{u} \subseteq \{1, \dots, s\}\}$ is a set of non-negative weights given to the projections [5] and \tilde{t} is a real-valued function of the t -value and other parameters of the net (some examples are given below).

For two integers $1 \leq m_0 \leq k$ and a vector \mathcal{S} of generator matrices of size $k \times k$, we also want to compute the figure of merit for embedded nets:

$$D_{\tilde{t},\gamma,m_0:k}^{(\mathbf{q})}(\mathcal{S}) = \max_{m_0 \leq m \leq k} D_{\tilde{t},\gamma}^{(\mathbf{q})}(\mathcal{S}_m). \quad (3)$$

These general forms leave a lot of freedom for the choice of the function \tilde{t} and the weights γ . Certain choices of \tilde{t} , γ and \mathbf{q} in (1) or (2) yield bounds on the integration error for a large class of integrands; see Subsection 2.3. The weights γ should reflect the importance of having good uniformity in the corresponding projections: the more relevant the projections, the greater the weights. This is discussed in [4] for digital nets and studied in [15, 16, 17] for lattice rules. The choice of \mathbf{q} determines the norm used to combine the values for the different projections. This choice does not affect our algorithm. The most common values are $\mathbf{q} = 2$ and $\mathbf{q} = \infty$. The max in (3) can also be replaced by a weighted sum or a more general norm.

2.3. Discrepancy-based error bounds

An important special case of the figure of merit (2) gives a bound on the *weighted star discrepancy* of \mathcal{P} , which in turns provides a bound on the integration error for a large class of integrands, via the Koksma-Hlawka inequality [4, 10]. For any $\mathbf{z} \in [0, 1]^s$, let $\mathbf{z}_{\mathbf{u}}$ be the projection of \mathbf{z} on the coordinates in \mathbf{u} and $(\mathbf{z}_{\mathbf{u}}, \mathbf{1})$ be the point \mathbf{z} in which the coordinates whose indices are not in \mathbf{u} have been replaced by 1. For positive weights $\gamma_{\mathbf{u}}$, we can define the *weighted Hardy-Krause variation* of a function $f : [0, 1]^s \rightarrow \mathbb{R}$ by

$$V_{\text{HK},\gamma}(f) = \sum_{\emptyset \neq \mathbf{u} \subseteq \{1, \dots, s\}} \gamma_{\mathbf{u}}^{-1} \int_{[0,1]^{|\mathbf{u}|}} \left| \frac{\partial^{|\mathbf{u}|}}{\partial \mathbf{z}_{\mathbf{u}}} f(\mathbf{z}_{\mathbf{u}}, \mathbf{1}) \right| d\mathbf{z}_{\mathbf{u}}, \quad (4)$$

assuming that the partial derivatives exist and are integrable. Moreover, when the integral is 0 for some \mathbf{u} , one can take $\gamma_{\mathbf{u}} = 0$. We also define the *weighted star-discrepancy* of \mathcal{P} as

$$D_{n,\gamma}^*(\mathcal{P}) = \sup_{\mathbf{z} \in [0,1]^s} \max_{\emptyset \neq \mathbf{u} \subseteq \{1, \dots, s\}} \gamma_{\mathbf{u}} \left| \text{vol}[\mathbf{0}_{\mathbf{u}}, \mathbf{z}_{\mathbf{u}}] - \frac{|\mathcal{P}_{\mathbf{u}} \cap [\mathbf{0}_{\mathbf{u}}, \mathbf{z}_{\mathbf{u}}]|}{n} \right|,$$

where $[\mathbf{0}_{\mathbf{u}}, \mathbf{z}_{\mathbf{u}}]$ is the box $[\mathbf{0}, \mathbf{z}]$ projected to the coordinates in \mathbf{u} , $\text{vol}[\mathbf{0}_{\mathbf{u}}, \mathbf{z}_{\mathbf{u}}]$ is the volume of that box, and $|\mathcal{P}_{\mathbf{u}} \cap [\mathbf{0}_{\mathbf{u}}, \mathbf{z}_{\mathbf{u}}]|/n$ is the proportion of the points that are in that box.

The generalized *Koksma-Hlawka inequality* [25] states that

$$\left| \int_{[0,1]^s} f(\mathbf{z}) d\mathbf{z} - \frac{1}{n} \sum_{i=0}^{n-1} f(\mathbf{x}_i) \right| \leq V_{\text{HK},\gamma}(f) \cdot D_{n,\gamma}^*(\mathcal{P}). \quad (5)$$

The weighted star discrepancy can in turn be bounded as follows [4, Theorems 5.12 and 5.26]:

$$D_{n,\gamma}^*(\mathcal{P}) \leq D_{\tilde{t},\gamma,k}^{(\infty)}(\mathcal{S}) \quad (6)$$

with

$$\tilde{t}(\mathcal{S}_{\mathbf{u}}) = b^{t(\mathcal{S}_{\mathbf{u}})-k} \sum_{\ell=0}^{|\mathbf{u}|-1} \binom{k-t(\mathcal{S}_{\mathbf{u}})}{\ell} (b-1)^\ell \quad (7)$$

where b is the base of the digital net and k the size of the generator matrices in $\mathcal{S}_{\mathbf{u}}$. As another special case, [8] made searches for good direction numbers for Sobol' sequences using a heuristic figure of merit that turns out to be a special case of (3) with

$$\tilde{t}(\mathcal{S}_{\mathbf{u}}) = \frac{t(\mathcal{S}_{\mathbf{u}})^p}{k-t(\mathcal{S}_{\mathbf{u}})+1} \text{ and } \gamma_{\mathbf{u}} = \begin{cases} 0.9999^{\min(j_1, j_2)-1} & \text{if } \mathbf{u} = \{j_1, j_2\}, \\ 0 & \text{otherwise,} \end{cases}$$

where k denotes the size of the generator matrices in $\mathcal{S}_{\mathbf{u}}$ and $p > 0$ is a real-valued parameter. Only the two-dimensional projections matter for this measure, and those whose coordinates are both large matter less.

3. Computing t -values for digital nets and their projections

Computing the t -value is a combinatorial problem: one must verify the equidistribution for a large number of possibilities of (q_1, \dots, q_s) , especially if k and/or s are large. Computing the t -value for all or many projections, and for several embedded subnets, as required in (1), (2), and (3), is even more challenging.

When s is large, it can be impractical to use positive weights for all the $2^s - 1$ non-empty subsets \mathbf{u} of coordinates. In the rest of the paper, we assume *finite-order weights* of maximal order $d_{\max} \leq s$, which means that $\gamma_{\mathbf{u}} = 0$ whenever $|\mathbf{u}| > d_{\max}$ [4, page 95]. Note that when $d_{\max} = s$, there is actually no restriction at all on the weights. The choice of d_{\max} is related to the computation budget: the higher the budget, the larger d_{\max} can be.

