

1 A random number generator for lightweight authentication protocols: xorshiftR+

2 Umut Can ÇABUK, Ömer AYDIN*, Gökhan DALKILIÇ

3 Department of Computer Engineering, Engineering Faculty, Dokuz Eylül University,

4 Izmir, Turkey

5 *Correspondence: omer.aydin@deu.edu.tr

6 **Abstract:** This paper presents results of a research that aims to find a suitable, reliable
7 and lightweight pseudo random number generator for constrained devices used in the
8 Internet of things. Within the study, three reduced versions of the xorshift+ generator are
9 built. They are tested using the TestU01 suite (as well as the NIST suite) to measure their
10 ability to produce randomness and performance values among with some other existing
11 generators. The best of our reduced variations according to our tests, called the
12 xorshiftR+, demonstrated great suitability for lightweight devices considering its
13 randomness, performance and resource usage.

14 **Key words:** TestU01, xorshift, lightweight cryptography, Internet of things

15 1. Introduction

16 The rapidly emerging concept of the Internet of things (IoT) brings new approaches to
17 everyday problems as well as industrial applications. These approaches rely on bundles
18 of cheap, efficient and dedicated networked devices that work and communicate
19 continuously. These so-called lightweight devices of IoT have limited power, space and
20 computation resources, hence there is a huge need of developing suitable security
21 protocols and methodologies tailored for those. Furthermore, most of the contemporary
22 security protocols are not optimized for lightweight environments and require more
23 sources than IoT devices may efficiently provide. For example, almost all authentication
24 protocols use random number generators to function and improvements on these may

1 boost the usage of IoT devices in areas, where security concerns exist. In this paper, we
2 propose a lightweight random number generator recommendation for security
3 applications in constrained devices.

4 While working with random number generators (RNGs), the most important thing to take
5 into consideration is that, there is no perfect generator that fits to all conditions [1]. This
6 is mostly because, true randomness is a non-deterministic process, which can't be
7 synthesized using mathematical methods on a software environment, and thus leads the
8 concept of pseudo randomness. On the other hand, even by using hardware sources, it is
9 practically very hard to produce series of numbers, which have anticipated characteristics
10 of true randomness. In any case, quality of a generator is correlated with its proximity to
11 true randomness and its computational requirements. True or high degree randomness can
12 be very expensive or inefficient or just unnecessary in some cases. Hence, the best option
13 should be chosen according to the needs of the application under development.

14 The main work that leads to this study is to develop a radio frequency identification
15 (RFID) authentication protocol, which runs on RFID tags without an intervention of the
16 reader or another computer, for particularly, but not limited to, wireless sensing and
17 identification platform (WISP). A very accurate use case scenario is given in another
18 study [2]. In this context, there is a need of a lightweight and reliable RNG. Besides, there
19 are hardware sensors, like thermometer, accelerometer, etc. on WISP and many other type
20 of tags those may be used to feed or seed the RNG algorithms. However, these hardware
21 inputs can't be used solely to produce random numbers, since in certain stable
22 environmental conditions, the outcomes will usually follow certain patterns. Thus, we
23 can't trust hardware sources, but a pseudo RNG, seeded with hardware sources or external
24 inputs may provide a satisfactory result, if suitable RNG algorithms are used.

