Generating good pseudo-random numbers

B. A. Wichmann
National Physical Laboratory*, Teddington, UK.
I. D. Hill
Chorleywood, UK.

December 5, 2005

Abstract

In 1982, the authors produced a pseudo-random number generator that has been widely used, but now has been shown to be inadequate by today's standards. In producing a revised generator, extensive use has been made of a test package TestU01 for random number generators. Using this, criteria have been devised for the revised generator— also other high quality generators have been identified. Facilities have been devised to allow the new generator to be used in a highly parallel environment, which is likely to be a feature of many future applications.

Keywords Pseudo-random numbers; Tests for randomness; Non-overlapping sequences; Parallel applications

Randomness and chaos are anathema to the mathematician.
Music of the Primes: Marcus du Sautoy, 2003.

1 Introduction

In the early 1980s there seemed to be a need for a pseudo-random generator that would have good statistical properties, could easily be implemented in any programming language, would give the same results on any computer, and could run on 16-bit computers without overflow problems.

With suitably chosen constants, multiplicative congruential generators were known to do well, but with the 16-bit restriction, it would be difficult to find any constants that would give good statistical properties, so we investigated whether it would be possible to combine more than one, relatively poor, generator in some way that would give better properties than those of the individual components.

Having found two such generators, we tried combining them by adding their results and taking the fractional part of the answer. Although still inadequate,

*and The Open University

the results were sufficiently encouraging as to suggest that if a third component were added, it would give what we were seeking. Algorithm AS 183, [Hill and Wichmann(1982)] and [Wichmann(1982)] resulted.

It has had a 'good innings' but its cycle length of about 7×10^{12} is now considered inadequate for some purposes, and it has been reported [McCullough and Wilson(2005)] as having failed some tests at a probability level of less than 10^{-15}.

Computing developments over the last quarter of a century now make a better version both possible and desirable. In particular, there does not now seem to be a need for the 16-bit restriction, as 32-bit availability is almost universal. In view of the widespread use that the original version has enjoyed, it seems wise to retain the same underlying plan, and to regard the passing of a suitable barrage of empirical tests as adequate justification. There are many different constants that could have been put into the plan with success and we make no claim that those we recommend are necessarily better than others that might have been selected, but only that they have survived stringent testing.

2 Testing a generator

In 1982, the work required to test the generator was very much larger than that required to write it. Fortunately, there are now publicly available test suites for random number generators, and that substantially reduces the effort involved. Moreover, the choice of statistical tests has been made independently of ourselves.

Two such test suites are: DIEHARD, [Marsaglia(2000)] and TestU01, [L'Ecuyer(2005)]. [McCullough and Wilson(2005)] report that our old generator passed DIEHARD, but failed the more recent tests in TestU01. The TestU01 package is very comprehensive with many individual tests but also three batteries of test: Small Crush, Crush, and Big Crush. Our aim with any generator is to 'pass' the Big Crush tests. A review of these tests, [McCullough(2006)], was particularly helpful. Big Crush uses 2^{38} random values and can take over 24 hours to execute, depending mainly upon the speed of the generator being tested.

Big Crush reports P-values for all its tests, and signals those that come outside the [0.01..0.99] range. As it produces 124 P-values altogether, 1 or 2 would be expected to fall outside this range even for perfect randomness, though it must be remembered that not all the tests will be independent. In addition, TestU01 indicates catastrophic failures (defined as outside the $[10^{-15}..1-10^{-15}]$ range) — these should clearly not arise with any high quality generator

The Big Crush test should be run at least twice, with different seeds, as an insurance against an exceptional single run. Big Crush does not specify a 'pass criterion' as such. It is not sufficient merely to have few results outside the [0.01..0.99] range. For a reasonable degree of randomness, the P-values should themselves be more-or-less uniformly distributed between 0 and 1. So to judge a new generator, we actually use three requirements which must all be satisfied:

