Efficient Jump Ahead for $\mathbb{F}_{2}$-Linear Random Number Generator<br>H. Haramoto, M. Matsumoto, T. Nishimura, F. Panneton, P. L'Écuyer<br>G-2006-62<br>October 2006

Les textes publiés dans la série des rapports de recherche HEC n'engagent que la responsabilité de leurs auteurs. La publication de ces rapports de recherche bénéficie d'une subvention du Fonds québécois de la recherche sur la nature et les technologies.

# Efficient Jump Ahead for $\mathbb{F}_{2}$-Linear Random Number Generators 

Hiroshi Haramoto, Makoto Matsumoto<br>Department of Mathematics<br>Hiroshima University<br>Kagamiyama 1-3-1<br>Higashi-Hiroshima, Hiroshima 739-8526, Japan<br>m-mat@math.sci.hiroshima-u.ac.jp<br>\section*{Takuji Nishimura}<br>Department of Mathematics<br>Yamagata University<br>Yamagata 990-8586, Japan<br>nisimura@sci.kj.yamagata-u.ac.jp<br>François Panneton, Pierre L'Écuyer*<br>Département d'informatique et de recherche opérationnelle<br>Université de Montréal<br>C.P. 6128, Succ. Centre-ville<br>Montréal (Québec) Canada, H3C 3J7<br>\{panneton;lecuyer\}@iro.umontreal.ca<br>* and GERAD

October 2006

Les Cahiers du GERAD
G-2006-62
Copyright © 2006 GERAD


#### Abstract

The fastest long-period random number generators currently available are based on linear recurrences modulo 2. So far, software that provides multiple disjoint streams and substreams has not been available for these generators because of the lack of efficient jump-ahead facilities. In principle, it suffices to multiply the state (a $k$-bit vector) by an appropriate $k \times k$ binary matrix to find the new state far ahead in the sequence. However, when $k$ is large (e.g., for a generator such as the popular Mersenne twister, for which $k=19937$ ), this matrix-vector multiplication is slow and a large amount of memory is required to store the $k \times k$ matrix. In this paper, we provide a faster algorithm to jump ahead by a large number of steps in a linear recurrence modulo 2 . The method uses much less than the $k^{2}$ bits of memory required by the matrix method. It is based on polynomial calculus modulo the characteristic polynomial of the recurrence and uses a sliding window algorithm for the multiplication.


Key Words: simulation, random number generation, jumping ahead, multiple streams.

## Résumé

Les générateurs de nombres aléatoires à longue période qui sont les plus rapides sont basés sur une récurrence linéaire modulo 2. À ce jour, il n'y a pas de logiciel qui offre de flots disjoints pour ces générateurs à cause de l'inefficacité des méthodes qui permettent de sauter en avant dans la récurrence. En principe, il suffit de multiplier l'état (un vecteur de $k$ bits) par une matrice binaire $k \times k$ pour obtenir un nouvel état qui se trouve beaucoup plus loin dans la séquence. Malheureusement, quand $k$ est grand (par exemple, pour un générateur comme le populaire Mersenne twister pour lequel $k=19937$ ), cette multiplication matrice-vecteur est lente et une quantité importante de mémoire est nécessaire pour emmagasiner la matrice $k \times k$. Dans cet article, nous introduisons un algorithme plus rapide pour sauter en avant par un grand nombre d'itérations dans une récurrence linéaire modulo 2. La méthode utilise beaucoup moins que les $k^{2}$ bits de mémoire requis par la méthode matricielle. Elle est basée sur l'arithmétique dans l'extension du corps fini à deux éléments défini par le polynôme caractéristique de la récurrence et utilise un algorithme à fenêtre glissante pour la multiplication.

Acknowledgments: This study is partially supported by JSPS/Ministry of Education Grant-in-Aid for Scientific Research No. 18654021 and 16204002, JSPS Core-toCore Program No. 18005, NSERC-Canada grant No. ODGP0110050, and a Canada Research Chair to the last author. The paper was written while the last author was enjoying hospitality of IRISA-INRIA in Rennes, France.

