LATTICE COMPUTATIONS FOR RANDOM NUMBERS

RAYMOND COUTURE AND PIERRE L'ECUYER

Abstract. We improve on a lattice algorithm of Tezuka for the computation of the k-distribution of a class of random number generators based on finite fields. We show how this is applied to the problem of constructing, for such generators, an output mapping yielding optimal k-distribution.

1. Introduction

Extensive classes of random number generators have the following structure. The state space is a finite field F of characteristic 2. We denote by d its degree over $\mathbf{F}_2$, and sometimes refer to it as the order of the generator. Any state $y \in F$ evolves into a state xy, where the distinguished element, $x \in F$, completely determines the evolution of the generator. Finally, the generator in state y outputs a w-bit vector $\Phi(y) = (\phi(y_1 y), \ldots, \phi(y_w y)) \in \mathbf{F}_2{}^w$, where $\phi : F \to \mathbf{F}_2$ is any non-zero linear form over $\mathbf{F}_2$, and where $y_1, \ldots, y_l$ are suitably chosen non-zero elements of F.

The study of the k-distribution of the output sequence involves the computation, for all $l \le w$ and $k \le d$, of the rank of the mapping $F \to \mathbf{F}_2{}^{lk}$ defined by

$$y \mapsto \begin{pmatrix} \phi(y_1 y) & \phi(y_1 x y) & \ldots & \phi(y_1 x^{k-1} y) \\ \phi(y_2 y) & \phi(y_2 x y) & \ldots & \phi(y_2 x^{k-1} y) \\ \vdots & & & \\ \phi(y_l y) & \phi(y_l x y) & \ldots & \phi(y_l x^{k-1} y) \end{pmatrix}. \tag{1}$$

One might naturally use gaussian elimination, as is done for instance in [2, 4], but there are other methods which are more efficient in terms of both time and space. The efficiency issue becomes critical if the order d of the generator is chosen large. One such method is proposed by Tezuka [7]. He computes the rank of (1), for a given value of l and all k, by means of an l-dimensional lattice Λ_l in the space $\mathbf{F}_2[X]^l$ of l-tuples of polynomials with $\mathbf{F}_2$ coefficients. We improve on this method by using instead a 'dual' lattice $\Lambda'_l \subset \mathbf{F}_2[X]^l$ which has the advantage that it has basis vector coordinates which are generally much smaller than those of Λ_l, and that a simple relationship between Λ'_l and Λ'_{l+1} allows for recursive computation. We will show how these features are well suited to the problem of constructing, for given F and $x \in F$, an output mapping Φ with optimal k-distribution.

1991 *Mathematics Subject Classification.* 65C10.

Key words and phrases. Random number generation, lattice structure .

This work has been supported by NSERC-Canada grant no. OGP0110050 and SMF0169893 to the second author.

2. Lattices

We will assume that our distinguished element x generates F as a ring so that, as a vector space over $\mathbf{F}_2$, F admits the basis $1, x, \ldots, x^{d-1}$. For $0 \leq k \leq d$, let $F_k \subset F$ denote the $\mathbf{F}_2$-subspace generated by the first k elements in this basis.

Consider the mapping $\mathbf{F}_2[X]^l \to F^l$ given by

$$(P_1(X), \ldots, P_l(X)) \mapsto (P_1(x), \ldots, P_l(x)). \tag{2}$$

The inverse image by this mapping of any F-linear subspace V of F^l, is a sublattice Λ_V of $\mathbf{F}_2[X]^l$. If $V = 0$, Λ_V is the kernel of (2), and we will denote it by K_l. Clearly, $K_l = K_1^l$, and K_1 is an ideal of $\mathbf{F}_2[X]$. This ideal is generated by a degree d polynomial, $P_{\mathrm{ch}}(X)$. We define the *absolute value* of $P(X) \in \mathbf{F}_2[X]$ to be 2^δ if δ is the degree of $P(X)$, and the *length* (resp. *degree*) of $(P_1(X), \ldots, P_l(X)) \in \mathbf{F}_2[X]^l$ to be the maximum absolute value (resp. degree) of the components.