In this section, we start by explaining the reduction of the computation of the t -value to a simpler combinatorial linear algebra problem, linked to the determination of the ranks of a large number of matrices. After reviewing how this problem has already been tackled in the literature, we describe our proposed method. First we focus on the non-embedded setting, starting from the computation of the t -value of the net and then generalizing to the computation of the t -value over many projections. Finally, we extend our method to an embedded setting where we aim at computing the t -values of many projections for embedded nets.

3.1. Computing the linear independence parameter and the t -value of the net

Given a digital net \mathcal{P} of b^k points with vector of generator matrices $\mathcal{S} = (C_1, \dots, C_s)$, we aim at computing $t(\mathcal{S})$. In other words, we search for the largest q such that \mathcal{P} is (q_1, \dots, q_s) -equidistributed for all tuples (q_1, \dots, q_s) of non-negative integers that sum to q . This q is denoted $\rho(\mathcal{S})$, and is called the *linear independence parameter* of the vector of generator matrices \mathcal{S} . We have the relation $t(\mathcal{S}) = k - \rho(\mathcal{S})$.

As detailed in [4], \mathcal{P} is (q_1, \dots, q_s) -equidistributed if and only if the $q \times k$ *composition matrix*

$$C = \begin{bmatrix} C_1^{(q_1)} \\ \vdots \\ C_s^{(q_s)} \end{bmatrix}$$

formed by the first q_1 rows of C_1 , the first q_2 rows of C_2, \dots , the first q_s rows of C_s , has full rank q . To find $\rho(\mathcal{S})$, we need to find the largest value of q

for which all the composition matrices of size $q \times k$ have full rank q . For any given q , there are $\binom{q+s-1}{s-1}$ such composition matrices [9]. Thus, in the worst case, we have to compute the rank of $\sum_{q=1}^k \binom{q+s-1}{s-1}$ matrices. We can do this either by increasing or decreasing values of $q \in \{1, \dots, k\}$, and stop as soon as we find $\rho(\mathcal{S})$. The best choice of direction depends on the actual value of t . Proposition 5 gives exact expressions for the number of vector additions for each direction.

The algorithms reported in the literature compute the ranks of the composition matrices independently (without reusing the computations from one composition matrix to another), until the t -value can be determined. For the special case of base $b = 2$, Schmid [24] visits the composition matrices by increasing value of q . For each composition matrix, he visits all linear combinations of the rows to check if they are nonzero, by using a Gray code so only one row is changed at each step. Pirsic and Schmid [23] proposed a different method that applies Gaussian elimination [1] independently to each composition matrix to compute its rank. They found that this second method is typically faster than the first for large b and when $k - t$ is not too large, and the first method was usually faster otherwise. We will denote these two methods by S and PS, respectively.

These two approaches are fine to compute a single t -value, but they do not scale well when dealing with the figures of merit in (1) to (3). The latter involve several t -values, one for each of $\sum_{d=1}^{d_{\max}} \binom{s}{d}$ projections \mathbf{u} such that $|\mathbf{u}| \leq d_{\max}$. All in all, we have to compute the rank of up to $\sum_{d=1}^{d_{\max}} \sum_{q=1}^k \binom{s}{d} \binom{q+d-1}{d-1}$ matrices. For $k = 10$ and $s = d_{\max} = 5$, this represents over 1.2×10^4 matrices, cumulating 9.1×10^4 rows. In use cases, k and s easily go up to 20. In this case, the number of matrices equals 2.6×10^{14} , cumulating 4.9×10^{15} rows. In the following, we explain how to optimize such a daunting computation, first on a single projection (optimizing the sum over q), then on many projections (optimizing the sum over d).

Our strategy takes advantage from the fact that one can enumerate the set of composition matrices in a way that any two consecutive composition matrices differ by exactly one row. This relies on the enumeration of weak compositions as proposed in [9]. We build upon the algorithm of Pirsic and Schmid [23], but instead of computing from scratch the Gaussian elimination for each composition matrix, we reuse computations. We detail this algorithm in the following subsection.

3.2. Efficient computation of the rank of all composition matrices

We now explain the reduction used to compute the rank of the composition matrices, which is weaker than the reduction in row echelon form [1]. We call it a *reduction in almost row echelon form* (RAREF). Then, we describe how a RAREF can be efficiently computed for a matrix C' which differs by exactly one row from a matrix C for which a RAREF is already known.

Definition 1. In \mathbb{F}_b , given a matrix C of size $q \times k$ such that $q \leq k$, a triplet (L, T, \mathbf{p}) where L is a $q \times q$ matrix, T is a $q \times k$ matrix and \mathbf{p} is a subset of $\{1, \dots, k\}$, is said to be a RAREF of C if the following properties hold:

- (a) L is a non-singular matrix such that $L \cdot C = T$;
- (b) The columns of T whose indexes are in \mathbf{p} are distinct vectors of the canonical base of \mathbb{F}_b^q . The nonzero coefficients of these columns are called pivots;
- (c) The cardinality of \mathbf{p} equals the rank of T .

Algorithm 1 Computing a RAREF

Let L be the identity matrix of size $q \times q$, $T = C$ and $\mathbf{p} = \emptyset$.

For each row i of T , from the first to the last:

1. Use transvections to zero all the elements of the i th row of T which belong to columns in \mathbf{p} . Apply the same transvections to L .
2. Locate the first nonzero coefficient of the i th row of T , say in the j th column of T , and multiply this row so that this leading nonzero coefficient equals 1. This coefficient is a new pivot. Let $\mathbf{p} = \mathbf{p} \cup \{j\}$. Apply the same multiplication to L . If such a coefficient does not exist, skip the row.
3. Otherwise, use transvections to zero all the elements of the j th column of T other than the pivot. Apply the same transvections to L .

Return (L, T, \mathbf{p}) .

Given a matrix C , one finds a RAREF by applying row operations on each row i of C from the first to the last, as described in Algorithm 1, which is a variation of the Gauss-Jordan elimination [1]. We refer to this as *pivoting*. Since L is non-singular, the rank is invariant to a left-multiplication by L , so

the rank of C must equal the number of pivots in T , i.e., the cardinality of \mathbf{p} . The following proposition proves the correctness of Algorithm 1.

