REAMC Report-008
Issue 1, June 2025

Statistical Performance of ACORN Generators: Further Evidence for Orders 8 – 10

Roy S Wikramaratna

REAMC Limited (Reservoir Engineering and Applied Mathematics Consultancy)

4 Nuthatch Close, Poole, Dorset BH17 7XR, United Kingdom

Website: https://www.reamc-limited.com

Email: rwikramaratna@gmail.com

Telephone: +44(0)7968 707062

REAMC®
Limited

Statistical Performance of ACORN Generators: Further Evidence for Orders 8 – 10

Roy S Wikramaratna

Abstract

The Additive Congruential Random Number (ACORN) generator is a method for generating uniformly distributed pseudo-random numbers which is straightforward to implement for arbitrarily large order and modulus (where the modulus is a sufficiently large power of 2, typically up to 2^{120}); it has been demonstrated in previous papers to give rise to a family of sequences with long period which, for a *k*-th order ACORN generator with modulus a power of 2, can be proven from theoretical considerations to approximate in a particular defined sense to the desired properties of uniformity in up to *k* dimensions.

In 2021 two conjectures were proposed in REAMC Report-003. These assert that for order 8 (or larger) and modulus 2^{120} almost every choice of odd seed together with an arbitrary set of initial values (including the case with all initial values set to zero) leads to a different sequence that can reasonably be expected to pass all the tests in the current Version 1.2.3 of the standard empirical test package known as TestU01. Supporting results were included for ACORN generators of modulus 2^{120} in Report-003 (for order 8, 9 and 10) and later extended in REAMC Report-004 (orders 11 to 15 inclusive) and in REAMC Report-006 (selected orders in the range 16 to 101 inclusive).

All these existing published results were for the set of cases R1-000 to R1-999 (each with a different randomly chosen seed and all initial values set to zero) and for the cases S1-*xxx* (each case S1-*xxx* having the same seed as the corresponding case R1-*xxx*, and each having a different specified set of non-zero initial values). The present REAMC Report-008, together with four Appendices A to D, provides further supporting results for ACORN generators of modulus 2^{120} and orders 8 to 10. Each Appendix includes results obtained using a different set of 1000 randomly chosen seed values to define the cases. It is anticipated that five further Appendices E to I may be published in the future containing results for additional sets of 1000 randomly chosen seeds; however, the report itself is expected to remain unchanged unless there are unexpected results that significantly change some aspect of the conclusions.

Each set of results provides additional support for the two conjectures originally proposed in Report-003. The author is not aware of any other family of pseudo-random number generators for which comparable results have been demonstrated for such a wide range of initialisations, now amounting to some five thousand cases per conjecture for each of the orders 8, 9 and 10.

1 INTRODUCTION

The Additive Congruential Random Number (ACORN) generator is a method for generating uniformly distributed pseudo-random numbers which gives rise to sequences with long period which can be proven from theoretical considerations to approximate to uniformity in any specified number of dimensions. Extensive empirical testing has previously demonstrated the excellent statistical performance of the ACORN generators with appropriately chosen parameters over a very wide range of initialisations.

A previous report [1] proposed two conjectures on the very wide range of conditions (specifically relating to the choice of seed and initial values) under which ACORN sequences having modulus 2^{120} and order 8 or larger can be expected to reliably pass all the TestU01 tests. Results were presented in [1] to support the conjectures for a selection of ACORN generators with 1000 different randomly chosen seed values and initial values chosen either as all zero, or initial values chosen at random, for modulus 2^{120} and for orders 8, 9 and 10. These results were subsequently extended for the same modulus and seed values and corresponding initial values, to cover ACORN generators having orders 11 to 15 inclusive [2]. The results were further extended in [3] which covers cases with selected orders between 16 and 101 (specifically 16, 24. 29, 39, 49, 59, 69, 79, 89, 99 and 101).

In the present report the results that were included in [1] (for modulus 2^{120} and for orders 8, 9 and 10) are extended to further selections of ACORN generators with 1000 different randomly chosen seed values and initial values either all set to zero, or with non-zero initial values selected as before. Nine such new selections of 1000 seed values have been chosen, and the TestU01 tests are in the process of being run for these cases (designated R*n*-*xxx* and S*n*-*xxx*, where *n* runs from 2 to 10 and *xxx* runs from 000 to 999). Summary results are included in the main body of the report (see Table 1 and Table 2) for the first four of these selections (with *n* running from 2 to 5), together with the original cases from reference [1] (which were designated R1-*xxx* and S1-*xxx*). More detailed results for these original cases have already been included in reference [1], while those for the remaining cases completed to date are presented in four Appendices A to D, which are being published contemporaneously with the main report. Five further Appendices E to I will be published later, at approximately 6-monthly intervals, as the cases for each selection of seed values are completed.