1. There must be no catastrophic failure (P value outside the $[10^{-15}..1 - 10^{-15}]$ range);

2. For any component test with a value outside the [0.01..0.99] range, we repeat the test a further 4 times. Of these further four runs, we require not more than one to indicate a value outside the above range. This criterion was suggested by Richard Simard, one of the Big Crush authors;

3. For each run of Big Crush we consider the distribution of the P-values by calculating Greenwood's statistic ([Greenwood(1946)]) and finding its approximate tail-area probability using the technique given in [Hill(1979)]. We require the two-tailed value to come within the [0.01..0.99] range. Preferably it should come within the [0.1..0.9] range, but even for perfect randomness, such limits would be violated on 20% of occasions, so it would be unreasonable to insist upon it.

We are now in a position to know when we have an acceptable generator (which need not be our own).

3 Constructing a revised generator

The obvious way to proceed was simply to enhance the existing generator by using three components with suitable constants for 32-bit arithmetic rather than the 16-bit arithmetic we used in 1982. Unfortunately, such a generator failed according to the criteria above. Firstly, a multinomial distribution test failed with Big Crush with a P-value outside the [0.01..0.99] range. When repeated 4 times, the same test failed a further three times. Secondly the observed values of the Greenwood statistic gave P=0.076 on a first test and P=0.050 on a second, not actually failing but too low for comfort. In consequence, we decided to add a further cycle to make our new generator a 4-cycle system.

Using the same design method, we need four primes p_1, p_2, p_3 and p_4, such that $p_i < 2^{31}$. In order to ensure that the generator has the maximum period, we select the $p_i - 1$ to have no common factor other than 2. It is straightforward to write a program to find suitable primes, the candidate ones being: 2147483579, 2147483543, 2147483423 and 2147483123.

The constituent linear congruence generators have no additive constant, but we need to choose multipliers with a range of values up to $\sqrt{p_i - 1}$ such that each value is a primitive root of $p_i - 1$. Suitable values can again be found with the aid of a short program. The candidate multipliers are: 11600, 47003, 23000 and 33000 respectively.

Combining these four generators in a simple way would then require 64-bit integer arithmetic, which is as follows:

$ix := 11\_600 \times ix \bmod 2147483579;$
$iy := 47\_003 \times iy \bmod 2147483543;$
$iz := 23\_000 \times iz \bmod 2147483423;$
$it := 33\_000 \times it \bmod 2147483123;$
$\mathrm{W} := ix/2\_147\_483\_579.0 + iy/2\_147\_483\_543.0$
$\quad + iz/2\_147\_483\_423.0 + it/2\_147\_483\_123.0;$
return $W - \lfloor W \rfloor;$

Note that the four constituent generators are combined by taking each as a fraction of its prime, summing them and taking the fractional part of the result.

We avoid 64-bit arithmetic in the same way as with the old generator. For the first constituent, we have $2147483579/11600 = 185127.89\ldots$ and $2147483579 - 185127 \times 11600 = 10379$. Hence the resulting computation becomes:

$$ix := 11600 \times (ix \bmod 185127) - 10379 \times (ix \div 185127)$$

However, if this result is negative, then 2147483579 must be added.

The algorithm in the variant suitable for 32-bit arithmetic is:

$ix := 11\_600 \times (ix \bmod 185\_127) - 10\_379 \times (ix \div 185\_127);$
$iy := 47\_003 \times (iy \bmod 45\_688) - 10\_479 \times (iy \div 45\_688);$
$iz := 23\_000 \times (iz \bmod 93\_368) - 19\_423 \times (iz \div 93\_368);$
$it := 33\_000 \times (it \bmod 65\_075) - 8\_123 \times (it \div 65\_075);$
if $ix < 0$ **then**
    $ix := ix + 2\_147\_483\_579;$
if $iy < 0$ **then**
    $iy := iy + 2\_147\_483\_543;$
if $iz < 0$ **then**
    $iz := iz + 2\_147\_483\_423;$
if $it < 0$ **then**
    $it := it + 2\_147\_483\_123;$
W := $ix/2\_147\_483\_579.0 + iy/2\_147\_483\_543.0$
        $+iz/2\_147\_483\_423.0 + it/2\_147\_483\_123.0;$
return $W - \lfloor W \rfloor;$