Proposition 1. *For every matrix C , Algorithm 1 returns a RAREF.*

Proof. We check each of the RAREF properties:

- (a) L is non-singular as a product of elementary row-operation matrices. Because we apply the same operations to L and to T , we have $L \cdot C = T$.
- (b) The columns of T indexed by \mathbf{p} belong to the canonical base of \mathbb{F}_b^q thanks to step 3. They are distinct because at most one pivot is selected on each row.
- (c) At the end of the algorithm, each row of T contains exactly one pivot or is all-zero, hence the rank of T equals the number of pivots. \square

Algorithm 2 Updating the RAREF

Let $L' = L$, $T' = T$ and $\mathbf{p}' = \mathbf{p}$.

1. Locate a nonzero coefficient in the i th column of L' . Such a coefficient exists as L' is non-singular. Say it is located in the j th row of L' and it equals $a \in \mathbb{F}_b$. Using a row permutation, exchange the i th and j th rows, both in L' and T' .
2. Using the nonzero coefficient a at position (i, i) of L' and row transvections, zero all the other coefficients in the i th column of L' . Apply the same transvections to T' .
3. Remove from \mathbf{p}' the pivot in the i th row of T' . Set to zero all the coefficients of the i th row of L' , except for (i, i) which remains equal to a . The element (i, i) of L' is now completely isolated. Thus, it is now straightforward to update $T' = L' \cdot C$ into $T' = L' \cdot C'$: one has to replace the i th row of T' by a times the i th row of C' .
4. Pivot on the i th row by applying Algorithm 1.

Return (L', T', \mathbf{p}') .

Given two composition matrices C and C' of size $q \times k$ which differ by exactly one row, we explain how to derive a RAREF (L', T', \mathbf{p}') of C' by reusing a known RAREF (L, T, \mathbf{p}) of C . We can assume that C has full

rank, because otherwise \mathcal{P} would not be (q_1, \dots, q_s) -equidistributed for the tuple (q_1, \dots, q_s) that corresponds to the composition matrix C .

As stated by Proposition 2, Algorithm 2, whose operations are illustrated in Figure 1, updates the RAREF so that all the properties listed in Definition 1 hold. The key advantage of this algorithm is that it requires only one row pivoting (step 4), whereas a full recomputation would require q pivotings.

Proposition 2. *If C and C' are two matrices of the same shape that differ only by one row, C has full rank, and (L, T, \mathbf{p}) is a RAREF of C , then the triplet (L', T', \mathbf{p}') returned by Algorithm 2 is a RAREF of C' .*

Proof. Let i be the index of the row that differs between C and C' . We check each of the RAREF properties:

- (a) Steps 1, 2 and 4 perform only row operations, so they preserve property (a). At step 3, the i th column of L' is all-zero except for the (i, i) coefficient. This ensures that replacing the i th row of L' by $(0, \dots, 0, a, 0, \dots, 0)$ preserves the non-singularity of L' . It is then clear that replacing the i th row of T' by ac'_i implies $L' \cdot C' = T'$.
- (b) Row permutations leave the canonical base invariant so step 1 preserves property (b). After step 2, the only column indexed by \mathbf{p}' which may not belong to the canonical base must have its pivot on the i th row. As we remove this column from \mathbf{p}' at step 3, property (b) is verified after step 3. Step 4 preserves property (b), due to step 3 of Algorithm 1. The distinctness stems from the fact that we may only add a single pivot on the i th row, which had no pivot before step 4.
- (c) Steps 1 and 2 change neither the rank of T' nor the cardinality of \mathbf{p}' . Because C has full rank, the removal of the i th row in step 3 amounts to decreasing by 1 the number of pivots and the rank. After adding back the new row, a new pivot is added during step 4 iff the row is linearly independent from the others (i.e., C' has full rank). \square

3.3. Computing the t -values over many projections

In this section, we extend our proposed methodology to the computation of the t -values of all or many projections. More precisely, given a vector \mathcal{S} of s generator matrices of size $k \times k$, we aim at computing $t(\mathcal{S}_{\mathbf{u}})$ (or equivalently $\rho(\mathcal{S}_{\mathbf{u}})$) for all projections \mathbf{u} such that $|\mathbf{u}| \leq d_{\max}$. We propose

Step 1: Find a row with a nonzero coefficient a in column i of L' , and exchange this row (blue row j) with row i , in L' and in T' .

$$\begin{array}{c}
 \boxed{L : q \text{ rows, } q \text{ columns}} \\
 \left(\begin{array}{cccc}
 l_{11} & \cdots & l_{1i} & \cdots & l_{1q} \\
 \vdots & & \vdots & & \vdots \\
 l_{j1} & \cdots & a \neq 0 & \cdots & l_{jq} \\
 \vdots & & \vdots & & \vdots \\
 l_{i1} & \cdots & l_{ii} & \cdots & l_{iq} \\
 \vdots & & \vdots & & \vdots \\
 l_{q1} & \cdots & l_{qi} & \cdots & l_{qq}
 \end{array} \right)
 \end{array}
 \times
 \begin{array}{c}
 \boxed{C : q \text{ rows } \{c_1, \dots, c_q\}} \\
 \left(\begin{array}{c}
 c_1 \\
 \vdots \\
 c_i \\
 \vdots \\
 c_q
 \end{array} \right)
 \end{array}
 =
 \begin{array}{c}
 \boxed{T = L \cdot C : q \text{ rows } \{t_1, \dots, t_q\}} \\
 \left(\begin{array}{c}
 \vdots \\
 t_j \\
 \vdots \\
 t_i \\
 \vdots
 \end{array} \right)
 \end{array}$$

Step 2: Use the nonzero element a in L' to zero all coefficients on column i of L' , using transvections.

$$\begin{array}{c}
 \left(\begin{array}{cccc}
 l_{11} & \cdots & l_{1i} & \cdots & l_{1q} \\
 \vdots & & \vdots & & \vdots \\
 l_{i1} & \cdots & l_{ii} & \cdots & l_{iq} \\
 \vdots & & \vdots & & \vdots \\
 l_{j1} & \cdots & a & \cdots & l_{jq} \\
 \vdots & & \vdots & & \vdots \\
 l_{q1} & \cdots & l_{qi} & \cdots & l_{qq}
 \end{array} \right)
 \end{array}
 \times
 \begin{array}{c}
 \left(\begin{array}{c}
 c_1 \\
 \vdots \\
 c_i \\
 \vdots \\
 c_q
 \end{array} \right)
 \end{array}
 =
 \begin{array}{c}
 \left(\begin{array}{c}
 t_1 \\
 \vdots \\
 t_j \\
 \vdots \\
 t_i \\
 \vdots \\
 t_q
 \end{array} \right)
 \end{array}$$

Step 3: Set all coefficients of row i of L' to 0, except the diagonal coefficient. Replace row i of T' with the new row ac'_i (in green). Since the i th row of C' (green row) interacts only with the green column of L' in the matrix multiplication, we have $L' \cdot C' = T'$.