If Λ is any sublattice of $\mathbf{F}_2[X]^l$, its *fundamental volume* $|\Lambda|$ is the absolute value of the determinant of any one of its bases, and we have

$$|\Lambda| = \prod_{i=1}^{l} \sigma_i(\Lambda), \tag{3}$$

where $\sigma_i(\Lambda)$ is the length of the ith vector of a Minkowski-reduced basis of Λ. The fundamental volume $|\Lambda|$ is also equal to the group theoretical index $[\mathbf{F}_2[X]^l : \Lambda]$ which, in case $\Lambda = \Lambda_V$, is simply $[F^l : V]$. For instance Tezuka's lattice Λ_l is equal to $\Lambda_{V^{(l)}}$ with $V^{(l)} = F \cdot (y_1, \ldots, y_l)$ (see def. 3 of [7]), and its fundamental volume is thus equal to $2^{d(l-1)}$.

We propose to use instead of Λ_l, the lattice Λ_l', given by $\Lambda_{W^{(l)}}$, where we take $W^{(l)}$ to be the ortho-complement of $V^{(l)}$ with respect to the standard F-bilinear scalar product defined, for $v = (x_1, \ldots, x_l)$ and $v' = (x_1', \ldots, x_l') \in F^l$, by

$$\langle v, v' \rangle = \sum_{i=1}^{l} x_i x_i'. \tag{4}$$

The fundamental volume of Λ_l' is equal to 2^d, and is thus much smaller than that of Λ_l' unless l is small. Because of (3), a lattice with a smaller fundamental volume will have, in the mean, smaller successive minima. We will show how to take advantage of this in Section 4. Note that the lattices Λ_l and Λ_l' depend only on the first l values of the sequence $y_1, \ldots, y_w$. We will occasionally indicate this dependence by writing $\Lambda_l(y_1, \ldots, y_l)$ and $\Lambda_l'(y_1, \ldots, y_l)$ respectively.

We will denote by C_k the set of all $(P_1(X), \ldots, P_l(X)) \in \mathbf{F}_2[X]^l$ of length smaller than 2^k. The following lemma establishes further connections between a subspace $V \subset F^l$ and the lattice Λ_V.

Lemma 1. (i) *The restriction of (2) to C_d is one to one, and its image is F^l.*
(ii) *For $0 \leq k \leq d$, (2) maps C_k onto F_k^l.*
(iii) *For any F-linear subspace V of F^l, (2) maps $\Lambda_V \cap C_d$ onto V.*

From this, and Theorem 2 of [1] we obtain, for any F-linear subspace V of F^l,

$$\dim_{\mathbf{F}_2}(V \cap F_k^l) = \sum_{i=1}^{l} (k - \lg \sigma_i(\Lambda_V))^+, \quad 0 \leq k \leq d. \tag{5}$$

3. The kernel of the adjoint

The rank of (1) is equal to $kl - \dim_{\mathbf{F}_2} R_{l,k}$, where $R_{l,k}$ denotes the vector space over $\mathbf{F}_2$ of all systems $(\alpha_{i,j})_{i,j} \in {\mathbf{F}_2}^{lk}$, $1 \leq i \leq l$, $0 \leq j < k$ such that

$$\sum_{i,j} \alpha_{i,j}\phi(y_i x^j y) = 0, \quad y \in F. \tag{6}$$

Since the rank of (1) does not depend on the choice of ϕ, we will take it to be that $\mathbf{F}_2$-linear form over F which has its kernel equal to F_{d-1}. The image of $R_{l,k}$ by the correspondence ${\mathbf{F}_2}^{lk} \to F^l$ given by

$$(\alpha_{i,j})_{i,j} \mapsto (\sum_j \alpha_{i,j}x^j)_i \tag{7}$$

can then be described as follows. We define, in addition to the standard scalar product (4), a $\mathbf{F}_2$-bilinear scalar product by

$$\langle v, v'\rangle_2 = \phi(\langle v, v'\rangle), \quad v, v' \in F^l. \tag{8}$$

Note that the orthocomplement of an F-subspace of F^l is the same for both scalar products (4) and (8). Thus, $W^{(l)}$ is also the orthocomplement of $V^{(l)}$ with respect to (8).