## 1 Introduction

Random number generators (RNGs) with multiple disjoint streams and substreams are an important component of any good general-purpose simulation or statistical software. They are very handy, for example, to obtain parallel RNGs and to support the implementation of variance reduction techniques (Kelton, 2006; Law and Kelton, 2000; L'Ecuyer et al., 2002). The most convenient way of getting these streams and substreams is to start with a backbone RNG having a huge period, and partition its output sequence into long disjoint subsequences and subsubsequences whose starting points are at equidistant lags. When a new stream is needed, we find its starting point by jumping ahead from the starting point of the current subsequence to the starting point of the next one. Substreams are obtained from subsequences in a similar way. To make sure that no overlap occurs, the streams and substreams must be very long, so that they cannot be exhausted even with days of computing time. To implement this, we need to know how to quickly jump ahead by large lags in the sequence of numbers produced by the generator.

Most generators used for simulation are based on linear recurrences. For these generators, the state $\mathbf{x}_{n}$ at step $n$ is a vector of $k$ integers in $\{0, \ldots, m-1\}$ for some integer $m$ called the modulus, and it evolves as $\mathbf{x}_{n}=\mathbf{A} \mathbf{x}_{n-1} \bmod m$ where $\mathbf{A}$ is a $k \times k$ matrix with elements in $\{0, \ldots, m-1\}$. To jump ahead by $\nu$ steps from any state $\mathbf{x}_{n}$, regardless of how large is $\nu$, it suffices to precompute the matrix $\mathbf{A}^{\nu} \bmod m$ (once for all) and then compute $\mathbf{x}_{n+\nu}=\left(\mathbf{A}^{\nu} \bmod m\right) \mathbf{x}_{n} \bmod m$ by a simple matrix-vector multiplication. This technique is used to provide streams and substreams in the random number package of L'Ecuyer et al. (2002), based on combined multiple recursive generators (CMRG), and which has been adopted in several simulation and statistical software products such as Arena, Automod, Witness, SSJ, SAS, etc.

There are faster generators than the CMRG, based on linear recurrences modulo 2, with extremely long periods and good statistical properties. The Mersenne twister and the WELL (Matsumoto and Nishimura, 1998; Panneton et al., 2006), for example, belong to that class. But efficient software that provides multiple disjoint streams and substreams for them is lacking. Since these generators are linear, the technique just described applies in principle (with $m=2$ ). But the matrix-vector multiplication is slow if implemented naively and an excessive amount of memory is required to store the matrix when $k$ is very large, which is typical. For example, the Mersenne twister generator has $k=19937$. Then, the $k \times k$ binary matrix occupies 47.4 MBytes of memory!

We propose a more efficient technique to perform this multiplication. It uses a representation of the recurrence in a space of polynomials. For a given step size $\nu$, the state $\mathbf{x}_{n+\nu}$ is expressed as a polynomial in $\mathbf{A}$ of degree less than $k$, say $g(\mathbf{A})$, multiplied by $\mathbf{x}_{n}$. A key ingredient is that we use the implementation of the original recurrence (i.e., of the generator) to compute the product $g(\mathbf{A}) \mathbf{x}_{n}$. The most expensive operations in this computation turn out to be the $k$-bit vector additions modulo 2 . We use a sliding window technique to reduce the number of these additions, e.g., by a factor of about 4 when $k=19937$. For
this particular value of $k$, with the proposed method, the generator can jump ahead for an arbitrary lag in less than 5 milliseconds on a 32 -bit 3.0 GHz computer.

The next section gives a framework for $\mathbb{F}_{2}$-linear generators and states the jump-ahead problem. In Section 3, we examine how to jump ahead, explain our proposed technique, and analyze its computational efficiency. The algorithm is stated in Section 3.4 and some timings are given in Section 3.5.

## $2 \quad \mathbb{F}_{2}$-Linear Generators

All along this paper, arithmetic operations are assumed to be performed in $\mathbb{F}_{2}$, the finite field with two elements, represented as 0 and 1 . This corresponds to doing arithmetic modulo 2. Note that in $\mathbb{F}_{2}$, subtraction and addition are equivalent; so we can always write "+" instead of "-" and we do so everywhere in this paper. The RNGs considered obey the general $\mathbb{F}_{2}$-linear recurrence

$$
\begin{equation*}
\mathbf{x}_{n}=\mathbf{A} \mathbf{x}_{n-1} \tag{1}
\end{equation*}
$$