Section 2 of the present report provides the definition of an ACORN sequence and an overview of the existing theoretical analysis and empirical test results that have been published to date. Section 3 gives a brief overview of the standard TestU01 package of empirical tests of uniformity and randomness that has been used in this work. Section 4 includes a statement of the two conjectures. It should be noted that Sections 2, 3 and 4 are largely repeated from the corresponding sections in references [1], [2] and [3]; they have been included here so that this

report can be read in isolation without needing to refer back to the earlier reports for this overview.

Section 5 includes a specification of the test cases together with a high-level summary of the results obtained to date and a discussion of the significance of the results. Detailed results for the new cases are being included in the Appendices, which are listed at the end of Section 5.

Section 6 reviews the ongoing work that is currently planned, together with a discussion of some possible avenues of further work (which may provide scope for collaboration with interested parties).

Finally, Section 7 details our overall conclusions based on the work completed to date.

2 OVERVIEW - ACORN SEQUENCES AND ACORN GENERATORS

Let k be a finite, strictly positive integer. A k -th order ACORN sequence is defined from an integer modulus M, an integer seed $Y^0{}_0$ satisfying $0 < Y^0{}_0 < M$ and an arbitrary set of k integer initial values $Y^m{}_0$, $m = 1, ..., k$, each satisfying $0 \leq Y^m{}_0 < M$ by the equations

$$Y^0{}_0 = Y^0{}_{n-1} \quad n \geq 1 \tag{1}$$

$$Y^m{}_n = [Y^{m-1}{}_n + Y^m{}_{n-1}]_{\mathrm{mod}M} \quad n \geq 1, m = 1, \ldots, k \tag{2}$$

where by $[Y]_{\mathrm{mod}\,M}$ we mean the (integer) remainder on dividing Y by M.

The k -th order Additive Congruential Random Number (ACORN) generator is defined by Wikramaratna [4,5] from equations (1) and (2) together with the observation that the sequence of numbers $Y^k{}_n$ can be normalised to the unit interval by dividing by M

$$X^k{}_n = Y^k{}_n/M \quad n \geq 1 \tag{3}$$

The numbers $X^k{}_n$ defined by equations (1) - (3) approximate to being uniformly distributed on the unit interval in up to k dimensions, provided a few simple constraints on the initial parameter values are satisfied. In short, the modulus M needs to be a prime power, with powers of 2 offering the most straightforward implementation, while the seed $Y^0{}_0$ and the modulus should be chosen to be relatively prime (two numbers are said to be relatively prime if they have no prime factors in common, which means that their greatest common divisor is 1). This is the approach that we have adopted in most of our previous experiments with the ACORN generator, and it appears to work very successfully.

The original implementation proposed in [4] first combined equations (1) – (3) to eliminate the $Y^k{}_n$ and then used real (floating point) arithmetic modulo one to calculate the $X^k{}_n$ directly. That

implementation suffered from a number of conceptual and practical limitations (in particular, although the statistical properties of the sequences were unaffected, the sequences generated with any specific initialisation could not be guaranteed reproducible on different hardware or with different compilers). These limitations could be overcome [5] through the use of the integer implementation based on equations (1) – (3). Theoretical analysis given by Wikramaratna [5] has shown that the numbers $Y^m{}_n$ are of the form

$$Y^m{}_n = \left[\sum_{i=0}^{m} Y^i{}_0 Z^{m-i}{}_n\right]_{\mathrm{mod}M} \tag{4}$$

where for any integer values of a (non-negative) and b (positive) we define $Z^a{}_b$ by

$$Z^a{}_b = \frac{(a+b-1)!}{a!(b-1)!} \tag{5}$$

More extensive theoretical analysis and empirical testing of the algorithm have been described in subsequent papers, including [6] and [7].