Lemma 2. *For $k \leq d$, the restriction of (7) to $R_{l,k}$ is one to one, and onto $W^{(l)} \cap F^l_k$.*

Proof. First, the image of ${\mathbf{F}_2}^{lk}$ by (7) is F^l_k. From (6), a system $(\alpha_{i,j})_{i,j} \in {\mathbf{F}_2}^{lk}$ belongs to $R^{(l)}_k$ if and only if $(\sum_j \alpha_{i,j}x^j)_i$ is orthogonal to $V^{(l)}$ with respect to (8), that is, if and only if $(\sum_j \alpha_{i,j}x^j)_i$ belongs to $W^{(l)}$. The Lemma follows. □

The main result shows how the computation of the rank of (1) is reduced to the computation of the quantities $\sigma_i(\Lambda'_l)$.

Theorem 1. *The rank of (1) is equal to*

$$lk - \sum_{i=1}^{l}(k - \lg \sigma_i(\Lambda'_l))^+, \quad 0 \leq k \leq d. \tag{9}$$

Proof. This follows from (5) and Lemma 2. □

The quantities $\sigma_i(\Lambda'_l)$ can be computed by applying the Lenstra reduction algorithm [5] to a suitably chosen basis of Λ'_l. We digress shortly to establish a remarkable connection between the quantities $\sigma_i(\Lambda_l)$, and $\sigma_i(\Lambda'_l)$. This is closely connected to a result of Mahler (see §10 of [6]). We first establish the following relation.

Proposition 1.

$$\dim_{\mathbf{F}_2}(V^{(l)} \cap F^l_{d-k}) - \dim_{\mathbf{F}_2}(W^{(l)} \cap F^l_k) = d - lk, \quad 1 \leq k \leq d. \tag{10}$$

Proof. We have

$$\begin{aligned} \dim_{\mathbf{F}_2}(V^{(l)} + F^l_{d-k}) + \dim_{\mathbf{F}_2}(V^{(l)} \cap F^l_{d-k}) &= \dim_{\mathbf{F}_2} V^{(l)} + \dim_{\mathbf{F}_2} F^l_{d-k} \\ &= d + (d-k)l. \end{aligned}$$

Since F^l_k (resp. $W^{(l)}$) is the orthocomplement of F^l_{d-k} (resp. $V^{(l)}$) with respect to (8), we also have

$$\dim_{\mathbf{F}_2}(W^{(l)} \cap F^l_k) + \dim_{\mathbf{F}_2}(V^{(l)} + F^l_{d-k}) = dl.$$

The proposition follows by combining these two equations. □

Corollary 1. *We have, for* $1 \le i \le l$,

$$\lg \sigma_i(\Lambda'_l) + \lg \sigma_{l-i+1}(\Lambda_l) = d, \quad 1 \le i \le l.$$

Proof. We abreviate $\lg \sigma_i(\Lambda_l)$ to s_i, and $\lg \sigma_i(\Lambda'_l)$ to s'_i. Using (5), we can then write (10) as

$$\sum_{i=1}^{l}(d-k-s_i)^+ - \sum_{i=1}^{l}(k-s'_i)^+ = d - lk.$$

Combining this with

$$\sum_{i=1}^{l}(k-s'_i)^+ - \sum_{i=1}^{l}(s'_i-k)^+ = \sum_{i=1}^{l}(k-s'_i) = lk - d,$$

we obtain

$$\sum_{i=1}^{l}((d-s_i)-k)^+ - \sum_{i=1}^{l}(s'_i-k)^+ = 0.$$

Since $0 \le s_i, s'_i \le d$, this implies that, for $0 \le k \le d$, the sets $\{i \mid s'_i = k\}$ and $\{i \mid d - s_i = k\}$ have the same cardinality. The statement of the Corollary follows. □

4. Recursivity

From theorem 1 and its corollary, the rank of (1) can be obtained, simultaneously for all k, by computation of the quantities $\sigma_i(\Lambda_l)$, or $\sigma_i(\Lambda'_l)$. This is achieved by use of Lenstra's reduction algorithm [5], applied to a suitable basis of Λ_l or Λ'_l and, as we shall now show, it is advantageous for this to use the latter lattice rather than the former. Assume $1 < l \le w$. The F-linear mappings $\iota : F^{l-1} \to F^l$, and $\rho : F^l \to F^{l-1}$, defined by addition of a lth coordinate taken equal to zero, and deletion of the lth coordinate respectively, are mutually adjoint, that is,

$$\langle \iota(w), v\rangle = \langle w, \rho(v)\rangle, \quad w \in F^{l-1}, v \in F^l. \tag{11}$$

Lemma 3. *For* $1 < l \le w$, *we have*

(i) $\rho(V^{(l)}) = V^{(l-1)}$.
(ii) $W^{(l)} = \iota(W^{(l-1)}) \oplus F(y_l, 0, \ldots, 0, y_1)$.