where $\mathbf{x}_{n}=\left(x_{n, 0}, \ldots, x_{n, k-1}\right)^{\mathrm{t}} \in \mathbb{F}_{2}^{k}$ is the $k$-bit state vector at step $n$ and $\mathbf{A}$ is the $k \times k$ transition matrix with elements in $\mathbb{F}_{2}$. The output can be defined by any transformation $\mathbf{x}_{n} \mapsto u_{n}$; the exact form of this transformation is irrelevant for the remainder of the paper. Usually, the output $u_{n} \in[0,1)$ at step $n$ is defined by $u_{n}=\sum_{\ell=1}^{w} y_{n, \ell-1} 2^{-\ell}$ for some positive integer $w$, where $\mathbf{y}_{n}=\left(y_{n, 0}, \ldots, y_{n, w-1}\right)^{\mathrm{t}}=\mathbf{B} \mathbf{x}_{n}$ and $\mathbf{B}$ is a $w \times k$ matrix with elements in $\mathbb{F}_{2}$. Several types of RNGs fit this framework, including the Tausworthe or linear feedback shift register (LFSR), the generalized feedback shift register (GFSR), the twisted GFSR (TGFSR), the Mersenne twister, the WELL, and xorshift generators (Tezuka, 1995; Matsumoto and Nishimura, 1998; L'Ecuyer and Panneton, 2005; Panneton and L'Ecuyer, 2005; Panneton et al., 2006).

Our aim is to compute

$$
\begin{equation*}
\mathbf{x}_{n+\nu}=\mathbf{A}^{\nu} \mathbf{x}_{n} \tag{2}
\end{equation*}
$$

for a large value of $\nu$, say larger than $2^{100}$ or even more. We assume that $\nu$ is fixed in advance and that (2) must be computed for several arbitrary vectors $\mathbf{x}_{n}$ unknown in advance. This is what we need to implement multiple streams and substreams. The algorithm also works if $\nu$ is not fixed, but then the computationally expensive setup must be repeated each time.

## 3 Jumping Ahead

### 3.1 Matrix Method

A first method to jump ahead is the standard one, described in the introduction: We start by precomputing the matrix $\mathbf{J}=\mathbf{A}^{\nu}$ in $\mathbb{F}_{2}$. By a standard square-and-multiply
exponentiation technique (Knuth, 1998), this requires $O\left(k^{3} \log \nu\right)$ operations, and we need $k^{2}$ bits to store $\mathbf{J}$. Then, whenever jumping-ahead is required from state $\mathbf{x}$, we compute the vector $\mathbf{J} \mathbf{x}$. To obtain the $i$ th element of $\mathbf{J} \mathbf{x}$, we compute the componentwise product of the $i$ th row of $\mathbf{J}$ by the (transposed) vector $\mathbf{x}$, by a bitwise AND, and add the bits of the resulting vector, modulo 2. A straightforward implementation of this on a $w$-bit computer requires $k\lceil k / w\rceil$ AND operations, followed by $k^{2}$ operations to count the bits.

But the work to add the bits modulo 2 can be reduced as follows. Observe that we only need the parity of the sum of bits in the vector, which can be obtained by xoring all its bits. This can be achieved as follows: partition the $k$-bit vector into $w$-bit blocks (this is how it is stored), xor all these blocks together (for a given vector, this requires $\lceil k / w\rceil$ xor operations), then xor the bits in the resulting $w$-bit block ( $w$ operations). The total number of operations with this approach is $2 k\lceil k / w\rceil+k w$.

Nevertheless, for $k=19937$, for instance, storing $\mathbf{J}$ takes around 47.4 megabytes of memory, and computing it by squaring and multiplying the binary matrices is impractical (each squaring takes $O\left(k^{3}\right)$ time).

### 3.2 Using the Polynomial Representation

A more efficient approach, when $k$ is large, works with the polynomial representation of the recurrence, as follows. Write the characteristic polynomial of the matrix A as

$$
p(z)=\operatorname{det}(z \mathbf{I}+\mathbf{A})=z^{k}+\alpha_{1} z^{k-1}+\cdots+\alpha_{k-1} z+\alpha_{k}
$$

where $\mathbf{I}$ is the identity matrix and $\alpha_{j} \in \mathbb{F}_{2}$ for each $j$, and recall that

$$
p(\mathbf{A})=z^{k}+\alpha_{1} \mathbf{A}^{k-1}+\cdots+\alpha_{k-1} \mathbf{A}+\alpha_{k} \mathbf{I}=0
$$