2. Related works

RNGs are mostly studied from the point of view of their proximity to true randomness and reliability until last decade since they were in use in relatively powerful computers. But when smart mobile devices gained in popularity and eventually the IoT concept introduced, the idea of moving RNGs to mobile devices with (sometimes very) limited capabilities brought performance issues to the sight. In example, the standardization document of randomness recommendations for security, RFC 1750, suggests data encryption standard (DES), cryptographic hash functions and some other (not so simple) methods to provide randomness however, there is no concern for lightweight devices [3]. Before proceeding to the RNG algorithms, it would be accurate to take a glance at the review criteria for these RNGs. There are several popular randomness tests (or more accurately, test suites) used to examine generator functions, those are supposed to have random outcomes. The ones we'll refer to are, the Diehard battery of Marsaglia [4], TestU01 suite of L'ecuyer and Simard [5], which consists of 6 test batteries, and the NIST test suite having 15 tests [6]. The more recent publications of NIST, namely SP800-90 (a, b, c), are also taken into account [7-9]. The generators we proposed and the TestU01 suite comply with most (if not all) of the instructions given in these papers. In example, SP800-90b [9] implies the need of a dataset containing at least 1,000,000 values, yet TestU01 produces around 50,000,000 numbers to test the generators [5]. Our ultimate generator, can be used in the frameworks mentioned by NIST [7], with no or very little modification depending on the scenario.

In fact, TestU01 suite contains, Small Crush, Crush, Big Crush, Alphabit, Rabbit batteries and a pseudo-NIST battery. This NIST battery contains 13 original tests out of 15

1 included in the real NIST suite, and the remaining 2 are replaced by more advanced
2 variations [5]. Therefore, it will be the main measure of our experiments.

3 Marsaglia in his paper [10] has shown a class of RNGs, called “xorshift RNGs”, consists
4 of consecutive bitwise xor and shift operations using seeds. He claimed that the
5 algorithms are extremely fast and reliable in terms of randomness. Since the speed of this
6 algorithm family is proven, they will be in our scope in this research. However, that
7 reliability claim is invalidated in some later studies [11]. Figure 1 shows the main
8 members of the xorshift family of RNGs according to their development timeline.

9 Apart from the TestU01 suite (and the NIST suite), there is another measure of
10 randomness for some (especially for RFID-compatible) lightweight RNGs, defined in the
11 EPC™ UHF RFID Generation-2 standardization document. It is used to test the suitability
12 of our ultimate generator for UHF RFID devices.

13 One of the parallel researches by Vigna [12] revealed another xorshift based RNG, called
14 xorshift*. The RNG performed well when compared to its predecessors, but was worse
15 than the successors. However, it showed that xorshift family is flexible and is very open
16 to further development.

17 After the famous and widespread “Mersenne Twister” [13], in a follow-up study, Saito
18 and Matsumoto came up with another similar algorithm “XSadd”, which is claimed to be
19 more reliable. This claim is verified according to the results of the tests of TestU01 suite
20 [14], however it is reported that, inverse of this algorithm fails some of the tests of the
21 same suite.

22 A later and yet very recent study of Vigna again [14], introduced another xorshift based
23 RNG, “xorshift+” claimed to pass all tests of the TestU01 suite, and even inverse. Which
24 is not an invalidated claim in any newer publication we could analyze. For this obvious

1 reason, this RNG is chosen as a starting point in our study. In later sections of this paper,
2 it is referred as the original xorshift+.

3 **3. Candidate RNGs**

4 Since the computational limitations may be very stringent on lightweight devices, many
5 of the complex algorithms considered as successful in related works are not considered
6 as eligible for this purpose. Xorshift+, a recent algorithm of the known xorshift family is
7 in the focus with its simple structure [14]. It is also a main part of the study to force this
8 algorithm to run on a variety of input and output sizes. Moreover, the original xorshift+
9 algorithm is manipulated by removing some last steps to make it even more compact
10 O’Neill also mentioned this reduction idea [11], though not tested systematically.