This generator passed the Big Crush test according to our criteria. Using two runs, with different seeds, there was only one value outside the [0.01..0.99] range, and repeating the particular test four more times, the value was within the range each time. The Greenwood statistic gave P=0.22 on the first occasion and P=0.52 on the second, which can be regarded as highly satisfactory.

4 Some properties

The four $p_i - 1$ are:

$$2147483579 - 1 = 2 \times 1073741789,$$

$$2147483543 - 1 = 2 \times 3137 \times 342283,$$

$$2147483423 - 1 = 2 \times 7 \times 557 \times 275389,$$

$$2147483123 - 1 = 2 \times 1073741561.$$

This implies that the period of the generator is:

$$
\begin{aligned}
&2 \times 1073741789 \times 3137 \times 342283 \times 7 \times 557 \times 275389 \times 1073741561 \\
&= 265845484276162438938826670941211 1698 \\
&\approx 2.65 \times 10^{36} \\
&\approx 2^{121}
\end{aligned}
$$

The operations $\times m_i \bmod p_i$ for each of the four cycles have an inverse of the same form. To find the inverse of m_1, we compute the continued fraction for m_1/p_1:

$$\frac{11600}{2147483579} = 1/(185127 + 1/(1 + 1/(8 + 1/(1 + 1/(1 + \frac{1}{610})))))$$

Removing the $\frac{1}{610}$ and multiplying up we get

$$\frac{11600}{2147483579} \approx \frac{19}{3517430}$$

or

$$19 \times 2147483579 - 1 = 11600 \times 3517430$$

Hence the inverse of 11600 is $2147483579 - 3517430 = 2143966149$. The other three inverses are 197144682, 981586662 and 1289335852.

Using these four inverses as multipliers, a new generator could be constructed which would have essentially the same statistical properties as the original one.

For our old generator, [Zeisel(1986)] pointed out that the three cycles could be combined to re-write the generator in the form:

$$X_{n+1} = 16555425264690 \times X_n \bmod 27817185604309$$

Similarly, the new generator has a single-cycle version in which the modulus is the product of the four primes. To find the multiplier a we need to solve the four equations: $a = a_i \bmod p_i$ for the four individual multipliers 11600, 47003, 23000 and 33000 and the four primes 2147483579, 2147483543, 2147483423, and 2147483123. These equations can be solved using the Chinese Remainder Theorem, as pointed out by Zeisel. Producing an explicit solution would be of no real benefit since it is impractical to compute the random numbers this way.

McLeod has pointed out that the precision of the floating point arithmetic influences the values that can be produced ([McLeod(1985)]). His analysis was concerned with obtaining 0.0 with a computer with only 23 mantissa bits. It is doubtful that a precision as low as this should be used for serious computation, but the analysis is indicative in other ways. The old generator produced essentially 48 'bits' of randomness by combining three 16-bit generators. If the old generator was used to produce IEEE double length values which have 53 bits in the mantissa, then the three integers in the seeds could be computed from the result. This implies that the next value can be computed! For this

reason, our generator cannot be regarded as cryptographically strong. With the new 4-cycle generator, the number of random bits is roughly 121 implying that no problems should arise with IEEE double length arithmetic — although, as McLeod noted, the value 0.0 can be produced.