(this is a fundamental property of the characteristic polynomial). Let

$$
\begin{equation*}
g(z)=z^{\nu} \bmod p(z)=a_{1} z^{k-1}+\cdots+a_{k-1} z+a_{k} \tag{3}
\end{equation*}
$$

This $g(z)$ can be computed (once for all) in $O\left(k^{2} \log \nu\right)$ time by the square-and-multiply method (Knuth, 1998) in the space of polynomials modulo $p(z)$. Observe that $g(z)=$ $z^{\nu}+q(z) p(z)$ for some polynomial $q(z)$. Combining this with the fact that $p(\mathbf{A})=0$, we see that $g(\mathbf{A})=\mathbf{A}^{\nu}$ and thus

$$
\mathbf{J}=\mathbf{A}^{\nu}=g(\mathbf{A})=a_{1} \mathbf{A}^{k-1}+\cdots+a_{k-1} \mathbf{A}+a_{k} \mathbf{I} .
$$

Therefore, $\mathbf{J x}$ can be computed by

$$
\begin{align*}
\mathbf{J} \mathbf{x} & =\left(a_{1} \mathbf{A}^{k-1}+\cdots+a_{k-1} \mathbf{A}+a_{k} \mathbf{I}\right) \mathbf{x} \\
& =\mathbf{A}\left(\cdots \mathbf{A}\left(\mathbf{A}\left(\mathbf{A} a_{1} \mathbf{x}+a_{2} \mathbf{x}\right)+a_{3} \mathbf{x}\right)+\cdots+a_{k-1} \mathbf{x}\right)+a_{k} \mathbf{x} \tag{4}
\end{align*}
$$

where the latter represents Horner's method for polynomial evaluation. To compute this, we can simply advance the RNG by $k-1$ steps from state $\mathbf{x}$ and add (by bitwise exclusiveor) the states obtained at the steps that correspond to the nonzero $a_{j}$ 's. This computation requires running the RNG for $k-1$ steps and adding at most ( $k-1$ ) $k$-bit vectors. For a random polynomial $g(z)$ (whose coefficients $a_{1}, \ldots, a_{k}$ are drawn uniformly over the set of all $2^{k}$ possibilities), there is on average $k / 2$ nonzero coefficients, so $(k / 2)-1$ vector additions are required. This computation still demands $O\left(k^{2}\right)$ operations, but only $k$ bits of storage are needed for the coefficients of $g(z)$.

To examine more closely the cost of this implementation, let us suppose that the computer has $w$-bit words and that $g(z)$ has $k / 2$ nonzero coefficients. Each $k$-bit vector addition requires $\eta=\lceil k / w\rceil$ XOR operations on the computer, so we need $((k / 2)-1) \eta \approx k^{2} /(2 w)$ operations to add the vectors if we use a "standard" method.

It is important to recall here that the large-period $\mathbb{F}_{2}$-linear RNGs are normally designed so that $\mathbf{A x}$ can be computed with only a handful of binary operations (such as XORs, shifts, and bit masks). Suppose our RNG needs $c$ such operations at each step. Then we need ( $k-1$ )c operations to advance the RNG by $k-1$ steps. The total jump-ahead cost is thus $(k-1) c+((k / 2)-1) \eta$ operations.

As a typical illustration, take $w=32, k=19937$, and $c=10$. Then $(k-1) c \approx 2.0 \times 10^{5}$ whereas $((k / 2)-1) \eta \approx 6.2 \times 10^{6}$, so the vector additions dominate the cost. Our next improvement will reduce this number of additions in exchange for some additional storage.

### 3.3 Improvement via Decomposition and a Sliding Window

We choose a small positive integer $q$, say somewhere from 4 to 10 . Let $\mathbb{T}_{q}$ be the set of polynomials with coefficients in $\mathbb{F}_{2}$ and of degree exactly $q$, i.e., of the form $h(z)=$ $z^{q}+b_{1} z^{q-1}+\cdots+b_{q}$ where the $b_{j}$ 's are in $\mathbb{F}_{2}$. This set has cardinality $2^{q}$. We decompose $g(z)$ as

$$
\begin{equation*}
g(z)=h_{1}(z) z^{d_{1}}+\cdots+h_{m}(z) z^{d_{m}}+h_{m+1}(z)+z^{q} \tag{5}
\end{equation*}
$$