From a theoretical viewpoint [6] the ACORN generator was shown to be a very particular special case of a multiple recursive generator. When this formulation was written in a specified matrix form, it led in turn to the discovery of some special matrices (called centro-invertible matrices) which have some interesting and unusual properties [8]. The theoretical analysis in [7] led to a proof that a k-th order ACORN generator with modulus 2^{30p} approximates to being k-distributed in a particular sense that was defined in the paper.

Empirical tests carried out previously by the author, making use of the Diehard statistical test suite, Marsaglia [9], have been reported in [6]. Further empirical testing reported in [7], used Version 0.6.1 of the TestU01 package described by L'Ecuyer and Simard [10]. More recently, empirical testing has been carried out using the most current Version 1.2.3 of the TestU01 package as reported in [11] and [12]; that work has since been systematically extended to ACORN generators with much wider choices of initialisations.

Another paper [13] addressed the periodicity of ACORN sequences for any specified order, modulus and any choice of seed that is relatively prime with the modulus. As an example, every ACORN generator with order at least 8, modulus 2^{120} and any choice of odd seed has a period length in excess of 2^{123}. This period was considered far in excess of the maximum period that might be required in the largest conceivable computationally practicable (using hardware available at the time in 2020) application requiring a source of uniformly distributed pseudo-random numbers. We believe that this still holds true in 2025; however, we note that the ACORN algorithm extends naturally and very easily to even longer period lengths (and further improved statistical performance) simply by increasing the modulus to a larger power of 2.

See also further discussion on the website http://acorn.wikramaratna.org, which includes a more comprehensive list of relevant ACORN references as well as links to downloadable versions of the references. Recent ACORN references (including this report) are available for download from the REAMC Limited website, https://www.reamc-limited.com.

3 EMPIRICAL TEST PACKAGE - TESTU01 VERSION 1.2.3

The TestU01 package has been described by L'Ecuyer and Simard [10]. They considered the application of empirical tests of uniformity and randomness to sequences generated by a wide range of algorithms and developed a comprehensive set of empirical tests that were designed to detect undesirable characteristics in such sequences. L'Ecuyer and Simard presented results of applying the TestU01 tests (collectively called the BigCrush test battery) to a large number of different sequences, identifying specific generators that passed all of the tests, as well as identifying many generators (including some that are still widely used) that have serious deficiencies in respect of certain specific tests.

L'Ecuyer and Simard documented the testing of "a long list of well known or widely used generators" comprising using some 92 different methods and identified some 26 methods that gave rise to a sequence that passed all the tests in the TestU01 BigCrush test battery, according to the criteria used. They stated that the methods included in their paper represent a subset of the methods that they tested, and that the results presented were a representative subset of the overall results obtained. It is noted that L'Ecuyer and Simard typically only reported results from one particular initialisation for each pseudo-random number generator that they tested, and it is unclear whether they applied their tests more than a single time to each generator.

The results presented below for ACORN generators with modulus equal to 2^{120} (which were not included among generators considered by L'Ecuyer and Simard) were obtained using the latest version 1.2.3 of TestU01. For each sequence that is tested the BigCrush battery of tests calculates p-values for 180 different test statistics (some of the test statistics are calculated more than once using different parts of the sequence of pseudo-random numbers that is being tested, so that in practice there are a total of 254 different p-values that are calculated and output by the BigCrush test battery in version 1.2.3 of TestU01), using some 2^{38} pseudo-random numbers from each sequence in the process. We follow L'Ecuyer and Simard in defining a "failure" to be a test statistic with a p-value outside the range $[10^{-10}, 1-10^{-10}]$ and a "suspect" value to be a test statistic falling in one of the ranges $[10^{-10}, 10^{-4}]$ or $[1-10^{-4}, 1-10^{-10}]$. This is the same approach that was adopted in [11]. For avoidance of doubt, we note that that there were a small number of cases giving rise to a result with either p=0.9999 or p=0.0001, which we have counted as a suspect value rather than a pass.