11 Within this work, we have made random scramblings (based on our predictions) of the
12 original xorshift+ and produced many generators with slight differences, in order to
13 derive a better xorshift RNG. Nevertheless, we have eliminated many of these, which are
14 considerably weak or do not sufficiently mitigate the complexity. Our efforts to find a
15 suitable reduced xorshift+ version (which is to be named as xorshiftR+, where R stands
16 for reduced) resulted these 3 candidate algorithms (later called as variations or shortly
17 var), given as follows (in C notation);

```
18 uint64_t s[2]; // seeds
19 uint64_t xorshift128plus(void) {
20     uint64_t x = s[0];
21     uint64_t const y = s[1];
22     s[0] = y;
23     x ^ = x << 23; // a, shift & xor
24     x ^ = x >> 17; // b, shift & xor
```

```

1      //-----
2  Var1: x ^ = y ^ (y >> 26); // c, xor
3      s[1] = x;
4      return x + y;
5      //-----
6  Var2: x ^ = y ^ (y >> 26); // c, xor
7      s[1] = x + y;
8      return x + y;
9      //-----
10 Var3: x ^ = y ^ (y >> 26); // c, xor
11      s[1] = x + y;
12      return x + y; }

```

13 The body of the code above (from beginning, until the first dashed line) is untouched and
14 common for all variations (in fact step c is same for all, too), while the struck code
15 snippets are removals from and the black-highlighted parts are additions to the original
16 code of the xorshift+. The variations given under three consequent tags are applied singly
17 to the body of the code, so that a full RNG algorithm is obtained.

18 Ending variation 1 excludes 26 right shifts and 1 xor operation per generated random
19 number. Ending variation 2 excludes 26 right shifts and 1 xor but adds 1 addition. Lastly,
20 variation 3 excludes 26 right shifts and 1 xor, removes 1 addition and adds 1 addition,
21 which makes it identical with variation 1 in terms of operation count. These manipulations
22 supposedly produce a big difference in applications where many numbers are generated
23 sequentially. However, from the codes themselves, it is not possible to accurately estimate
24 the time gain without experiments. That's because the run-time highly depends on the

1 processor type [15]. Thus, we make use of our test completion times to measure the
2 algorithms' performance on a lightweight IoT device, given in Section 5.

3 Memory gain, on the other hand, is easier to calculate or estimate. The memory usage of
4 generators are mostly dependent on the period of their output and the number of their
5 input (seed) states. These two factors are determinative of the number and the size of the
6 variables used in their source code. To make a standard, we fixed the integer types to 32
7 bits as mentioned above. But please note, the original xorshift+ (and our reduced
8 variations) uses two seed inputs which makes it occupying 64 bits of input memory, while
9 most RNGs use only one, so it is implanted using 64 bits of seed state. Nevertheless,
10 because of its very poor performance, we didn't use 32-bit seed/input version of the linear
11 feedback shift register (LFSR), but 128-bit version instead. The output kept as 32-bit to
12 fix the output sizes of all subjects. In order to see the effect of the output or the period of
13 the RNGs on their randomness scores, please refer to another paper [16].

14 Other than our xorshift+ variations, we also tested an earlier xorshift based algorithm,
15 namely xorshift64* [13], a simple linear congruential generator (LCG) and an LFSR
16 implementation to compare the results and estimate the benefits of using each RNG.

17 XSadd is intentionally excluded, since it is already superseded by xorshift+. A
18 comparative breakdown is provided for a brand new xorshift based RNG, called
19 xoroshiro128+ that is published at the time of writing of this paper. LCG is included
20 because of its popularity, even though not recommended as a reliable RNG anymore
21 [1,11]. It is based on the one implemented in TestU01 suite as "CreateLCG()". The LFSR
22 is included because of its different bit-level structure and high speed nature. It is also
23 based on the one introduced in the suite with the name "CreateLFSR113()".

1 In fact, we had momentous efforts to functionalize DES as in such a purpose as
2 recommended in RFC 1750 [3] with relevant modifications. Nevertheless, later we
3 discarded it during our pre-evaluation phase because of its extremely poor performance
4 results. It wasn't even close to be "lightweight" when compared to other RNGs presented
5 here. Information that explains why and how much DES, advanced encryption standard
6 (AES) and derivatives are slower than these lightweight algorithms can be found
7 elsewhere [17].

8 **4. Test methodology**

9 The RNGs are tested using the Big Crush battery of tests, which is the most
10 comprehensive test battery in the TestU01 suite. It contains 106 different tests and makes
11 a total of 160 runs (some tests are repeated with different parameters) [5,18].