where $h_{j}(z) \in \mathbb{T}_{q}$ for $j=1, \ldots, m+1,0 \leq d_{m}<\cdots<d_{1}<k$, and $m$ is as small as possible. This decomposition is obtained as follows. We write the coefficients of $g(z)$ in a sequence:

$$
a_{1} a_{2} a_{3} \cdots a_{k-1} a_{k} .
$$

If $i_{1}=\min \left\{i>0: a_{i}=1\right\}$ is the index of the first nonzero coefficient in the sequence, we set $d_{1}=k-q-i_{1}$, we define

$$
h_{1}(z)=z^{q}+a_{i_{1}+1} z^{q-1}+\cdots+a_{i_{1}+q},
$$

and we remove $a_{1}, \ldots, a_{i_{1}+q}$ from the sequence. Then we repeat the same process with the sequence that starts with $a_{i_{1}+q+1}$ to define $d_{2}$ and $h_{2}(z)$, and so on. In general, if
$i_{j}=\min \left\{i>i_{j-1}+q: a_{i}=1\right\}$ is the index of the first nonzero element in the sequence $a_{i_{j-1}+q+1}, \ldots, a_{k}$ and if $k-i_{j} \geq q$, we put $d_{j}=k-q-i_{j}$ and

$$
h_{j}(z)=z^{q}+a_{i_{j}+1} z^{q-1}+\cdots+a_{i_{j}+q} .
$$

As soon as $k-i_{j}<q$, we put $m=j-1$ and

$$
h_{m+1}(z)=z^{q}+a_{i_{j}} z^{k-i_{j}}+\cdots+a_{k} .
$$

This completes the decomposition (5). Note that this decomposition does not depend on x .

Now, from (5), we can rewrite

$$
\begin{align*}
\mathbf{J} \mathbf{x}= & g(\mathbf{A}) \mathbf{x} \\
= & \mathbf{A}^{d_{m}}\left(\cdots\left(\mathbf{A}^{d_{2}-d_{3}}\left(\mathbf{A}^{d_{1}-d_{2}} h_{1}(\mathbf{A}) \mathbf{x}+h_{2}(\mathbf{A}) \mathbf{x}\right)+h_{3}(\mathbf{A}) \mathbf{x}\right)+\cdots\right. \\
& \left.\quad+h_{m}(\mathbf{A}) \mathbf{x}\right)+h_{m+1}(\mathbf{A}) \mathbf{x}+\mathbf{A}^{q} \mathbf{x} \tag{6}
\end{align*}
$$

To compute $\mathbf{J} \mathbf{x}$, we first compute the vectors $h(\mathbf{A}) \mathbf{x}$ for all $2^{q}$ polynomials $h(z)$ in $\mathbb{T}_{q}$, and store these vectors in a table. Then, we start the RNG from state $h_{1}(\mathbf{A}) \mathbf{x}$, advance it by $d_{1}-d_{2}$ steps, add $h_{2}(\mathbf{A}) \mathbf{x}$ to its state, advance it by $d_{2}-d_{3}$ steps, $\ldots$, add $h_{m}(\mathbf{A}) \mathbf{x}$ to the state, advance the RNG by $d_{m}$ steps, and finally add $h_{m+1}(\mathbf{A}) \mathbf{x}+\mathbf{A}^{q} \mathbf{x}$. We still need to advance the RNG by a total of $k-1$ steps, but $m+1$ vector additions now suffice instead of $k / 2$, where $m \leq\lceil k /(q+1)\rceil$.

This method is a direct adaptation of the sliding window algorithm used for exponentiation in a group (Möller, 2005). We illustrate it by an example.

Example 1 Let $k=18, q=3$, and suppose that the coefficients of $g(z)$ are the following:


In this case, we have $i_{1}=3, d_{1}=15-i_{1}=12, h_{1}(z)=z^{3}+z^{2}+z+1, i_{2}=8$, $d_{2}=15-i_{2}=7, h_{2}(z)=z^{3}, i_{3}=12, d_{3}=15-i_{3}=3, h_{3}(z)=z^{3}+z^{2}+z$, and $h_{4}(z)=z^{3}+1$.