Proof. Statement (i) is immediate from the definition of $V^{(l)}$. To prove (ii), notice that (11) implies that $\iota(W^{(l-1)})$ is an F-linear subspace of $W^{(l)}$. It is in fact of codimension 1 in $W^{(l)}$, since it has dimension $l-1$, while $W^{(l)}$ has dimension l. The statement now follows since $(y_l, 0, \ldots, 0, y_1)$ belongs to $W^{(l)} \setminus \iota(W^{(l-1)})$. □

We deduce from this lemma, the recursivity properties of the lattices Λ_l and Λ'_l. Denote again by ι and ρ the similarly defined $\mathbf{F}_2[X]$-linear mappings $\iota : \mathbf{F}_2[X]^{l-1} \to$

$\mathbf{F}_2[X]^l$, and $\rho : \mathbf{F}_2[X]^l \to \mathbf{F}_2[X]^{l-1}$. Take, $Q_i(X) \in \mathbf{F}_2[X]$ of degree less than d, and such that $y_1 Q_i(x) = y_i$, $2 \leq i \leq l$.

Proposition 2. *For $1 < l \leq w$, we have*

(i) $\rho(\Lambda_l) = \Lambda_{l-1}$.
(ii) $\Lambda'_l = \iota(\Lambda'_{l-1}) \oplus \mathbf{F}_2[X](Q_l(X), 0, \ldots, 0, 1)$.

Proof. Note that ι and ρ commute with (2). Therefore statement (i) of Lemma 3 implies our first statement. Also, since the vector $(Q_l(X), 0, \ldots, 0, 1)$ is mapped by (2) to the vector $(y_l/y_1, 0, \ldots, 0, 1)$, statement (ii) of Lemma 3 implies that

$$\Lambda'_l = \iota(\Lambda'_{l-1}) + \mathbf{F}_2[X](Q_l(X), 0, \ldots, 0, 1) + K_l.$$

But $K_l = \iota(K_{l-1}) + \mathbf{F}_2[X](0, \ldots, 0, P_{\mathrm{ch}}(X))$ so that our second statement follows from the previous equation. □

The starting point for the Lenstra reduction algorithm is a lattice basis B for an l-dimensional sublattice Λ of $\mathbf{F}_2[X]^l$. The algorithm transforms this basis into another basis of Λ, which is *Lenstra-reduced* and, in particular, Minkowski-reduced. We associate with the basis B the quantities $d_{\mathrm{s}}(B)$ and $d_{\mathrm{m}}(B)$, which are defined respectively as the sum and the maximum of the basis vector degrees. The storage requirement for the algorithm is then measured by $l d_{\mathrm{s}}(B)$, and an upper bound for the execution time (the required number of bit operations) is given by

$$C l^3 d_{\mathrm{m}}(B)(d_{\mathrm{s}}(B) - \lg|\Lambda| + 1), \tag{12}$$

for some absolute constant C (see prop. 1.14 of [5]).

In case of Λ_l, one uses the basis B_l composed of the vector $(1, Q_2(X), \ldots, Q_l(X))$, and $P_{\mathrm{ch}}(X)\delta_j^{(l)}$, $2 \leq j \leq l$, where $\delta_j^{(l)} \in \mathbf{F}_2[X]^l$ has all its components equal to 0, except for the jth which is equal to 1. In case of Λ'_l we may, by (ii) of Proposition 2, take a basis B'_l, composed of the images by ι of the vectors belonging to a Lenstra-reduced basis of Λ'_{l-1}, and of the vector $(Q_l(X), 0, \ldots, 0, 1)$. The required space to reduce the basis B'_l is significantly less than for B_l, as we see from the following lemma.