12 L'ecuyer and Simard defined 3 failures out of 15 tests of Small Crush battery and 7
13 failures out of 38 tests of Rabbit battery as "clearly unacceptable". Which make 80% and
14 81.6% respectively. So, in order to be called as successful, any RNG must obtain a clearly
15 higher score than these, preferably 100% depending on the scenario. Nevertheless, the
16 rate of passed tests is not solely enough to decide, but the result values of each failed test
17 should also be taken into account. TestU01 checks p-values to determine if RNGs pass
18 certain tests or not, for all tests, with a pass interval of 0.001 to 0.999. It also introduces
19 two epsilon values, namely ϵ and ϵ_1 , which represent very small numbers (latter
20 being much smaller) and are practically equal to 0 [5,18]. Hence, if an RNG fails at least
21 one test with a value of ϵ (or ϵ_1) then that RNG will most likely to fail this test in
22 each run and possibly for all seeds, which means a systematic fail. But if the p-value is
23 very close to the limits, (while the term "very close" is essentially vague), for instance
24 $0.999 < p < 0.9999$, then the result is not clear. Even though, L'ecuyer and Simard propose

1 repeating the suspicious tests is the best practice to get rid of the suspicious results. Our
2 experiments show that in the case when a p-value falls in the range given above, the test
3 can be accepted as a pass without repeating it, especially if it is the only fail or there are
4 very few, since a repeat with a different seed will most likely to result a pass. However,
5 if the p-value is not so close to the limits, though not eps, then repeating the tests might
6 be necessary. Note that, the same seed most likely will cause exactly the same result in
7 further repeats. If repeating the tests is not feasible, these tests should be taken as fails.
8 Remember, the number of total test runs are more than the number of different tests, since
9 some of the tests are repeated with various parameters in all batteries [18].

10 Other than the randomness levels, since we also have lightness concerns, run time, seed
11 and output sizes of each algorithm are taken into account to determine their memory
12 usage. Memory usage is not in the center of the focus since no tests are seen as required
13 besides the register size calculations, which can be done using function definitions.

14 For evaluating the time performance, two measures are considered; total time lapsed
15 during the tests and time period required to generate 100 million numbers (not bits). Total
16 duration of the tests is a good indicator of speed of the generators because it represents a
17 comprehensive usage scenario. However, it is not always very consistent as number of
18 tests applied may differ for some RNGs depending on their algorithmic structure. Thus,
19 we also recorded the time required to create 100 million numbers to be sure about their
20 pure speed. This is done using the native timer function of the TestU01 suite, with the
21 default method GetU01. In order to improve the solidity of our tests and
22 recommendations, we have run most of the algorithms, namely xorshift based ones and
23 LFSR113, also on a WISP device ten thousand times sequentially and recorded the total

1 creation time likewise. This additional step made us sure, our claims, verified on a PC are
2 also valid for a real lightweight IoT device and vice versa.

3 When an RNG is forced to use different seed numbers without completing its regular run
4 for its full period, i.e. seed is changed after each random number output, then the RNG
5 will highly probable lose its ability to produce randomness (except for some, i.e. LFSR
6 derivative scramblers). In this case, degree of the randomness of the output will depend
7 on the randomness of the seed string, because every seed input will cause the RNG to
8 produce the same set of random numbers in its each run. This prevented us from using
9 hardware generated so-called true random numbers as a seed chain. Instead, we used only
10 a single seed (duplicated when RNG required multiple seeds) for each subject generator.
11 A seed of an RNG is only changed to renew a test if there are suspicious results.

12 **5. Empirical results**

13 In order to make the final evaluation, all subject generators are tested with Big Crush
14 battery of the TestU01 suite, which is the biggest and toughest battery in the suite, while
15 other batteries (like Small Crush and NIST) are used to make preliminary tests. As a
16 recall, Big Crush contains 106 tests and makes a total of 160 runs, since some of the tests
17 are repeated with different parameters. Furthermore, in order the find out the average time
18 required for generation of a single random number for each algorithm; 100 million
19 numbers on PC and 10,000 numbers on WISP are generated. Hence, the average time
20 required for 1 number is calculated and shown in the Table.