We also need an efficient method to compute the $2^{q}$ vectors $h(\mathbf{A}) \mathbf{x}$, because this has to be redone for each new vector $\mathbf{x}$. These $2^{q}$ vectors can be computed efficiently by using a Gray code (Savage, 1997) to represent the elements of $\mathbb{T}_{q}$ : we enumerate these elements as $t_{0}(z), t_{1}(z), \ldots, t_{2^{q}-1}(z)$ so that $t_{0}(z)=z^{q}$ and any two successive elements in the sequence differ by a single coefficient. This is Gray code enumeration. The first vector $t_{0}(\mathbf{A}) \mathbf{x}=\mathbf{A}^{q} \mathbf{x}$ is computed automatically when we advance the RNG, and then each vector $t_{i}(\mathbf{A}) \mathbf{x}$ is computed from the previous one, $t_{i-1}(\mathbf{A}) \mathbf{x}$, by adding or subtracting (in $\mathbb{F}_{2}$ this is the same) a single vector of the form $\mathbf{A}^{j} \mathbf{x}$ for $0 \leq j<q$. These $q$ vectors are precomputed when we advance the generators by $q$ steps.

In the next subsection, we put together these ingredients to define our algorithm.

### 3.4 Algorithm

The algorithm has two parts: (a) a one-time setup for each jump size $\nu$ and (b) the jumping from $\mathbf{x}_{n}$ to $\mathbf{x}_{n+\nu}$. It is described in Figure 1.

$$
\begin{aligned}
& \text { Preliminary setup for a given } \nu \text {. } \\
& \text { Compute the polynomial } g(z) \text { in (3) and store its coefficients in an array. } \\
& \text { Select } q>0 \text { and choose a Gray code to enumerate the polynomials of } \mathbb{T}_{q} \\
& \text { (equivalently, the integers } 0,1, \ldots, 2^{q}-1 \text { ). } \\
& \text { Compute } m, d_{1}, \ldots, d_{m}, h_{1}(z), \ldots, h_{m}(z), h_{m+1}(z), c_{1}, \ldots, c_{m}, c_{m+1} \text {, } \\
& \text { where } c_{j} \text { is the Gray code of } h_{j}(z) \text {, i.e., } h_{j}(z)=t_{c_{j}}(z) \text {, for each } j \text {. } \\
& \text { Compute } \mathbf{A x}, \mathbf{A}^{2} \mathbf{x}, \ldots, \mathbf{A}^{q} \mathbf{X} \text { by running the RNG for } q \text { steps. } \\
& \mathbf{y}_{0} \leftarrow t_{0}(\mathbf{x})=\mathbf{A}^{q} \mathbf{x} \text {. } \\
& \text { For } i=1, \ldots, 2^{q}-1 \text { do } \\
& \text { Compute } \mathbf{y}_{i+1}=t_{i+1}(\mathbf{A}) \mathbf{x} \text { from } \mathbf{y}_{i}=t_{i}(\mathbf{A}) \mathbf{x} \text {; this requires a single vector XOR. } \\
& \mathrm{x}^{\prime} \leftarrow \mathbf{y}_{c_{1}} \text {. } \\
& \text { For } j=2, \ldots, m \text { do } \\
& \begin{array}{c}
\mathbf{x}^{\prime} \underset{\mathbf{A}^{d_{m}}}{\leftarrow} \mathbf{A}^{\mathbf{x}^{\prime}+\mathbf{y}_{c_{m+1}-d_{j}}+\mathbf{y}_{0} .} . . \mathbf{y}_{c_{j}} . \\
\mathbf{x}^{\prime}
\end{array} \\
& \mathbf{x}^{\prime} \leftarrow \mathbf{A}^{d_{m}} \mathbf{x}^{\prime}+\mathbf{y}_{c_{m+1}}+\mathbf{y}_{0} . \\
& \text { Return } \mathbf{x}^{\prime} \text {. }
\end{aligned}
$$

Figure 1: The jump-ahead algorithm
To summarize the computing costs in the second part, we need $(k-1) c$ operations to advance the RNG by $k-1$ steps, then $2^{q}-1$ vector additions to compute all the vectors $t_{i}(\mathbf{A}) \mathbf{x}$, and a further $m+1 \leq 1+\lceil k /(q+1)\rceil$ additions to compute (6). Since each vector addition requires $\eta$ operations, the total cost is at most

$$
(k-1) c+\left(2^{q}-1+m+1\right) \eta \leq(k-1) c+\left(2^{q}+\lceil k /(q+1)\rceil\right) \eta
$$

operations. The value of $q$ can be chosen to minimize this number, i.e., minimize the upper bound $n_{\mathrm{a}}(k, q)=2^{q}+\lceil k /(q+1)\rceil$ on the number of vector additions (since the term $(k-1) c$ does not depend on $q$ ). As an illustration, Table 1 gives the value of $n_{\mathrm{a}}(k, q)$ as a function of $q$ for $k=19937$ and $w=32$. The minimum is attained for $q=8$. The case $q=0$ refers to the computation via the ordinary Horner method given in (4), for which the table gives the expected number of vector additions for a random polynomial. With $q=8$, the number of additions is reduced approximately by a factor of 4 .