Step 4 (not shown): Pivot on row i of T' to obtain the RAREF T' .

$$\begin{array}{c}
 \left(\begin{array}{cccc}
 l_{11} & \cdots & 0 & \cdots & l_{1q} \\
 \vdots & & \vdots & & \vdots \\
 0 & 0 & a & 0 & 0 \\
 \vdots & & \vdots & & \vdots \\
 l_{q1} & \cdots & 0 & \cdots & l_{qq}
 \end{array} \right)
 \end{array}
 \times
 \begin{array}{c}
 \left(\begin{array}{c}
 c_1 \\
 \vdots \\
 c'_i \\
 \vdots \\
 c_q
 \end{array} \right)
 \end{array}
 =
 \begin{array}{c}
 \left(\begin{array}{c}
 t_1 \\
 \vdots \\
 ac'_i \\
 \vdots \\
 t_q
 \end{array} \right)
 \end{array}$$

$$\begin{array}{ccc}
 \boxed{L' : q \text{ rows, } q \text{ columns}} & \boxed{C' : q \text{ rows, } k \text{ columns}} & \boxed{T' = L' \cdot C' : q \text{ rows, } k \text{ columns}}
 \end{array}$$

Figure 1: Procedure to update a reduced almost row echelon form when row i changes

a dynamic programming method that achieves this much more efficiently than computing the t -values independently across all the projections. It exploits a recurrence that expresses $\rho(\mathcal{S}_u)$ in terms of the linear independence parameters of all strict subvectors \mathcal{S}_v , with $v \subset u$.

We define the *partial linear independence parameter* $\tilde{\rho}(\mathcal{S})$ of a vector of d generator matrices \mathcal{S} as the maximum integer $q \geq d$ such that for every tuple of *positive integers* (q_1, \dots, q_d) that sum to q , the composition matrix

$$C = \begin{bmatrix} C_1^{(q_1)} \\ \vdots \\ C_d^{(q_d)} \end{bmatrix}$$

formed by the first q_1 rows of C_1 , the first q_2 rows of C_2, \dots , the first q_d rows of C_d , has full rank q . By convention, we set $\tilde{\rho}(\mathcal{S}) = d - 1$ whenever there is no such q . The difference with the linear independence parameter is that we consider here only tuples of strictly positive integers (q_1, \dots, q_d) and not non-negative integers as before. Note that

$$0 \leq \rho(\mathcal{S}) \leq \tilde{\rho}(\mathcal{S}) \leq k.$$

With these ingredients, we can formulate the computation of $\rho(\mathcal{S})$ as a dynamic programming problem, as follows.

Proposition 3. *For each subset $u \subseteq \{1, \dots, s\}$ of cardinality greater than 2*

$$\rho(\mathcal{S}_u) = \min \left(\tilde{\rho}(\mathcal{S}_u), \min_{\emptyset \neq v \subset u} \rho(\mathcal{S}_v) \right) = \min \left(\tilde{\rho}(\mathcal{S}_u), \min_{j \in u} \rho(\mathcal{S}_{u \setminus \{j\}}) \right). \quad (8)$$

Proof. The first equality stems from the definition of ρ and $\tilde{\rho}$, the second equality comes from the fact that $\rho(\mathcal{S}_{v'}) \geq \rho(\mathcal{S}_v)$ when $v' \subseteq v$. \square

To compute $\rho(\mathcal{S}_u)$, for every u such that $|u| \leq d_{\max}$, we first compute $\rho(\mathcal{S}_{\{j\}})$ for each $j \in \{1, \dots, s\}$. Thanks to the assumption of full projection regularity, we already know that $\rho(\mathcal{S}_{\{j\}}) = k$. Then, we use the recurrence (8) in the following manner: enumerate the subsets of $\{1, \dots, s\}$ by increasing cardinality until reaching cardinality d_{\max} . For each subset u , we first compute $q_{\max} = \min_{j \in u} \rho(\mathcal{S}_{u \setminus \{j\}})$. We know that $\rho(\mathcal{S}_u) \leq q_{\max}$. If $q_{\max} < |u|$, we know that $\rho(\mathcal{S}_u) = q_{\max}$ as $\tilde{\rho}(\mathcal{S}_u) \geq |u| - 1$. Otherwise, we test by decreasing values of $q \in \{q_{\max}, \dots, |u|\}$ whether $\tilde{\rho}(\mathcal{S}_u) \geq q$, by adapting the

procedure described in the two previous subsections. The main technicality is that we need to modify the enumeration of composition matrices to stay in the subset of composition matrices which contain at least one row from each generator matrix of $\mathcal{S}_{\mathbf{u}}$. To update the RAREF when the composition matrix changes by a single row, we use the same method as before.

We visit the values of q in decreasing order, from q_{\max} to $|\mathbf{u}|$, because this turns out to be faster when computing figures of merit that involve a large number of low-dimensional projections: for low-dimensional projections, the chances are that t is small, so $\rho(\mathcal{S}_{\mathbf{u}})$ is nearer to q_{\max} than to $|\mathbf{u}|$. Another reason is that we implement an early stopping criterion: for a given q , as soon as we encounter a composition matrix which does not have full-rank, we stop and move on to $q - 1$. This does not work in the increasing order, as we have to check that all matrices are full-rank for a given q before moving on to $q + 1$.

3.4. Computation of the t -values for embedded nets

We now extend our proposed methodology to the case of embedded nets. More precisely, for given integers $1 \leq m_0 \leq k$ and a vector \mathcal{S} of s generator matrices of size $k \times k$, we aim at computing $t(\mathcal{S}_{\mathbf{u},m})$ for all projections \mathbf{u} such that $|\mathbf{u}| \leq d_{\max}$ and all $m_0 \leq m \leq k$.

The procedure described in the paragraph that follows Proposition 3 can be adapted as follows. For a given $q \in \{q_{\max}, \dots, |\mathbf{u}|\}$, let l_q be the minimal number of columns such that all the composition matrices of shape $q \times m$ for $l_q \leq m \leq k$ have full rank. By convention, if such a number does not exist, we set $l_q = k + 1$. The composition matrices of interest to compute $t(\mathcal{S}_{\mathbf{u},m})$ are the restrictions to their first m columns of the composition matrices used to compute $t(\mathcal{S}_{\mathbf{u},k})$. Thus, for any $m \in \{m_0, \dots, k\}$, $\tilde{\rho}(\mathcal{S}_{\mathbf{u},m}) \geq q$ iff $m \geq l_q$. Thus computing l_q for decreasing values of $q \in \{q_{\max}, \dots, |\mathbf{u}|\}$ yields the linear independence parameter for all m , hence solving the problem.