21 **5.1. Testbeds**

22 Our PC testbed was a generic server computer with Intel Xeon E5540 processor @2.53
23 GHz running on Windows utilizing 4 GB of RAM. Nevertheless, we observed that the
24 software of TestU01 only occupies a single core during the whole test runs. This leads a

1 more linear base time comparison and prevented us to use our bed at full rate, but this
2 might also be positive since a lightweight system, highly probable, will only have single
3 core dedicated to random number generation. Our IoT testbed is a standalone WISP5
4 device which includes a built-in MSP 430 processor. Results of the tests are as follows;

5 **5.2. xorshift+**

6 The original xorshift+, not so surprisingly, demonstrated excellent randomness according
7 to Big Crush tests as Vigna claimed [14]. It passed 105 of the 106 different tests and 159
8 of the 160 test runs in our initial attempt.

9 Here, the p-value for this test is very close to the limit and as explained in Section 4, it
10 can be marked as a false positive and be ignored. In fact, repeated attempts show that this
11 fail was not systematic, but a coincidence. So, the ultimate scores might also be taken as
12 100% for both measures. We still kept this fail in our records for the sake of temporal
13 consistency. It was however, somewhat slow in terms of speed on the PC and the slowest
14 on WISP, when compared to other RNGs, as can be seen in the Table, which makes our
15 efforts more significant.

16 **5.3. xorshift***

17 The original xorshift* obtained, not surprisingly, slightly lower randomness scores on Big
18 Crush suite, though not bad. It failed four tests systematically. On PC, its speed was
19 unsatisfactory, actually slightly the worst among all. However, on the WISP device it was
20 (only) better than xorshift+. Please see the Table. Hence, it cannot be recommended as a
21 lightweight IoT solution, but might be used on PC platforms depending on the scenario.

22 **5.4. Reduced xorshift+ var1 (s[1] = x, return x + y)**

23 Our first variation obtained test scores very close to the original xorshift*, they also share
24 81. LinearComp test as a fail. On the other hand our variation was super-fast. On PC, it

1 was the fastest (along with the variation 3) among other xorshift based subjects. It failed
2 five tests systematically.

3 From the p-values of the failed tests, we can say that these are clear fails, because they
4 are far out of the limits. However, considering its speed, this might still be a useful
5 generator since its score is much higher than the (bad) examples given in Section 4 (which
6 were 80% and 81.6%).

7 **5.5. Reduced xorshift+ var2 (s[1] = x + y, return x + y)**

8 Our second variation got excellent scores (which is 100%) from Big Crush suite by
9 passing all the tests in all runs with no exception. Even though it wasn't the fastest, it was
10 still clearly faster than the original xorshift+ and the xorshift*. This agility difference is
11 much more obvious on WISP, rather than the PC. On PC, all our variations showed very
12 similar results.

13 **5.6. Reduced xorshift+ var3 (s[1] = x + y, return x)**

14 This third variation (which we have named as xorshiftR+) demonstrated excellent scores
15 (again 100%) by passing all the tests in all runs without any exception, too. In fact,
16 variation 3 was slightly faster than the variation 2 both on PC and on WISP. It was also
17 slightly faster than the variation 1 on WISP, while having the same speed on PC. Hence,
18 it is the fastest of its kind. Considering its perfect and consistent randomness scores and
19 very good time performance, it can safely and properly replace the original xorshift+.
20 Because, there is no statistics, in which the original xorshift+ is better than our variation
21 3 as can be seen in the Table. A bitmap graph using TestU01's PlotUnif function is seen
22 in Figure 2. It shows how the bits are randomly distributed under function's drawing rules.
23 Even though the TestU01 suite is already superior to the well-known NIST suite, because
24 of its widespread use and fame, we have tested xorshiftR+ using the NIST suite, too. The