4 TWO CONJECTURES ON ACORN STATISTICAL PERFORMANCE

Based on our experience of testing ACORN generators (including results discussed in references [14], [11], [12] as well as other unpublished results) we proposed the following two conjectures concerning the results of applying the TestU01 version 1.2.3 BigCrush tests to certain ACORN generators (see reference [1]). A few preliminaries are appropriate before stating the conjectures, in particular to define some terms that will be required.

i. The conjectures are stated for ACORN generators with modulus 2^{120}. We believe that the same conjectures would hold for ACORN generators with modulus equal to any larger power of two (after making some obvious generalisations to the definitions).
ii. It is straightforward to verify (as was shown previously in [1]) by counter-example that the conjectures do not hold for order k=7, so that k=8 is the smallest order for which the conjectures might hold.
iii. The phrase 'select an odd seed at random' is interpreted to mean that each of the first 119 bits in the binary representation of the seed is assigned to be zero or one with probability 0.5, and the final bit is assigned to be one.
iv. The phrase 'select the initial values at random' is interpreted to mean that each bit in the binary representation of the seed is assigned to be zero or one with probability 0.5.
v. The phrase "the sequence … will almost certainly pass all the tests in version 1.2.3 of the TestU01 BigCrush test suite" is interpreted to mean that the likelihood of failure on one or more of the tests on each run is vanishingly small; based on our experience this probability is certainly very much less than 1 in a thousand and we believe in practice it will turn out to be less than 1 in a million.
vi. As already noted, 254 different p-values are calculated in each run of the BigCrush test battery. For a truly random UD [0,1) sequence, the theoretical likelihood of a single 'suspect value' arising by chance in an individual run of the BigCrush test battery would be 254/5000 = 0.0508 (5.08% or roughly 1 in 20) while the corresponding theoretical likelihood of a single 'failure' would be 5.08×10^{-8} (roughly 1 in 20 million).

CONJECTURE 1. If we define an ACORN generator with modulus 2^{120} and order k of 8 or larger, select an odd seed "at random" and set all k initial values to be zero, then the resulting ACORN sequence will almost certainly pass all the tests in version 1.2.3 of the TestU01 BigCrush test suite. □

CONJECTURE 2. If we define an ACORN generator as in Conjecture 1, again select an odd seed "at random" (in the sense of conjecture 1) together with any specified set of initial values (which may either be selected "at random", or assigned particular chosen values) then the resulting ACORN sequence will almost certainly pass all the tests in version 1.2.3 of the TestU01 BigCrush test suite. □

It should be noted that Conjecture 1 is a particular special case of Conjecture 2. If Conjecture 1 fails, then Conjecture 2 must also fail. On the other hand, if Conjecture 2 is correct then Conjecture 1 must also be correct.

It is clearly not possible to conclusively 'prove' the correctness of either of the conjectures, given the number of different possible choices of seed and the time it would take to exhaustively test all the possible choices; in addition, the probabilistic nature of the TestU01 tests, together with the requirement that the initialisation should be made "at random" means that a definitive answer cannot be achieved. However, it is possible to test the validity of the conjectures by randomly selecting a sufficiently large set of seeds and initial values (using any appropriate method to assign a value to each of the bits) and then running the TestU01 BigCrush tests for each resulting case.

The result of such testing will be either to disprove the conjectures (if there is a significant fraction of the randomly selected initialisations that result in one or more failures) or to lend support to the conjectures (if there are either no failures or only an extremely small fraction of randomly selected initialisations result in a failure). In either case, further support for the conjectures can be gained by calculating the proportion of cases giving rise to 'suspect values' and confirming that these arise with approximately the same frequency that would be expected for a truly random sequence.

We note the importance of selecting the seeds "at random". It is clear that there are certain specific seed values for which the conjecture does not hold (for example, one could achieve this by choosing the seed to be equal to 119 zeroes followed by a single 1, and setting all the initial values to be zero; on the other hand, the probability of randomly selecting this particular seed value would be 1 in 2^{120} or about 1 in 10^{40}). However, it is possible to choose any appropriate and convenient method - including ACORN (as here), any other good quality standard prn generator, or an unbiassed physical method such as coin tossing to assign the individual bits in the seeds to be used for each of the test cases. The important point is that the bits take the values 0 and 1 with equal likelihood and that the values assigned to successive bits are uncorrelated.