Algorithm 2 yields l_q as the index of the rightmost column where a pivot is placed when cycling through all composition matrices, or $k + 1$ if it is not possible to add a new pivot at some RAREF update. Let l'_q be the latter quantity and let us prove that $l_q = l'_q$. First, $l_q = k + 1$ iff $l'_q = k + 1$, because all composition matrices have full rank ($l_q \leq k$) iff we add one pivot at each RAREF update ($l'_q \leq k$). Next we focus on the case where we achieve to add one pivot at each update. If a submatrix contains all the pivots, it has full rank, hence $l_q \leq l'_q$. During the update where a pivot was placed for the first time on column l'_q , because we choose the pivot as the *first* nonzero

coefficient of the new row, the restriction of the matrix to its first $l'_q - 1$ columns contains an all-zero row, so it does not have full rank. Hence we cannot have $l_q < l'_q$. Thus, computing the figure of merit (3) for embedded nets does not add significant overhead compared to computing the simpler figures (1) or (2), as we only have to keep tab on the right-most column where a pivot has been placed.

4. Complexity analysis

In this section, we perform a time complexity analysis of our method (MGL) and compare it with methods S and PS from [24] and [23].

For each method, we obtain formulas that provide bounds on the total number of vector additions in \mathbb{F}_b^k required to perform any of the following three computations, given a vector \mathcal{S} of s generator matrices of size $k \times k$: (a) computing the t -value of the net, as described in Subsections 3.1 and 3.2; (b) computing the t -value over many projections, as described in Subsection 3.3; and (c) computing the t -value over several projections for embedded nets, as described in Subsection 3.4. For $b > 2$, scalar multiplications are also used in some algorithms (PS and MGL). We do not count them as it would complicate the formulas without adding further insight and because the number of scalar multiplications is normally less than the number of vector additions.

We also compare our method with the one of Dick and Matsumoto (DM) in [3]. The latter uses weight enumerator polynomials to compute the t -value. The authors also outline an extension to compute the t -values over all projections by using multivariate polynomials, but we do not consider this extension in this paper because we do not have an implementation and it appears non-trivial to implement in practice. Due to the different nature of their method, they express its complexity in terms of integer operations.

For $b = 2$, if we assume that k does not exceed the computer word length, which is usually the case in the context of QMC, then vector additions in \mathbb{F}_b^k can be implemented as XOR operations between two integers. In this case, it makes sense to compare the complexity of DM to the others, because the elementary operation is the same. However, for $b > 2$, comparing vector additions in \mathbb{F}_b^k and integer operations does not really make sense. Moreover, vector additions must be computed as k independent additions in \mathbb{F}_b , so they are essentially k times slower than for $b = 2$.

The memory space required by our method is not at all an issue, as the most significant overhead is to store at each step two matrices of size $k \times k$.

Therefore we only study time complexity and not space complexity.

4.1. Computing the t -value of the net

We first look at the number of additions needed to compute or to update the RAREF:

Proposition 4. *For general b and a matrix of size $q \times k$, with $q \leq k$, computing the RAREF by Algorithm 1 requires at most $4q^2$ vector additions, whereas updating the RAREF by Algorithm 2 requires at most $6q$ vector additions.*

Proof. For Algorithm 1, at most $4q$ vector additions are done for each of the q pivotings ($2q$ at step 1 and $2q$ at step 3), for a total of $4q^2$ operations. For Algorithm 2, $6q$ vector additions are needed ($2q$ at step 2 and $4q$ at step 4), for a total of $6q$ operations. \square

Next we bound the total number of vector additions required to compute a single t -value, for each method. For MGL, we give two bounds, depending on the enumerating order (q increasing or q decreasing), as explained in Subsection 3.1. We also provide simplified asymptotic expressions for the bounds when k tends to infinity with s fixed.

Proposition 5. *The number of operations to compute a t -value with the four methods considered is bounded as follows:*

S	$\sum_{d=1}^s \sum_{q=d}^{k-t} \binom{s}{d} \binom{q-1}{d-1} (b-1)^d b^{q-d} = \mathcal{O}((b+1)^{k-t+s-1})$
PS	$\sum_{q=1}^{k-t} \binom{q+s-1}{s-1} \binom{q}{2} = \mathcal{O}\left(\frac{(k-t)^{s+2}}{s!}\right)$
DM	$\mathcal{O}(ksb^k)$
MGL (q decreasing)	$\sum_{q=k-t}^k [4q^2 + [\binom{q+s-1}{s-1} - 1] 6q] = \mathcal{O}\left(\frac{k^{s+1}}{s!}\right)$
MGL (q increasing)	$\sum_{q=1}^{k-t} [4q^2 + [\binom{q+s-1}{s-1} - 1] 6q] = \mathcal{O}\left(\frac{(k-t)^{s+1}}{s!}\right)$

Proof. The upper bounds for S and PS are derived in [23], in terms of the number of vector additions. The bound for DM can be found in [3].

To prove the bounds for MGL, we first note from Subsection 3.1 that the number of composition matrices whose ranks are to be computed is at most $\sum_{q=k-t}^k \binom{q+s-1}{s-1}$ if we choose to visit the values of q in decreasing order. For each $q \in \{k, \dots, k-t\}$, the first RAREF is computed entirely, followed by $\binom{q+s-1}{s-1} - 1$ updates, which takes $\sum_{q=k-t}^k [4q^2 + [\binom{q+s-1}{s-1} - 1] 6q]$ vector additions in total, according to Proposition 4. By using the fact that $q \leq \binom{q+s-1}{s-1}$ in the first inequality, this expression simplifies as:

$$\begin{aligned}
& \sum_{q=k-t}^k \left[4q^2 + \left[\binom{q+s-1}{s-1} - 1 \right] 6q \right] \\
& \leq 10 \sum_{q=k-t}^k \binom{q+s-1}{s-1} q = 10 \sum_{q=k-t}^k \left[\binom{q+s}{s} - \binom{q+s-1}{s} \right] q \\
& \leq 10 \sum_{q=k-t}^k \left[\binom{q+s}{s} q - \binom{q+s-1}{s} (q-1) \right] \\
& = 10k \binom{k+s}{s} - 10(k-t-1) \binom{k-t-1+s}{s}.
\end{aligned}$$

If the values of q are visited in increasing order, the computation above can be updated as follows:

$$\sum_{q=1}^{k-t} \left[4q^2 + \left[\binom{q+s-1}{s-1} - 1 \right] 6q \right] \leq 10(k-t) \binom{k-t+s}{s}$$

The asymptotic expressions come from

$$\binom{k+s}{s} = \frac{(k+s)(k+s-1)\dots(k+1)}{s!} \sim \frac{k^s}{s!} \text{ when } k \rightarrow \infty.$$

The above calculation can be adapted to prove the asymptotic expression for PS, which is tighter than the one given in [23]:

$$\sum_{q=1}^{k-t} \binom{q+s-1}{s-1} \binom{q}{2} \leq (k-t)^2 \binom{k-t+s}{s} = \mathcal{O} \left(\frac{(k-t)^{s+2}}{s!} \right). \quad \square$$

There is no uniform ranking between those methods, except that S is exponential in both s and k , DM is exponential only in k , and PS and MGL

are exponential only in s . For MGL, the best choice in the enumerating order depends on the value of t : visiting the values of q by decreasing order is the best choice when the t -value is small, whereas increasing order wins when t is far from 0. For $b = 2$ and typical values of s and k , the threshold lies around $t = 2$, so the decreasing order wins only when t is very close to 0.