Timings have been made of this generator on an Apple 1.6 GHz Power PC G5 as follows:

Generator	*Millions of calls per second*
Ada GNAT	1.31
Old generator, 32-bit	1.91
Old generator, 16-bit	1.76
3-cycle generator, 64-bit	0.81
3-cycle generator, 32-bit	1.93
4-cycle generator, 64-bit	0.65
4-cycle generator, 32-bit	1.57
C coding of new 32-bit version	3.97

The C generator times are not strictly comparable with the others as the timing methods were different — it seems that the C generator is about 20% faster than the Ada ones. The 16-bit old generator and the 32-bit new 3-cycle one are roughly comparable, the only difference being the size of the operands. The timing shows that even when 64-bit integer arithmetic is available, the 32-bit version can be significantly faster. Of course, our statistical testing implies that only the 4-cycle generators are acceptable — with the 32-bit one being the fastest (at least in this case).

In 1982, the old generator took 0.85 ms on the PDP11 of its day [Hill and Wichmann(1982)]. This implies that the old generator would repeat after 187 years. The new generator, based upon the timings above, would repeat in about 8,000 times the age of the earth! In other words, the increase in the period of the new generator seems to be adequate to cater for the likely increase in computer speeds over the next 20 years. (In contrast, the old generator on an Apple G5 machine can execute the entire sequence in 49 days, which shows that the period is indeed inadequate.)

5 A generator package

Programming languages and implementations typically provide a random number facility. In the case of the C language ([C(1999)]), this provides integers only in the 0..`RAND_MAX` range, and with a means of resetting the seed. Since the integers have type `int`, the range need only be 16 bits.

In contrast, Ada 95 ([Ada(1995)]) provides two comprehensive packages for random numbers. These are very similar, one being for type `Float` and the other generic for any discrete type. In both cases, several sequences can be used, and the state of any generator can be saved or restored. For handling the setting of seeds from external information, the state can be saved or restored

to/from a string. The standard makes the observation 'No one algorithm for random number generation is best for all applications'. Two problems with the existing Ada facilities are worth noting: random numbers of type Long_Float are not available, and the required period for the generator when the numerics annex is implemented is only $2^{31} - 2$.

The Ada packages can provide a random number by means of a function. However, an Ada function cannot change its parameter, which implies that the side-effect that the function must perform to advance the cycle must be undertaken indirectly. For those concerned with program proof, typically for highly critical situations, such behaviour is not allowed. Hence the SPARK Ada subset ([Barnes(2002)]) could not be used to write a random number generator in the functional style. These considerations led to the formulation in Ada different from that in the standard library. Other changes from the Ada 95 standard specification is to produce a result of type Long_Float and to have the Initiator value to the Reset procedure to be positive. The reason for the latter change is for the initialization to align to the proposals for handling multiple sequences.

Abstractly, one would like to hide the state, which in Ada is achieved by means of a private type. However, one does need to set the seeds and hence some form of visibility is needed, which is undertaken by means of conversion to a string. The 'size' of the state is given by the length of the string, which for the generator here is given by the four integers in decimal.

Both Ada and Java provide a simple mechanism to set the seeds. The Ada GNAT implementation uses a 32-bit integer value to set the seed, although this is not adequate to produce all values of the state (the period is about 2^{49}). (The string facility can be used to cover all values.) With Java, the situation is reversed with only 48 of the 64 bits of the value provided being effective in setting the seed. Note that changing the actual algorithm for random numbers could easily alter the relationship between the seed size and the state size.

In Ada 95 and Java it would be possible to undertake an implicit initialization, perhaps using the clock, on the declaration of a generator. We have not done this with the Ada 95 implementation, since we will show later that when many sequences are required, additional care is required with the initialization.

Java shares with Ada the need to provide random numbers in the presence of tasking which implies that the state data must be separated from the code and be able to be placed within the data associated with a task. `Random` is a constructor class, while the methods are of the form `next...`. In fact, the Java class provides very extensive facilities, see [Java(2002)]. Here, the one generator has methods for providing uniform random values of all the major Java types, and well as a Gaussian for `double`. (There is another random facility in the class `Math` which we do not consider here.)