Lemma 4. *We have*

(i) $(l-1)d \leq d_{\mathrm{s}}(B_l) \leq ld - 1$,
(ii) $d \leq d_{\mathrm{s}}(B'_l) \leq 2d - 1$,

Proof. Statement (i) follows from the fact that $P_{\mathrm{ch}}(X)$ has degree equal to d, while all $Q_i(X)$ have it less than d. Using (3) we obtain that the sum of the degrees of the first $l-1$ vectors of B'_l is equal to d, and this proves statement (ii). □

We say that an l-dimensional lattice $\Lambda \subset \mathbf{F}_2[X]^l$ is *regular* if

$$\sigma_l(\Lambda)/\sigma_1(\Lambda) \leq 2.$$

Clearly the rank of (1) is bounded by $\min(d, lk)$.

Proposition 3. *For a given l, the rank of (1) is equal to* $\min(d, lk)$ *for all k if and only if Λ'_l is regular.*

Proof. By Theorem 1, when $lk \leq d$ (resp. $lk > d$), the rank of (1) is equal to lk (resp. d) if and only if, for all i, $\lg\sigma_i(\Lambda'_l) \geq k$ (resp. $\lg\sigma_i(\Lambda'_l) \leq k$). Thus the rank of (1) is equal to $\min(d, lk)$ for all k if and only if

$$[d/l] \le \lg \sigma_i(\Lambda_l') \le [d/l] + 1, \quad 1 \le i \le l.$$

But, this is equivalent to $\lg \sigma_l(\Lambda_l') - \lg \sigma_1(\Lambda_l') \le 1$ since we have, from (3), that $\sum_{i=1}^{l} \lg \sigma_i(\Lambda_l') = d$. □

Note, by Corollary 1, the equivalence of the regularity of the lattices Λ_l and Λ_l'.

Theorem 2. *If the lattice Λ_{l-1}' is regular, then the Lenstra basis reduction algorithm applied to the basis B_l' has running time not exceeding*

$$C_l' l(d + l - 1)^2, \quad l \ge 2,$$

where $C_l' = (l/(l-1))^2 C + 1/l$, and C is the constant appearing in (12).

Proof. Since Λ_{l-1}' is assumed regular, the first $l-1$ vectors of B_l' have their degree bounded by $d/(l-1) + 1$. In a first phase, the algorithm will reduce (in length) the lth vector by the repeated operation of adding to it one of the first $l-1$ vectors, pre-multiplied by a suitable power of X. Each such operation requires at most $d/(l-1) + 2$ bit operations. We thus need at most $d + 2l - 2$ bit operations to diminish by 1 the degree of the lth vector and at most

$$\left(\frac{l-2}{l-1}\right) d(d + 2l - 2), \tag{13}$$

to diminish its degree to a value bounded by $d/(l-1)$. After termination of this first phase, the algorithm terminates, according to (12), using at most

$$Cl^3 \left(\frac{d}{l-1} + 1\right)^2 \tag{14}$$

further bit operations. The sum of (13) and (14) is bounded by $C_l' l(d+l-1)^2$, and the theorem follows. □

For given F, $x \in F$, and a subset $E \subset F^w$, it is a problem of interest to determine $(y_1, \ldots, y_w) \in E$, such that the rank of (1) is equal to $\min(d, lk)$ for all $l \le w$, and all $k \le d$, that is, such that the lattices $\Lambda_l(y_1, \ldots, y_l)$ (or, equivalently, $\Lambda_l'(y_1, \ldots, y_l)$) are regular for all $l \le w$. This type of question arises when one wants to construct an optimally equidistributed output mapping $\Phi(y) = (\phi(y_1 y), \ldots, \phi(y_w y))$ for a generator based on the field F. Consider the rooted tree $T = T(E)$ whose vertices of depth l (or l-vertices for short) are those l-tuples $(y_1, \ldots, y_l) \in F^l$ for which there exists $y_{l+1}, \ldots, y_w$ such that $(y_1, \ldots, y_w) \in E$, and whose edges link an $(l-1)$-vertex to an l-vertex if and only if these have the same first $l-1$ components. We associate with an l-vertex the lattices $\Lambda_l = \Lambda_l(y_1, \ldots, y_l)$ and $\Lambda_l' = \Lambda_l'(y_1, \ldots, y_l)$. We will say that an l-vertex $(y_1, \ldots, y_l)$ of T is *regular* if its associated lattice Λ_l (or, equivalently, Λ_l') is regular. A *regular* path in T is a path visiting only regular vertices. One may then reformulate our problem as the determination of a regular path in T joining the root to an w-vertex.