This algorithm can still be applied when the value of $\nu$ is not fixed, but then the (costly) preliminary setup must be repeated each time.

Table 1: Value of $n_{\mathrm{a}}(k, q)$ as a function of $q$ for $k=19937$.

| $q$ | 0 | $\ldots$ | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $n_{\mathrm{a}}(k, q)$ | 9968 | $\ldots$ | 4004 | 3355 | 2913 | 2621 | 2472 | 2506 | 2837 | 3710 | 5630 |

### 3.5 Timings

We made the following experiment to measure the CPU time required by the proposed algorithm to jump ahead by an arbitrary number of steps, from an arbitrary state, for $k=19937$ with both the Mersenne twister and a WELL generator, $k=1024$ with a WELL generator, and various values of $q$. We generated a polynomial $g(z)$ and a state $\mathbf{x}$ at random, uniformly over the set of possible values, and measured the time to compute $g(\mathbf{A}) \mathbf{x}$. Thus, $g(z)$ was a polynomial of degree $\leq k$ and $\mathbf{x}$ was a $k$-bit vector. We replicated this 1000 times and computed the average CPU time for jumping ahead, in milliseconds (msec). Table 2 reports these CPU times and the required memory size for each method, on the following computers: (a) An Intel Pentium 4 at 3.0 GHz , with 1.0 GB of memory, using the gcc compiler with the -02 option, under Linux; and (b) a 64-bit AMD-Athlon 64 3200+, with 2.0 GB of memory, also using Linux and the same compiler. The tested generators are the 32-bit MT19937 (Matsumoto and Nishimura, 1998), the 32-bit WELL19937 (Panneton et al., 2006), and a 64-bit Mersenne Twister named MT19937-64 (Nishimura, 2000).

The timings show that the proposed jump-ahead algorithm is viable even with $q=0$. For $k=19937$, the sliding window with a good value of $q$ provides a speedup by a factor of about 3 on a 32 -bit computer. It also requires 312 Kb of memory, but for most practical applications this is not a serious drawback given the memory sizes currently available. There is more improvement on the Pentium than on the AMD Athlon, and this especially true for the MT19937-64 generator. Note that the speed does not necessarily double when going from a 32 -bit to a 64 -bit processor, for several reasons (memory access is not twice faster, the Pentium and Athlon are different, etc.). Even for $k=1024$ (a small value) the sliding window remains advantageous.

Table 2: Required memory size and average CPU time (in milliseconds) for a random jump ahead, for $k=19937$, on Pentium 4 (32-bit) and AMD Athlon (64-bit) computers.

|  | $q$ | 0 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Memory $(\mathrm{Kb})$ | 2.5 | 39 | 78 | 156 | 312 | 624 | 1248 | 2496 |
| 32-bit Pentium | MT19937, $w=32$ | 15.9 | 5.8 | 5.1 | 4.7 | 4.3 | 4.5 | 5.1 | 6.8 |
|  | WELL19937, $w=32$ | 15.9 | 6.3 | 5.4 | 4.9 | 4.6 | 4.7 | 5.4 | 7.4 |
|  | MT19937-64, $w=64$ | 19.8 | 7.1 | 6.2 | 5.7 | 5.2 | 5.3 | 6.1 | 8.4 |
| 64-bit Athlon | MT19937, $w=32$ | 9.0 | 3.9 | 3.6 | 3.3 | 3.2 | 3.3 | 4.5 | 6.7 |
|  | WELL19937, $w=32$ | 9.7 | 4.0 | 3.8 | 3.5 | 3.4 | 3.5 | 4.8 | 7.1 |
|  | MT19937-64, $w=64$ | 5.6 | 2.4 | 2.4 | 2.3 | 2.1 | 2.3 | 2.9 | 5.0 |

Table 3: Required memory size and average CPU time (in milliseconds) for a random jump ahead, for $k=1024$, on Pentium 4 (32-bit) and AMD Athlon (64-bit) computers.