4.2. Computing the t -value over many projections

We now compare the methods in terms of the computational effort to obtain the t -value over all projections over no more than d_{\max} coordinates. This corresponds to the complexity of computing the figures of merit (1)–(2). Here the values of q are enumerated by decreasing order, which is more effective than increasing order when many low-order projections have to be examined.

Proposition 6. *The number of vector additions to compute the t -value of all projections \mathbf{u} such that $|\mathbf{u}| \leq d_{\max}$ is bounded as follows:*

S	$\sum_{d=1}^{d_{\max}} \sum_{q=d}^k \binom{s}{d} \binom{q-1}{d-1} (b-1)^d b^{q-d}$
PS	$k^2 \sum_{d=1}^{d_{\max}} \binom{s}{d} \binom{k+d}{d}$
DM	$\mathcal{O}\left(kb^k \sum_{d=1}^{d_{\max}} \binom{s}{d} d\right)$
MGL	$k \sum_{d=1}^{d_{\max}} \binom{s}{d} [4k(k-d+1) + 6\binom{k}{d}]$

Proof. For S, the complexity is very similar to that in Proposition 5, because the algorithm to compute the t -value actually yields the t -value over all projections. The outer sum is restricted to the projections of cardinality $d \leq d_{\max}$.

For PS and DM, the complexity is the sum of the complexities of computing the t -values of each projection, which is respectively $k^2 \binom{k+d}{d}$ and $\mathcal{O}(kdb^k)$ for a projection of size d (by Proposition 5).

For MGL, for general b , the algorithm described in Subsection 3.3 requires at most

$$\sum_{d=1}^{d_{\max}} \binom{s}{d} \sum_{q=d}^k \left[4q^2 + \left[\binom{q-1}{d-1} - 1 \right] 6q \right] \quad (9)$$

vector additions, because for each projection \mathbf{u} of cardinality d , we compute by decreasing values of $q \in \{k, \dots, d\}$ the ranks of all composition matrices.

As in Proposition 5, the cost of computing these ranks is $4q^2 + \left[\binom{q-1}{d-1} - 1 \right] 6q$, hence the formula (9). It can be simplified as in the proof of Proposition 5:

$$\begin{aligned} \sum_{q=d}^k \left[4q^2 + \left[\binom{q-1}{d-1} - 1 \right] 6q \right] &\leq 4(k-d+1)k^2 + 6 \sum_{q=d}^k \binom{q-1}{d-1} q \\ &\leq 4(k-d+1)k^2 + 6k \binom{k}{d}. \quad \square \end{aligned}$$

These expressions are difficult to compare in full generality for the parameters d_{\max} , k and s . We can observe two benefits of MGL compared to PS: the algorithm described in Subsection 3.2 replaces the k^2 factor by k , and the algorithm described in Subsection 3.3 replaces the $\binom{k+d}{d}$ factor by approximately $k^2 + \binom{k}{d}$.

We give simplified expressions for three choices of d_{\max} , namely $d_{\max} = 2$, 3, and s . Small values of d_{\max} can be used when the computational budget is limited and s can be large, as in [8] for instance, where $d_{\max} = 2$ was used. Taking $d_{\max} = s$ is the other extreme, and would be doable only if s is small.

Corollary 7. *For the special cases of $d_{\max} = 2$, $d_{\max} = 3$, and $d_{\max} = s$, the number of vector additions to compute the t -value of all projections \mathbf{u} such that $|\mathbf{u}| \leq d_{\max}$ is bounded as follows:*

	$d_{\max} = 2$	$d_{\max} = 3$	$d_{\max} = s$
S	$\mathcal{O}(s^2 b^k)$	$\mathcal{O}(s^3 k b^k)$	$\sum_{d=1}^s \sum_{q=d}^k \binom{s}{d} \binom{q-1}{d-1} (b-1)^d b^{q-d}$
PS	$\mathcal{O}(s^2 k^4)$	$\mathcal{O}(s^3 k^5)$	$k^2 \sum_{d=1}^s \binom{s}{d} \binom{k+d}{d}$
DM	$\mathcal{O}(s^2 k b^k)$	$\mathcal{O}(s^3 k b^k)$	$\mathcal{O}(s 2^{s-1} k b^k)$
MGL	$\mathcal{O}(s^2 k^3)$	$\mathcal{O}(s^3 k^4)$	$k \left[4k^2 2^s + 6 \binom{s+k}{s} \right]$

Proof. The expressions for $d_{\max} = 2$ and 3 follow from Theorem 6. For the case $d_{\max} = s$, for DM, we use the identity $\sum_{d=1}^s \binom{s}{d} d = s 2^{s-1}$. For MGL, we use Vandermonde's convolution $\sum_{d=1}^s \binom{s}{d} \binom{k}{d} = \binom{s+k}{s}$. \square

We see that for $d_{\max} = 2$ or 3, MGL has a better complexity than all the other methods. For $d_{\max} = s$, the formulas are not easy to compare. If we take the limit as $k \rightarrow \infty$, we obtain similar asymptotic expressions as in Proposition 5.

4.3. Computing the t -values for embedded nets

Given two integers $m_0 \leq k$, we now consider the number of elementary operations to compute $t(\mathcal{S}_{\mathbf{u},m})$ for all projections \mathbf{u} such that $|\mathbf{u}| \leq d_{\max}$ and all $m_0 \leq m \leq k$. This represents the complexity of computing the figures of merit (3). This may seem counter-intuitive, but for all four methods, considering embedded nets ($m_0 < k$) does not increase the workload compared to the non-embedded case ($m_0 = k$). The reason is that the computations for $m = k$ include the computations for smaller values of m , by reordering some computations or saving intermediate results if needed. The details for MGL are given in Subsection 3.4.