For any l-vertex $(y_1, \ldots, y_l)$ of T we may, as above, construct lattice bases B_l and B_l' for the associated lattices Λ_l and Λ_l'. We denote them by $B_l(y_1, \ldots, y_l)$ and $B_l'(y_1, \ldots, y_l)$, respectively. The regularity of an l-vertex $(y_1, \ldots, y_l)$ can be determined by application of Lenstra's basis reduction algorithm, either to $B_l(y_1, \ldots, y_l)$ or to $B_l'(y_1, \ldots, y_l)$. If we use $B_l(y_1, \ldots, y_l)$ then, according to (12), the execution time does not exceed $Cl^3 d^2$. It does not exceed $C_l' l(d+l-1)^2$ ($\sim Cld^2$ for l and d/l

large), according to Theorem 2, if we use $B'_l(y_1,\ldots,y_l)$ instead, and if the $(l-1)$-vertex $(y_1,\ldots,y_{l-1})$ is regular. Obviously, in the latter case, one needs a Lenstra-reduced basis of the lattice Λ'_{l-1} associated with the $(l-1)$-vertex $(y_1,\ldots,y_{l-1})$, but such a basis is already available when constructing a regular path, visiting a regular $(l-1)$-vertex before any adjacent l-vertex. Memorizing a reduced basis of Λ'_{l-1} for a regular $(l-1)$-vertex also permits one to verify the regularity of several l-vertices adjacent to it, without recomputing the reduced basis. We finally note that, given a regular path of length $l-1$ and an adjacent l-vertex $(y_1,\ldots,y_l)$, the regularity of the latter can be obtained by successively constructing and reducing (by Lenstra's algorithm), the bases $B'_2(y_1,y_2),\ldots,B'_l(y_1,\ldots,y_l)$, in a time which, by Theorem 2, does not exceed $C''_l(l^2/2)(d+l-1)^2$ ($\sim C(l^2/2)d^2$, for l and d/l large). Here the constants C''_l are given by

$$C''_l=\left(1+\frac{5}{l}+\frac{6\ln l+4}{l^2}\right)C+\frac{2(l-1)}{l^2}.$$

5. Computation of a random regular path

The advantage of using the lattices Λ'_l instead of Λ_l is confirmed by extensive computer experiments. We give a typical illustration. We take F to be the field of degree 19937 over $\mathbf{F}_2$, and $x\in F$ to be a root of $P_{\mathrm{ch}}(X)=X^{19937}+X^{9842}+1$. This trinomial is primitive (see the table in [3]). Let $w=32$ and $E=(F\backslash\{0\})^w$. We seek to determine a regular path in $T(E)$, recursively. Having found a regular $(l-1)$-vertex $(y_1,\ldots,y_{l-1})$, a regular l-vertex $(y_1,\ldots,y_l)$ is determined by randomly choosing $y\in F\backslash\{0\}$, each outcome being equaly likely, and taking for y_l the first value of y for which the vertex $(y_1,\ldots,y_{l-1},y)$ is regular. The regularity is determined by using either of the lattices $\Lambda_l(y_1,\ldots,y_{l-1},y)$ and $\Lambda'_l(y_1,\ldots,y_{l-1},y)$. In the first case, Lenstra's reduction algorithm is applied to the basis $B_l(y_1,\ldots,y_{l-1},y)$, while in the second case it is applied to the basis $B'_l(y_1,\ldots,y_{l-1},y)$ constructed with the help of the previously reduced basis for the lattice $\Lambda'(y_1,\ldots,y_{l-1})$.