|  | $q$ | 0 | 3 | 4 | 5 | 6 | 7 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Memory $(\mathrm{Kb})$ | 0.13 | 1.0 | 2.0 | 4.0 | 8.0 | 16.0 |
| 32-bit Pentium | WELL1024, $w=32$ | 0.098 | 0.069 | 0.068 | 0.075 | 0.077 | 0.092 |
| 64-bit Athlon | WELL1024, $w=32$ | 0.096 | 0.060 | 0.062 | 0.068 | 0.071 | 0.086 |

A similar experiment with a clever implementation of the matrix method of Section 3.1 gave the following results: For $k=19937$, the jump ahead took 24.5 msec on the 32 -bit Pentium and 17.0 msec on the 64 -bit Athlon, on average. For $k=1024$, the timings were 0.117 msec on the 32 -bit Pentium and 0.096 msec on the 64 -bit Athlon, on average. This is roughly five times slower than the proposed method for $k=19937$ and $50 \%$ slower for $k=1024$.

We also estimated the time to precompute $z^{\nu} \bmod p(z)$ using $\mathrm{C}++$ and the NTL library (http://www. shoup.net/ntl) for the polynomial calculations. For this, we generated 1000 values of $\nu$ randomly and uniformly in $\left\{0,1, \ldots, 2^{64}-1\right\}$, and measured the average CPU time to compute $z^{\nu} \bmod p(z)$. The average was 239 milliseconds for $k=19937$ and 1.9 milliseconds for $k=1024$. This is much more than the time required to jump ahead.

## 4 Conclusions

We have developed a viable jump-ahead algorithm for large linear RNGs over $\mathbb{F}_{2}$. With this technique, one can easily implement RNG packages with multiple streams and substreams, based on long-period generators such as the Mersenne twister and the WELL with period length of $2^{19937}-1$. For these generators, jumping ahead takes a few milliseconds with the proposed method. This is still significantly slower than for the MRG32k3a generator in L'Ecuyer et al. (2002), whose jump-ahead time is a few microseconds. On the other hand, MRG32k3a is slower to generate its numbers, by a factor of 2 or 3 on common 32 -bit computers, and has a much shorter period length. For applications where jumping ahead is not required too frequently and where a fast long-period RNG is desired, the new jumpahead algorithm comes very handy.

## References

Kelton, W. D. 2006. Implementing representations of uncertainty. S. G. Henderson, B. L. Nelson, eds., Simulation. Handbooks in Operations Research and Management Science, Elsevier, Amsterdam, The Netherlands, 181-191. Chapter 7.

Knuth, D. E. 1998. The Art of Computer Programming, Volume 2: Seminumerical Algorithms. 3rd ed. Addison-Wesley, Reading, Mass.

Law, A. M., W. D. Kelton. 2000. Simulation Modeling and Analysis. 3rd ed. McGraw-Hill, New York.

L'Ecuyer, P., F. Panneton. 2005. Fast random number generators based on linear recurrences modulo 2: Overview and comparison. Proceedings of the 2005 Winter Simulation Conference. IEEE Press, 110-119.

L'Ecuyer, P., R. Simard, E. J. Chen, W. D. Kelton. 2002. An object-oriented randomnumber package with many long streams and substreams. Operations Research 50 10731075.

Matsumoto, M., T. Nishimura. 1998. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. ACM Transactions on Modeling and Computer Simulation 83-30.

Möller, B. 2005. Sliding window exponentiation. H. C. A. van Tilborg, ed., Encyclopedia of Cryptography and Security. Springer-Verlag, New York, 588-590.

Nishimura, T. 2000. Tables of 64 -bit Mersenne twisters. ACM Transactions on Modeling and Computer Simulation 10 348-357.

Panneton, F., P. L'Ecuyer. 2005. On the xorshift random number generators. ACM Transactions on Modeling and Computer Simulation 15 346-361.

Panneton, F., P. L'Ecuyer, M. Matsumoto. 2006. Improved long-period generators based on linear recurrences modulo 2. ACM Transactions on Mathematical Software 32 1-16.

Savage, C. 1997. A survey of combinatorial Gray codes. SIAM Review 39 605-629.
Tezuka, S. 1995. Uniform Random Numbers: Theory and Practice. Kluwer Academic Publishers, Norwell, Mass.