5. Numerical implementation and experiments

Here we compare the average CPU times required to compute the t -value of a random net, and to compute the t -values of all projections of order up to d_{\max} that are needed to obtain the figure of merit (1) or (2), using our implementations of the four methods S, PS, DM, and MGL, discussed in the previous sections, for digital nets in base $b = 2$. All these methods are implemented efficiently in the *LatNet Builder* software [13]. All experiments were run on an Intel(R) Core(TM) i7-6500U CPU @ 2.50GHz processor.

5.1. Computing the t -value of a random net

Figure 2 compares the four methods in terms of the CPU time required to compute the t -value of a random fully projection-regular net. For each considered pair (s, k) , between 20 and 10^5 fully projection-regular nets were sampled uniformly, and the average CPU time to calculate the t -value was obtained. The number of samples was smaller for longer execution times. In the left panel of Figure 2, for each pair (s, k) , the symbol indicates which of the four methods is fastest on average, with our implementations. As expected from Proposition 5, our strategy is the fastest for a large number of points in low dimension, while DM performs better in large dimension with a small number of points. Method S wins for easy cases (a small number of points in small dimension). In 19 or more dimensions, S is never the fastest method. PS never wins, as it is always slower than MGL.

In the right panel, we fix $s = 12$ and plot the execution time (in log scale) as a function of k . We see that DM performs best for small point sets (when the required CPU time is under about 0.01 seconds), S performs better for

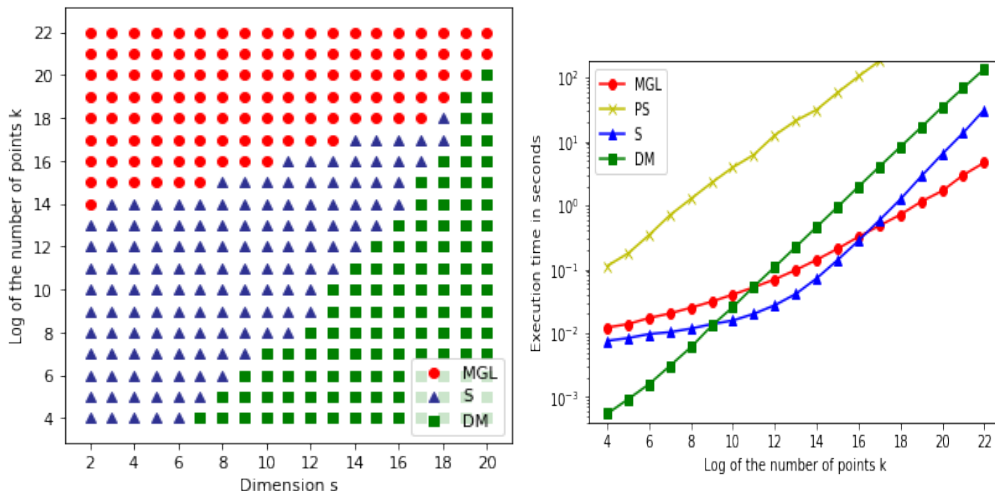


Figure 2: Speed comparison for four methods (S, PS, DM, and MGL) to compute the t -value of a digital net in base 2 with $n = 2^k$ points, in s dimensions. In the left panel, the symbol indicates which method is fastest as a function of (s, k) . In the right panel, we fix $s = 12$ and plot the execution time as a function of k , in log-log scale.

point sets of moderate size (CPU time between 0.01 and 0.5 seconds), and MGL performs better for large point sets (when it takes over 0.5 seconds).

5.2. Computing the t -values of low-dimensional projections

For large high-dimensional nets, e.g., if $s \geq 20$ and $k \geq 20$, looking at the t -values of low-dimensional projections is more reasonable and also more relevant than computing the t -value of the full net [8, 11]. For $s = k = 20$, for example, it takes 100 seconds to compute the t -value of the full net and only 0.5 second to compute the figure of merit (2) for $d_{\max} = 3$. Table 1 compares all four methods in terms of the CPU time required to compute the figure of merit (2) with $m = k$, $\tilde{t}(|\mathbf{u}|, m, t) = t$ (meaning we take the maximum of the t -values of each projections), and weights $\gamma_{\mathbf{u}} = 1$ for $|\mathbf{u}| \leq d_{\max}$ and 0 otherwise. We look at projections up to order $d_{\max} = 2, 3$ and 5, number of dimensions $s = 5, 20, 100$, and number of points n from 2^{10} to 2^{30} . Each experiment was repeated for a sample of 20 random independent uniformly-distributed fully projection-regular nets and we report the average CPU times. For each combination of d_{\max} , s , and k , the smallest value is in bold. For the entries marked “ $> 10^6$ ”, 1000 seconds were not sufficient on average. We see from the

		S method			MGL method		
d_{\max}	k	$s = 5$	$s = 20$	$s = 100$	$s = 5$	$s = 20$	$s = 100$
2	10	0.4	4	76	1	13	305
	15	34	58	1 522	2	31	790
	20	95	1 766	47 159	5	60	1 585
	25	3 538	71 554	$> 10^6$	7	101	2 552
	30	106 048	$> 10^6$	$> 10^6$	10	153	4 002
3	10	0.6	12	1 334	1	45	5 023
	15	4	221	25 832	3	186	23 188
	20	147	7 869	$> 10^6$	8	556	70 215
	25	5 495	306 162	$> 10^6$	17	1 286	168 941
	30	182 193	$> 10^6$	$> 10^6$	30	2 478	336 843
5	10	0.5	81	$> 10^6$	1	237	$> 10^6$
	15	5	1 110	$> 10^6$	4	2 013	$> 10^6$
	20	168	59 586	$> 10^6$	12	16 936	$> 10^6$

		PS method			DM method		
d_{\max}	k	$s = 5$	$s = 20$	$s = 100$	$s = 5$	$s = 20$	$s = 100$
2	10	4	52	1 052	24	417	9 858
	15	9	166	4 264	741	12 155	306 093
	20	26	492	12 607	25 440	413 915	$> 10^6$
	25	61	1 174	30 747	851 757	$> 10^6$	$> 10^6$
	30	139	2 432	63 433	$> 10^6$	$> 10^6$	$> 10^6$
3	10	7	509	66 444	62	4 647	589 036
	15	30	3 127	406 101	1 865	140 408	$> 10^6$
	20	130	12 280	$> 10^6$	63 412	$> 10^6$	$> 10^6$
	25	390	38 414	$> 10^6$	$> 10^6$	$> 10^6$	$> 10^6$
	30	1 011	99 078	$> 10^6$	$> 10^6$	$> 10^6$	$> 10^6$
5	10	10	16 626	$> 10^6$	99	154 950	$> 10^6$
	15	64	175 731	$> 10^6$	3 180	$> 10^6$	$> 10^6$
	20	335	$> 10^6$	$> 10^6$	106 845	$> 10^6$	$> 10^6$

Table 1: Average CPU times (in milliseconds) for computing the figure of merit (2) for $d_{\max} = 2, 3$ and 5, various s and k , with methods S, PS, DM, and MGL.

table that MGL performs best in all cases, except when the number of points is small ($n = 2^{10}$), in which case S performs better. For such a small number, the execution time is very small anyway. For large s and n ($n \geq 2^{25}$), MGL is more than ten times quicker than the second best algorithm (PS).