5 TEST CASES

All cases considered here have modulus equal to 2^{120}. In each case considered, the seed S is a 120-bit integer. Note that all seeds are chosen to be odd (so that the last digit in the binary representation of the seed is a one). Choosing an odd seed value guarantees that the period of the sequence will be at least eight times the modulus (which comes out as 2^{123} in this instance) for every ACORN generator of order 8 or more (see Wikramaratna [13]). When implementing an ACORN generator with modulus equal to 2^{120} (using 64-bit or 32-bit integer arithmetic), the

seed S can be represented as follows, using either two 60-bit integers, t_1 and t_2, or four 30-bit integers, s_1, s_2, s_3 and s_4.

$$S = 2^{60}t_1 + t_2 = 2^{90}s_1 + 2^{60}s_2 + 2^{30}s_3 + s_4 \qquad (6)$$

This representation is unique in the sense that it defines a one-one onto mapping between any two of S, (t_1, t_2) and (s_1, s_2, s_3, s_4). In presenting detailed results (in the Appendices) it has proved helpful to specify both the values of s_1, s_2, s_3 and s_4 as base 10 numbers in the range zero to (2^{30}-1), and the corresponding values of t_1 and t_2 as base 10 numbers in the range zero to (2^{60}-1).

5.1 Specification of Test Cases

A first set of test cases was specified in [1], as follows. Firstly, one thousand 120-bit integers were generated, the first 119 bits being assigned randomly and the final bit set to one; these integers were assigned Case Indices 1-000 to 1-999, and their values are in the tables of results for individual test cases (see the Tables 3 to 42 of reference [1]). In this study, an ACORN generator was called 119 times to generate the individual bits for each of the 1000 seed values; however, as already discussed the particular method used to generate the seeds is not significant, and similar results would be expected if we used different generators to construct the seeds (either changing the order, modulus and initialisation of the ACORN generator used to generate the individual bits, or selecting any other standard algorithm to generate the individual bits).

The cases designated R1-*xxx* were designed to test Conjecture 1. Each of the Cases R1-*xxx* (*xxx* from 000 to 999) was given a random seed, as selected above (with Case Index 1-*xxx*) and all k initial values set to zero. These cases were previously run with order k = 8 to 15 inclusive (in addition, a smaller number of cases, *xxx* from 000 to 099, were run with order k = 7, to demonstrate that the Conjecture 1 cannot hold for order k=7) – see references [1, 2]. The results were subsequently extended in [3] to cover selected orders between 16 and 101 (specifically, choosing k = 16, 24, 29, 39, 49, 59, 69, 79, 89, 99 and 101).

The cases designated S1-*xxx* were designed to test Conjecture 2. Each Case S1-*xxx* was given the same random seed as selected above for the corresponding Case R1-*xxx* and assigned non-zero initial values; the i-th initial value was the integer assigned to Case Index 1-*yyy*, where $yyy=[xxx+i]_{\text{mod } 1000}$. Thus, the initial values for Case S1-000 are equal to the integers assigned to Case Index 1-001, 1-002, 1-003, 1-004, etc; while for Case S1-997 they are equal to the seeds assigned to Case Index 1-998, 1-999,1-000, 1-001, …, etc. Each Case S1-*xxx* was previously run with the same choice of orders as the Cases R1-*xxx*, see references [1, 2, 3].

For the purpose of this report, additional sets of cases have been defined with the Case Indices *n-xxx*, with *n*=2, 3, 4, 5, …, 10, leading in a similar way to the Cases R*n-xxx* and S*n-xxx* respectively. Summary results for all the Cases R1-000 to R5-999 and S1-000 to S5-999 (but limited to orders 8, 9 and 10) are included in the next section. Detailed results were included in reference [1] for the cases with Case Index 1-*xxx*; detailed results for the remaining cases completed to date are included in the Appendices A through D as detailed below (which are being published in 2025 at the same time as this Report-008) and it is anticipated that further Appendices may be published from time to time in the future, at approximately 6 monthly intervals, as additional sets of test cases are completed.

Appendices published contemporaneously with the main report:

Report-008A Appendix A, Case Indices 2-*xxx*
Report-008B Appendix B, Case Indices 3-*xxx*
Report-008C Appendix C, Case Indices 4-*xxx*
Report-008D Appendix D, Case Indices 5-*xxx*

Appendices to be published at future dates, as the relevant cases are completed:

Report-008E Appendix E, Case Indices 6-*xxx* (anticipated publication 2026)
Report-008F Appendix F, Case Indices 7-*xxx* (anticipated publication 2026)
Report-008G Appendix G, Case Indices 8-*xxx* (anticipated publication 2027)
Report-008H Appendix H, Case Indices 9-*xxx* (anticipated publication 2027)
Report-008I Appendix I, Case Indices 10-*xxx* (anticipated publication 2028)

5.2 Summary Results

The results for the 5000 cases R*n-xxx* with values of *n* between 1 and 5 and values of *xxx* from 000 to 999 are summarised in Table 1 (numbers of ‘failures’ on the left and numbers of ‘suspect values’ on the right); the definition of ‘failures’ and ‘suspect values’ is as already discussed in the final paragraph of Section 3.