One property of Java is that of strict portability. For instance, the *sin* function must produce the nearest approximation to the mathematical result. For the random function, this implies that when initialized with a specific value, the sequence is determined. Hence a strict implementation cannot change the algorithm for the sequence generation — this is unfortunate since the simple

linear congruence generator used could otherwise be replaced by a generator passing Big Crush.

Producing random numbers is not quite the same as producing repeatable, random-like values in a sequence. Strict repeatability, as in Java, is useful in applying a technique of generating random test cases for software [Wichmann(2000)]. Each test, no matter how complex when generated, can be recorded merely by the seeds. Retesting can be undertaken by regeneration and regression testing by regenerating just those tests which failed.

6 Generating many sequences

Consider the problem of undertaking a Monte Carlo simulation on a highly parallel system with a hundred or more processors. One needs hundreds of different sequences which should not overlap at all.

Given an existing long period generator, even with a randomly chosen seed, there is a small risk that two sequences will overlap. What is the best approach to take under such circumstances? Should one accept the risk, which would certainly be small with the generator presented here?

In fact, for the generator proposed here, the solution appears to be quite simple. We assume each parallel process is given a unique number n. For each simulation, fixed values x, y, z are taken for the first three seeds of the generator. The fourth seed is set to n. For any sequence to overlap, the first three integers must be x, y and z, but this can only arise after about 2^{90} calls of the generator. In other words, we are splitting the generator by means of starting from fixed points on the first three cycles.

It is not always possible to obtain the same effect with the other generators which pass the Big Crush tests. One needs a means of splitting the entire sequence into subsequences which are further apart than the likely number of calls made to the generator.

Unfortunately, our simple solution above has a flaw. Assuming we have 100 parallel processes, when they start execution, the first random number produced will be very similar! We need therefore to devise a method such that the sequence of random numbers given by the first number from each of our processes themselves pass specific tests for randomness. This property may not be needed by some applications but could be important for others.

6.1 A list of sequences

The generation of multiple sequences is a special case of a more general problem of producing a matrix of random numbers:

$$\begin{array}{cccc} s_{11} & s_{12} & s_{13} & \cdots \\ s_{21} & s_{22} & s_{23} & \cdots \\ s_{31} & s_{32} & s_{33} & \cdots \\ \vdots & \vdots & \vdots & \\ \vdots & \vdots & \vdots & \end{array}$$

Our usual sequence of random numbers is represented by the rows. Of course, the rows are much longer than any likely use of the random values.

The solution given above was to set s_{n1} to (x, y, z, n), which we know is not adequate in some circumstances. It is inadequate because the columns do not give a statistically random sequence.

In 2001, a researcher in Spain, Pedro Gimeno, reported a problem to Knuth which showed the Knuth generator as giving unacceptable results. Using the notation above, the issue was that the sequence s_{n1} was not random. Does this matter? This is exactly the problem of producing a matrix of random values so that the columns as well as the rows are random.

If one requires a number of *independent sequences*, each one of which is random (say, passing Big Crush), then the proposal above is fine. Here, only the rows are relevant. However, if the application only uses a few random values from each sequence, and the ordering of the sequences is important, then the problem that was reported to Knuth may be critical.

Can we therefore adapt the generator to produce a list of sequences? The properties we require is that each row and column should be statistically sound and that none of these should overlap.

Using our four primes p_i, and the existing generator which gives the rows above, how can we produce the columns? The answer is simple. For two of the primes, say p_1 and p_2, we produce another two multipliers distinct from those used in the original generator (and their inverses). The method of obtaining the next row is by applying the multiplier to the first two values while leaving the other two fixed. (The operations of moving along the row or going down the column are commutative.) By using two new multipliers we ensure that the column sequences pass at least Small Crush, and, of course, each individual row passes Big Crush as before.