1 result was a complete pass for all instances of all tests in the suite, as expected. The test
2 was made with 1,000,000 generated 64-bit numbers. The detailed documentation
3 regarding the test results is given online at: <http://srg.cs.deu.edu.tr/publications/2017/xor/>
4 To assure the algorithm's compliance with the RFID tags' security standards, three
5 conditions mentioned in the EPC™ Gen-2 Class 1 document were tested, too. The first
6 implies that the probability of a single 16-bit random number should be $0.8/2^{16} < P(\text{RN16})$
7 $< 1.25/2^{16}$ for 2^{30} numbers. xorshiftR+ satisfies this condition, moreover, even for 2^{26}
8 numbers, the result was $0.852/2^{16} < P(\text{RN16}) < 1.18/2^{16}$. The second mentions that the
9 probability of simultaneously identical sequences for 10000 tags should be $< 0.1\%$. With
10 xorshiftR+, since there are two 64-bit seeds, this probability is calculated as $(1/2^{64}) \times$
11 $(1/2^{64}) = 1/2^{128}$ then the result is, $[10000 \times (1/2^{128})] * 100 = 2.94\% \times 10^{-34} < 0.1\%$. Lastly,
12 the third implies that an RN16 drawn from a tag's RNG shall not be predictable with a
13 probability greater than 0.025%. This is proven via the ENT suite, another old yet popular
14 test suite. There, we have also used the original xorshift+ for the sake of a consistent
15 comparison; and xorshiftR+ demonstrated a very high performance that is practically
16 equivalent to the original xorshift+. Detailed information regarding these conditions, as
17 well as the test results can also be found on the link given above.

18 Overall, for xorshiftR+, NIST and EPC Gen-2 Class 1 (incl. ENT) examinations were
19 successful, too, in addition to TestU01.

20 **5.7. Simple LCG**

21 The simple LCG, implemented with parameters (2147483647, 16807, 0, 12345); was
22 very weak in terms of randomness, since it can only pass a little more than the half of the
23 tests and test runs, which is far below the unacceptable examples given in Section 4 and
24 the lowest among our subjects. Figure 3, generated by scatter of TestU01, reveals a

1 repeating pattern and thus clearly shows that the generated bits are not random. Total
2 number of test runs was lower because some tests were not repeated as a result of the
3 RNG's poor performance as described in Section 4. Names and parameters of the failed
4 tests were not presented because of their huge count. All fails were clear considering their
5 p-values. Note that, the run time for the suite given in the Table (which is very short)
6 might be misleading since the number of test runs were fewer than the others (because of
7 its poor randomness performance). Anyway, number creation times show that this RNG
8 is very fast in the scale (not tested on WISP). Despite its speed, LCG as implemented
9 here, is not recommended in any scenario and should not be used in applications, where
10 high level of randomness is required.

11 **5.8. LFSR113**

12 LFSR113, with parameters (12345, 12345, 12345, 12345); was the fastest RNG on the
13 PC among our candidates with a decent difference when compared to xorshift based
14 subjects including our variations. It was approximately 8% faster than all our variations
15 on PC. Nevertheless, it is remarkable that it demonstrated even better performance on the
16 WISP. Conversely, LFSR has indecisive test scores, which can be found acceptable
17 depending on the scenario, but surely not better than xorshift family. The generator failed
18 six tests systematically. This RNG is recommendable on PC and in a case where the
19 randomness requirements are loose but time limitation is a more dominant concern.

20 **5.9. xoroshiro+**

21 After we have concluded our tests and most of the study as well; a brand new xorshift
22 based generator, namely xoroshiro128+, was announced by Vigna. It is claimed that, this
23 is the fastest of all. Hence, we could not ignore and made a comparative analysis, though
24 we could not run a new TestU01.