The results for the 5000 cases S*n-xxx* are summarised in Table 2 (numbers of ‘failures’ on the left and numbers of ‘suspect values’ on the right); again, the definition of ‘failures’ and ‘suspect values’ has been discussed in the final paragraph of Section 3.

These two Tables both have similar format, as follows: each of the first 5 lines in each table represents a set of 1000 cases. The final line of the table gives the average over the full set of 5000 cases that have been completed to date. Columns relating to numbers of errors are headed in pink, while those relating to numbers of suspect values are headed in green.

Table 1 shows that there were no failures on any of the tests for any of the five thousand cases R*n-xxx* with values of *n* between 1 and 5 and values of *xxx* from 000 to 999 for any of the ACORN generators tested with order 8, 9 or 10. On the other hand, suspect values were found in between 4.8% and 5.4% of the cases, consistent with the theoretical estimate of 5.08%.

Table 1 Summary results for the 5000 Cases R1-000 to R5-999

FAILURES - SUMMARY				SUSPECT VALUES - SUMMARY			
	Order 8	Order 9	Order 10		Order 8	Order 9	Order 10
Cases R1-000 to R1-999	0	0	0	Cases R1-000 to R1-999	63	47	52
Cases R2-000 to R2-999	0	0	0	Cases R2-000 to R2-999	57	54	44
Cases R3-000 to R3-999	0	0	0	Cases R3-000 to R3-999	54	55	45
Cases R4-000 to R4-999	0	0	0	Cases R4-000 to R4-999	54	44	49
Cases R5-000 to R5-999	0	0	0	Cases R5-000 to R5-999	42	57	51
TOTAL R1-000 to R5-999	0	0	0	TOTAL R1-000 to R5-999	270	257	241
AVERAGE R1-000 to R5-999	0.0000	0.0000	0.0000	AVERAGE R1-000 to R5-999	0.0540	0.0514	0.0482

Table 2 shows that there were no failures on any of the tests for any of the five thousand cases S*n-xxx* with values of *n* between 1 and 5 and values of *xxx* from 000 to 999 for any of the ACORN generators tested with order 8, 9 or 10. Once again, suspect values were found in between 4.8% and 5.4% of the cases, consistent with the theoretical estimate of 5.08%.

Table 2 Summary results for the 5000 Cases S1-000 to S5-999

FAILURES - SUMMARY				SUSPECT VALUES - SUMMARY			
	Order 8	Order 9	Order 10		Order 8	Order 9	Order 10
Cases S1-000 to S1-999	0	0	0	Cases S1-000 to S1-999	60	43	43
Cases S2-000 to S2-999	0	0	0	Cases S2-000 to S2-999	44	57	51
Cases S3-000 to S3-999	0	0	0	Cases S3-000 to S3-999	57	47	51
Cases S4-000 to S4-999	0	0	0	Cases S4-000 to S4-999	53	57	47
Cases S5-000 to S5-999	0	0	0	Cases S5-000 to S5-999	57	51	48
TOTAL S1-000 to S5-999	0	0	0	TOTAL S1-000 to S5-999	271	255	240
AVERAGE S1-000 to S5-999	0.0000	0.0000	0.0000	AVERAGE S1-000 to S5-999	0.0542	0.0510	0.0480

It is worth noting that the criterion for a "failure" (a single p-value less than or equal to 10^{-10}) to occur by chance is 1 million times less likely to occur than a "suspect value" (a single p-value less than or equal to 10^{-4}), assuming that the sequences being tested are in fact uniformly distributed in sufficient dimensions; thus, we might expect to get an average of approximately one "failure" for every 2×10^7 randomly chosen seeds that were tested (compared with

approximately one suspect value for every 20 randomly chosen seeds that were tested), so that the absence of any failures to date provides additional supporting evidence for the validity of the conjectures. Thus, it is clear that even with a doubling of the number of cases tested the chance of even a single failure arising across all the cases that are in the process of being run should still remain negligible.