The two new multipliers are 46340 and 22000 to give the cycles:

$$ix := 46340 \times ix \bmod 2147483579$$

and

$$iy := 22000 \times iy \bmod 2147483543$$

Taking the seeds for s_{11} in the matrix above as ix, iy, iz and it, we compute the seeds corresponding to s_{21} by:

$ix := 46\_340 \times (ix \bmod 46\_341) - 41\_639 \times (ix \div 46\_341);$
$iy := 22\_000 \times (iy \bmod 97\_612) - 19\_543 \times (iy \div 97\_612);$
if $ix < 0$ **then**
    $ix := ix + 2\_147\_483\_579;$
if $iy < 0$ **then**
    $iy := iy + 2\_147\_483\_543;$

where iz and it are unchanged. Repeating this operation we can compute the seeds for s_{31}, and so on.

Keeping iz and it fixed ensures that no overlap occurs for over 2.3×10^{18} values at the very minimum.

Since generators running in parallel on different processors, using this method, will then have two of the four components in common, it might seem likely that they would suffer from greater correlation between them than if all four were varying separately on each. In the event, this is not so. Taking 10 cases of 10000 pairs of numbers from non-overlapping parts of the sequence with all four components varying separately, 95% confidence limits for the mean correlation were -0.044 to 0.011. Doing the same with only two components varying separately and the other two varying together, using the technique for deriving seeds given above, 95% limits were -0.003 to 0.010. These both include zero, which is satisfying, and the latter limits are marginally narrower than the former ones. We do not for one moment suggest that the latter would actually give less correlation in general, but there is certainly no evidence here of it being greater.

7 Conclusions

We know of several generators which pass the Big Crush battery of statistical tests. We think that only such generators can be recommended for general use. These can be compared for basic properties (fastest first):

Name	*Period*	*Lines of code*	*Size of state (bytes)*	*Relative timing*
ISAAC	$\geq 2^{40}$	97	1024	1.0
AES	??	85	16	2.1
Mersenne twister [Matsumoto(1998)]	$2^{19937} - 1$	48	2,500	2.3
MRG32k3a [L'Ecuyer(1999)]	$\approx 2^{191}$	31	48	2.7
Knuth, TAOCP [Knuth(2002)]	$\approx 2^{129}$	90	404	4.9
CLCG4 [L'Ecuyer(1997)]	$\approx 2^{121}$	34	16	9.2
This paper — 4-cycle	$\approx 2^{120}$	26	16	10.0
MRG63k3a [L'Ecuyer(1999)]	$\approx 2^{377}$	40	48	14.3

The last three columns should only be taken as an indication of the basic characteristics since the generators operate in rather different ways which makes direct comparison problematic.

Our combined 4-cycle generator can be recommended for the following reasons:

1. Our generator is easy to code in any programming language. It does not depend upon bit manipulation used by several of the other generators.

2. The state is small and easy to handle.

3. It is possible to use the generator to provide multiple sequences needed for highly parallel applications.

We would not necessarily wish to advocate our generator, but rather any generator which satisfies our criteria for passing Big Crush and has a means of handling highly parallel systems.

8 Acknowledgements

In 1982, our original generator was produced by us almost in isolation. In contrast, the revision has benefited substantially from the published literature and help from the following: Bruce McCullough (Drexel University, Philadelphia), Pierre L'Ecuyer and Richard Simard (Université de Montréal) and support from the National Physical Laboratory under the Software Support for Metrology programme. The Open University electronic library was invaluable in obtaining details of many of the generators referenced here.

References

[Ada(1995)] ISO/IEC 8652: 1995. Programming languages — Ada. ISO. Geneva.

[Barnes(2002)] Barnes, J., High Integrity Software — the SPARK Approach to Safety and Security. Addison-Wesley. 2002.

[C(1999)] ISO/IEC 9899: 1999. Programming languages — C. ISO. Geneva.