For each value of l, from 2 to 32, the CPU time (in seconds) for the reduction required at the l-vertex, and the total cumulative CPU time to determine the first l vertices, are recorded in Table 1. In most cases, the first y that was tried already gave a regular vertex. When more than one value of y was needed, their number is indicated in parentheses, and the reduction time given is the mean reduction time for all these values of y. Since in both computations the same values of y are used, the same regular path is determined. It appears from this table that the reduction itself takes almost all of the CPU time, and that it is always much quicker to determine the regularity of a vertex using the lattice Λ'_l than Λ_l. In this instance, there is as much as a 10-fold time reduction for dimension $l=18$, and this increases with l up to a 16-fold time reduction for $l=32$.

Here, we have taken $E=(F\backslash\{0\})^w$. When dealing with the problem of constructing an output mapping $\Phi(y)=(\phi(y_1y),\ldots,\phi(y_wy))$ for some generator based on the field F, one must choose E such that each of its members $(y_1,\ldots,y_w)$ defines an *efficient* mapping Φ, when viewed as depending on a computer memory image of the state of the generator (i.e., an output mapping for which a fast computer implementation is available). A description of a specific case, with a new class of random number generators, will be the subject of a forthcoming paper.

TABLE 1. Efficiency comparison for a random regular path. The first column under Λ_l (resp. Λ'_l) gives the (mean) reduction time, and the second one, the total cumulative execution time.

l	Λ_l		Λ'_l	
2	.76	.84	.77	.83
(2)3	3.04	7.04	1.66	4.27
(2)4	6.77	20.70	2.57	9.53
(4)5	11.80	68.16	3.51	23.81
6	18.04	86.26	4.50	28.37
(2)7	25.68	137.74	5.54	39.58
8	34.68	172.49	6.98	46.62
9	44.61	217.16	7.75	54.43
10	55.62	272.84	8.99	63.49
11	68.37	341.27	10.29	73.84
12	82.09	423.42	11.41	85.31
13	97.00	520.49	12.70	98.08
(3)14	115.38	866.82	14.11	140.62
15	137.35	1004.23	15.81	156.50
(2)16	159.01	1322.37	17.40	191.44
17	183.66	1506.09	18.74	210.25
18	209.16	1715.32	20.27	230.59
19	237.23	1952.62	22.06	252.73
20	266.26	2218.95	23.48	276.29
21	298.51	2517.53	26.06	302.42
(2)22	331.43	3180.54	26.94	356.46
(2)23	366.08	3912.84	28.93	414.48
24	401.14	4314.05	30.84	445.41
25	438.63	4752.76	31.91	477.41
26	478.44	5231.28	33.83	511.32
27	520.87	5752.23	35.91	547.32
28	560.36	6312.67	39.02	586.44
29	602.05	6914.81	40.60	627.14
30	649.19	7564.08	42.10	669.33
31	696.55	8260.72	43.93	713.36
32	742.96	9003.76	46.75	760.21

REFERENCES

[1] R. Couture, P. L'Ecuyer, and S. Tezuka, *On the distribution of k-dimensional vectors for simple and combined Tausworthe sequences*, Mathematics of Computation **60** (1993), no. 202, 749–761, S11–S16.

[2] M. Fushimi and S. Tezuka, *The k-distribution of generalized feedback shift register pseudorandom numbers*, Communications of the ACM **26** (1983), no. 7, 516–523.

[3] J. R. Heringa, H. W. J. Blöte, and A. Compagner, *New primitive trinomials of Mersenne-exponent degrees for random-number generation*, International Journal of Modern Physics C **3** (1992), no. 3, 561–564.

[4] P. L'Ecuyer, *Maximally equidistributed combined Tausworthe generators*, Mathematics of Computation **65** (1996), no. 213, 203–213.

[5] A. K. Lenstra, *Factoring multivariate polynomials over finite fields*, Journal of Computer and System Sciences **30** (1985), 235–248.

[6] K. Mahler, *An analogue to Minkowski's geometry of numbers in a field of series*, Annals of Mathematics **42** (1941), no. 2, 488–522.

[7] S. Tezuka, *The k-dimensional distribution of combined GFSR sequences*, Mathematics of Computation **62** (1994), no. 206, 809–817.

Département d'Informatique et de Recherche Opérationnelle, Université de Montréal, C.P. 6128, Succ. Centre-Ville, Montréal, H3C 3J7, Canada

E-mail address: couture@iro.umontreal.ca and lecuyer@iro.umontreal.ca