Both tables show the total number of suspect values occurring; across the two sets of runs there were a small number of cases giving rise to two suspect values and an even smaller number with three suspect values, so the number of suspect values occurring is also a close approximation to the number of cases with one or more suspect values.

Taken together, the results summarised in these two tables provide strong supporting evidence for the correctness of Conjectures 1 and 2. For more detailed results (including the seed values chosen for each of the runs that were considered) see Tables 1 to 42 in reference [1] and the corresponding Tables A-1 to A-42, B-1 to B-42, C-1 to C-42 and D-1 to D-42 contained in the Appendices A to D respectively.

5.3 Significance of Results

In this report, and in the earlier report [1], we have identified tens of thousands of different ACORN sequences (of order 8, 9 and 10), with 5000 different seed values which were selected at random according to the criteria of the conjectures, which we have demonstrated to pass all the TestU01 BigCrush tests. More significantly, using this approach of random selection we have shown that ‘suspect values’ occurred in approximately the expected proportion of cases, while we have so far completely avoided selection of any of the ‘undesirable’ seed values that do give rise to failures.

We note that corresponding tests have also been undertaken using the first 1000 of these seed values for ACORN generators with orders 11 to15 [2] and for selected orders between 16 and 101 [3], showing very similar results for the higher orders as for orders 8, 9 and 10.

Taken together, all these results lend considerable support to the veracity of the Conjectures 1 and 2 concerning the ACORN generators.

6 FURTHER WORK

6.1 Work in Progress

Work is currently in progress to complete the runs required to compile the Appendices E to I. When completed this will double the number of cases that have been tested. At the present rate

of progress this work is expected to complete early in 2028. However, it is possible that this may be accelerated if additional computing resources can be deployed, either through access to cloud computing resources or through some form of research collaboration.

6.2 Related Topics

We have identified the potential to undertake some further work on related topics as outlined below, given access to the required time and resources. Any reader with a Mathematics, Computer Science or Statistics background who may be interested in collaboration on these (or related) topics is invited to contact the author using the details shown on the front cover of this report.

i. Further statistical analysis of existing results for each of the ACORN cases that have been completed. In particular, it may be instructive to look at the results for intermediate values of p (falling somewhere between a suspect value, $p=10^{-4}$, and a failure, $p=10^{-10}$). As an example, reducing the threshold value from $p=10^{-4}$ to $p=10^{-5}$ might be expected to result in a factor of 10 reduction in the number of occurrences. In this case we might expect to see on average just over 5 occurrences for every 1000 cases tested and just over 25 occurrences for the 5000 cases corresponding to either Table 1 or Table 2. While there is a vast amount of data to be trawled through, the process is amenable to automation and might be completed relatively rapidly.

ii. As stated, the Conjectures apply to ACORN generators having modulus 2^m with m equal to 120. It might be interesting to apply the same tests to ACORN generators with selected smaller values of m to understand the behaviour as m is reduced down from 120 towards 60 (based on preliminary results from a very small number of tests, it seems clear that a value of m equal to 60 is not large enough for either of the Conjectures to hold, and it may be possible to identify a threshold value of m below which the Conjectures begin to fail). A series of such tests could be run without needing to make any changes to the existing code, simply by setting the seed and initial values appropriately.

7 CONCLUSIONS

In a previous report [1] we proposed the Conjectures 1 and 2 concerning the results of applying the TestU01 version 1.2.3 BigCrush tests to ACORN generators having modulus 2^{120} and with order 8 or more.

i. Extensive new results have been presented here (for ACORN generators having modulus 2^{120} and with orders 8, 9 and 10) to further support the veracity of the conjectures. Overall, the results provide strong supporting evidence for the correctness of the two Conjectures 1 and 2.

ii. If correct, the conjectures that have been proposed lead to a method for generating vast numbers of different sequences which can each be relied on to pass all the TestU01 BigCrush tests with an extremely high probability. The number of such ACORN sequences with modulus 2^{120} and any given order greater than or equal to 8 is a very large multiple of 2^{119}; each of these sequences has a period length in excess of 2^{123}.

iii. It should be noted that two ACORN sequences selected in the manner described cannot strictly be guaranteed to be independent unless they have different orders. Despite this, if the seeds and the initial values are all chosen randomly (in the sense defined in the paper), there is minimal chance of any significant correlation between the two resulting sequences even if they both have the same order.

iv. In addition, we have identified a number of areas for possible future collaboration with mathematical, computer sciences or statistics researchers, including some that may form a suitable basis for undergraduate or postgraduate research projects. Any reader for whom such a collaboration might be of interest is invited to contact the author by e-mail (rwikramaratna@gmail.com).