6. Conclusion

Being able to compute t -values efficiently is critical for the practitioner who wants to find quickly (in real time) good digital nets with arbitrary values of s and k , arbitrary weights, etc., and for the researcher who wants to study numerically the properties of figures of merit based on t -values. We offer in this paper a new efficient algorithm to compute the t -value of a net and of its projections, and to do so for a series of embedded nets as well. We compare the performance of our algorithm with three other ones, proposed in [24], [3] and [23]. We also implement these four algorithms in the open-source software *LatNet Builder* [13] for the most commonly used base, $b = 2$.

According to our numerical comparisons, the method of [24] is the best performer when s and k are small, the method of [3] performs better when s is large and k is small, and our new method is the best performer when k is large and either s is small or we consider only the projections of relatively low order. The latter case is very relevant for practical applications. For example, we can easily compute a figure of merit that accounts for the t -values of all projections of order 3 or less for a net with 2^{25} points in $s = 100$ dimensions, whereas with all others methods, the cost of doing this is prohibitive.

7. References

- [1] R.P. Agarwal, E.C. Flaut, An introduction to linear algebra, Chapman and Hall, 2017.
- [2] S. Asmussen, P.W. Glynn, Stochastic Simulation, Springer-Verlag, New York, 2007.
- [3] J. Dick, M. Matsumoto, On the fast computation of the weight enumerator polynomial and the t value of digital nets over finite Abelian groups, *SIAM Journal on Discrete Mathematics* 27 (2013) 1335–1359.
- [4] J. Dick, F. Pillichshammer, Digital Nets and Sequences: Discrepancy Theory and Quasi-Monte Carlo Integration, Cambridge University Press, Cambridge, U.K., 2010.

- [5] J. Dick, I.H. Sloan, X. Wang, H. Woźniakowski, Liberating the weights, *Journal of Complexity* 20 (2004) 593–623.
- [6] H. Faure, Discr pance des suites associ es   un syst me de num ration en dimension s , *Acta Arithmetica* 61 (1982) 337–351.
- [7] F.J. Hickernell, H.S. Hong, P. L’Ecuyer, C. Lemieux, Extensible lattice sequences for quasi-Monte Carlo quadrature, *SIAM Journal on Scientific Computing* 22 (2001) 1117–1138.
- [8] S. Joe, F.Y. Kuo, Constructing Sobol sequences with better two-dimensional projections, *SIAM Journal on Scientific Computing* 30 (2008) 2635–2654.
- [9] P. Klingsberg, A Gray code for compositions, *Journal of Algorithms* 3 (1982) 41–44.
- [10] G. Larcher, F. Pillichshammer, K. Scheicher, Weighted discrepancy and high-dimensional numerical integration, *BIT Numerical Mathematics* 43 (2003) 123–137.
- [11] P. L’Ecuyer, Quasi-Monte Carlo methods with applications in finance, *Finance and Stochastics* 13 (2009) 307–349.
- [12] P. L’Ecuyer, Randomized quasi-Monte Carlo: An introduction for practitioners, in: P.W. Glynn, A.B. Owen (Eds.), *Monte Carlo and Quasi-Monte Carlo Methods: MCQMC 2016*, Springer, Berlin, 2018, pp. 29–52.
- [13] P. L’Ecuyer, M. Godin, A. Jemel, P. Marion, D. Munger, LatNet Builder: A general software tool for constructing highly uniform point sets, 2019. <https://github.com/umontreal-simul/latnetbuilder>.
- [14] P. L’Ecuyer, C. Lemieux, Recent advances in randomized quasi-Monte Carlo methods, in: M. Dror, P. L’Ecuyer, F. Szidarovszky (Eds.), *Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications*, Kluwer Academic, Boston, 2002, pp. 419–474.
- [15] P. L’Ecuyer, D. Munger, Constructing adapted lattice rules using problem-dependent criteria, in: *Proceedings of the 2012 Winter Simulation Conference*, IEEE Press, 2012, pp. 373–384.

- [16] P. L'Ecuyer, D. Munger, On figures of merit for randomly-shifted lattice rules, in: H. Woźniakowski, L. Plaskota (Eds.), Monte Carlo and Quasi-Monte Carlo Methods 2010, Springer-Verlag, Berlin, 2012, pp. 133–159.
- [17] P. L'Ecuyer, D. Munger, Algorithm 958: Lattice builder: A general software tool for constructing rank-1 lattice rules, ACM Transactions on Mathematical Software 42 (2016) Article 15.
- [18] W.J. Martin, Linear programming bounds for ordered orthogonal arrays and (t, m, s) -nets, in: H. Niederreiter, J. Spanier (Eds.), Monte Carlo and Quasi-Monte Carlo Methods 1998, Springer-Verlag, Berlin, 2000, pp. 368–376.
- [19] H. Niederreiter, Point sets and sequences with small discrepancy, Monatshefte für Mathematik 104 (1987) 273–337.
- [20] H. Niederreiter, Random Number Generation and Quasi-Monte Carlo Methods, volume 63 of *SIAM CBMS-NSF Reg. Conf. Series in Applied Mathematics*, SIAM, 1992.
- [21] H. Niederreiter, Constructions of (t, m, s) -nets and (t, s) -sequences, Finite Fields and Their Applications 11 (2005) 578–600.
- [22] A.B. Owen, Scrambled net variance for integrals of smooth functions, Annals of Statistics 25 (1997) 1541–1562.
- [23] G. Pirsic, W.C. Schmid, Calculation of the quality parameter of digital nets and application to their construction, Journal of Complexity 17 (2001) 827–839.
- [24] W.C. Schmid, The exact quality parameter of nets derived from Sobol' and Niederreiter sequences, Recent Advances in Numerical Methods and Applications (1999) 287–295.
- [25] I.H. Sloan, H. Woźniakowski, When are quasi-Monte Carlo algorithms efficient for high dimensional integrals, Journal of Complexity 14 (1998) 1–33.
- [26] I.M. Sobol', The distribution of points in a cube and the approximate evaluation of integrals, U.S.S.R. Comput. Math. and Math. Phys. 7 (1967) 86–112.