[Matsumoto(1998)] Matsumoto, M., Nishimura, T., 1998. Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.* 8 (1), 3–30.

[Greenwood(1946)] Greenwood, M. The statistical study of infectious diseases. *J.R.Statist.Soc.,* A, 109, 85–109. 1946.

[Hill and Wichmann(1982)] Hill I. D, and Wichmann B. A., A Pseudo-Random Number Generator. NPL Report, DITC 6/82 May 1982.

[Hill(1979)] Hill, I.D. Approximating the distribution of Greenwood's statistic with Johnson distributions. *J.R.Statist.Soc.,* A, 142, 378–380. 1979. (see also Corrigenda. *J.R.Statist. Soc.,* A, 144, 388. 1981)

[Java(2002)] Class Random. See: http://java.sun.com/j2se/1.4.2/docs/api/java/util/Random.html.

[Jenkins(2002)] Bob Jenkins Jr. ISAAC (Indirection, Shift, Accumulate, Add, and Count) generator. http://burtleburtle.net/bob/rand/isaacafa.html

[Knuth(2002)] Knuth, D.E., 1981. *The Art of Computer Programming*, vol. 2, Seminumerical Algorithms. Section 3.2.1. AddisonWesley, Reading, MA. (The 2002 printing is required.)

[L'Ecuyer(1994)] L'Ecuyer, P., 1994. Uniform random number generation. *Ann. Oper. Res.* 53, 77–120.

[L'Ecuyer(1997)] P. L'Ecuyer and T. H. Andres. A random number generator based on the combination of four LCGs. *Mathematics and Computers in Simulation*, 44:99–107, 1997.

[L'Ecuyer(1999)] P. L'Ecuyer. Good parameters and implementations for combined multiple recursive random number generators. *Operations Research*, 47(1):159–164, 1999.

[L'Ecuyer(2005)] L'Ecuyer, P., and Simard, R. TestU01: A Software Library in ANSI C for Empirical Testing of Random Number Generators. Laboratoire de simulation et d'optimisation. Université de Montréal IRO. Version 6.0, dated January 14, 2005.
`http://www.iro.umontreal.ca/~simardr/`

[Marsaglia(2000)] Marsaglia, G. (1985) The Marsaglia Random Number CDROM, with The Diehard Battery of Tests of Randomness, produced at Florida State University under a grant from The National Science Foundation, available at `http://www.cs.hku.hk/~diehard`

[McLeod(1985)] McLeod A. I. Remark on Algorithm AS 183, *Appl. Statist.,* 34, 198–200, 1985.

[McCullough and Wilson(2005)] McCullough, B. D., and Wilson, Berry. On the Accuracy of Statistical Procedures in Microsoft Excel 2003. *Computational Statistics & Data Analysis*. 49(4), 1244–1252, 2005.

[McCullough(2006)] McCullough, B D (2006), A Review of TESTU01, *Journal of Applied Econometrics* (to appear)

[NIST(2001)] NIST. Specification for the advanced encryption standard (aes). NIST special publication 197 (FIPS-197), National Institute of Standards and Technology (NIST), 2001. See http://csrc.nist.gov/encryption/aes/.

[Wichmann(1982)] Wichmann, B. A, Hill, I. D., 1982. Algorithm AS 183: an efficient and portable pseudo-random number generator. *Appl. Statist.* 31, 188–190. See also: Correction: algorithm AS 183: an efficient and portable pseudo-random number generator. *Appl. Statist.* 33, 123. Reprinted with the correction in Griffiths P. and Hill I. D. *Applied Statistics Algorithms*, Ellis Horwood. 1985.

[Wichmann(2000)] Wichmann, B. A., Some Remarks about Random Testing. Download from: `http://www.npl.co.uk/ssfm/download/stress.pdf`

[Zeisel(1986)] Zeisel H. Remark on Algorithm AS 183, *Appl. Statist.,* 35, 89, 1986.