REFERENCES

NOTE. Further discussion of ACORN sequences is available at the ACORN website http://acorn.wikramaratna.org/index.html. Included on that website there is a page with a more comprehensive list of relevant ACORN references as well as links, pointing to downloadable versions, or to other sites where those references can be accessed and downloaded (see http://acorn.wikramaratna.org/references.html). Recent ACORN references (including REAMC Limited reports, papers and presentations) are available for download from the publications page of the REAMC Limited website, https://www.reamc-limited.com.

---

1 R.S. Wikramaratna, Statistical Performance of Additive Congruential Random Number Generators Part 2 - Conjectures Concerning Seed Values Chosen Uniformly at Random, REAMC Report-003, Issue 2, August 2021, REAMC Limited, UK. *Note that this report was originally published as Issue 1 in January 2021; Issue 2 is unchanged apart from a few minor typographic corrections.* [Link for download is available at https://www.reamc-limited.com]

2 R.S. Wikramaratna, Two Conjectures on Statistical Performance of ACORN Generators: Evidence for Orders 11 - 15, REAMC Report-004, August 2021, REAMC Limited, UK. [Link for download is available at https://www.reamc-limited.com]

3 R.S. Wikramaratna, Statistical Performance of ACORN Generators: Evidence for Selected Orders 16 – 101, REAMC Report-006, May 2023. REAMC Limited, UK. [Link for download is available at https://www.reamc-limited.com]

4 R.S. Wikramaratna, ACORN - A New Method for Generating Sequences of Uniformly Distributed Pseudo-random Numbers, *J. Comput. Phys.*, **83**, pp16-31, 1989.

5 R.S. Wikramaratna, Theoretical Background for the ACORN Random Number Generator, Report AEA-APS-0244, AEA Technology, Winfrith, Dorset, UK, 1992.

6 R.S. Wikramaratna, The Additive Congruential Random Number Generator – A Special Case of a Multiple Recursive Generator, *J. Comput. and Appl. Mathematics*, **261**, pp371–387, 2008. [doi: 10.1016/j.cam.2007.05.018].

7 R.S. Wikramaratna, Theoretical and Empirical Convergence Results for Additive Congruential Random Number Generators, *J. Comput. Appl. Math.*, **233**, pp2302-2311, 2010. [doi: 10.1016/j.cam.2009.10.015].

8 R.S. Wikramaratna, The Centro-invertible Matrix: A New Type of Matrix Arising in Pseudo-random Number Generation, *Linear Algebra and Its Applications*, **434**, pp144-151, 2011. [doi: 10.1016/j.laa.2010.08.011].

9 G. Marsaglia, The Marsaglia Random Number CDROM including the Diehard Battery of Tests of Randomness, Florida State University, Florida, USA, 1995. (originally made available from http://stat.fsu.edu/pub/diehard ; since November 2019 has been available from https://github.com/jeffThompson/DiehardCDROM)

10 P. L'Ecuyer and R. Simard, TestU01: A C Library for Empirical Testing of Random Number Generators, *ACM Transactions on Mathematical Software*, **33**, 4, Article 22, 2007.

11 R.S. Wikramaratna, Statistical Testing of Additive Congruential Random Number (ACORN) Generators, Meeting on 'Numerical algorithms for high-performance computational science', April 2019, The Royal Society, London, UK. [Link for download is available at https://www.reamc-limited.com]

12 R.S. Wikramaratna, The Additive Congruential Random Number (ACORN) Generator - pseudo-random sequences that are well distributed in k dimensions, University of Oxford Numerical Analysis Group Internal Seminar, June 2019. [Link for download is available at https://www.reamc-limited.com]

13 R.S. Wikramaratna, Periodicity of ACORN Sequences with Arbitrary Order and Modulus, REAMC Report-001, March 2020. REAMC Limited, UK. [Link for download is available at https://www.reamc-limited.com]

14 R.S. Wikramaratna, Statistical Performance of Additive Congruential Random Number Generators Part 1 - Results of Testing Some Specific Seed Values, REAMC Report-002, November 2020. REAMC Limited, UK. [Link for download is available at https://www.reamc-limited.com]