1 xoroshiro128+ contains the following operations: 105 left shifts, 37 right shifts, 1 addition
2 (64 bit number), 2 or (64 bit numbers), and 2 xor (64 bit numbers). A total of 462 bitwise
3 operations. While our xorshiftR+ only contains, 23 left shifts, 17 right shifts, 2 xor (64
4 bit numbers), and 1 addition (64 bit number); which makes 232 bitwise operations. Since
5 types of the operations required are the same for both, and the number of operations for
6 xorshiftR+ are almost half of the ones for xoroshiro+; xorshiftR+ is clearly lighter and
7 expected to run faster. Thus, xorshiftR+ is apparently superior to xoroshiro+, too.

8 **5.10. Further notes**

9 As a further information, test results of the given xorshift based generators represent their
10 32-bit output versions, instead of the contemporarily used 64-bit (or more) versions,
11 where the integers and the outputs are 64-bit numbers. Because, on the lightweight family
12 of devices, 32-bit is a common upper limit. Presumably, 64-bit versions get the same or
13 better results from the tests. In contrast, 16-bit integer sized versions of the generators
14 (including the original xorshift+) fail almost all tests even in the Small Crush battery. This
15 is because, many tests of the TestU01 assume an output of at least 30 bits and fails the
16 subject generator otherwise [18]. Anyway, test results of these 16-bit versions are still
17 very weak (fail most of the tests in the suite) and are not statistically valuable. This also
18 explains why standard rand() function of C language (initiated by the srand(time(NULL))
19 code) fails all the tests in the Big Crush suite (obtaining a score of 0% in both measures).
20 Because of their very poor performance, rand() and simple LCG are not tested on WISP.
21 Please note, the runtime or number generation time difference between the RNGs may
22 look small or even negligible, but surely not. When slower machines are used as hosts or
23 extensive repeats will be made, these differences shall become more obvious. All three
24 of our reduced variations were faster than the xorshift+ itself. But, our variation 1 couldn't

1 pass all the tests and so is discarded during the final evaluation. Even though variation 2
2 also passed all the tests, yet variation 3 was the fastest among all our variations and
3 consequently selected as the xorshiftR+. It can be used wherever xorshift+ can be used
4 and it can safely replace the xorshift+ in further studies and applications.
5 Additionally, xorshift* can be considered as deprecated after xorshift+ (and xorshiftR+)
6 is introduced, this was also mentioned in [14]. The simple LCG was expectedly very weak
7 (in terms of randomness) and should not be used in any scenario without further
8 modifications. LFSR was very fast however, weaker than all the xorshift family members
9 given here. Lastly, bitmap graphs are not drawn for other subject generators, since it
10 wasn't possible for human eye to detect a pattern from them as well. So, only the best and
11 the worst generators' graphs are shown (See Figures 2 and 3).

12 **6. Conclusion**

13 According to our experiments, the xorshiftR+ (also mentioned as reduced xorshift+
14 variation 3, R stands for reduced, pronounced xorshifter plus) has demonstrated
15 outstanding results when compared to other candidates, in terms of both randomness and
16 speed performance. It passes all the tests in the Big Crush battery of TestU01, in all runs
17 and is slightly (but noticeably) superior to its ancestors, especially the xorshift+, with its
18 higher speed. This superiority becomes even more obvious on lightweight environments.
19 It also occupies less memory due to its compacted nature. Good to mention that it also
20 passes all tests of NIST and ENT suites and complies with EPC-Gen 2 Class 1 security
21 standards. This RNG, xorshiftR+, is strongly recommended to use in both powerful
22 computer environments and lightweight devices, where high level of randomness and
23 temporal performance are desired. As an addition to the conclusion given in [14], our
24 experiments showed that, even if integer sizes are set to 32 bits (this makes 64-bit

1 input/seed and 32-bit output), xorshift+ and xorshiftR+ can still pass all TestU01 tests in
2 all runs systematically.
3 LFSR113 on the other hand, was the fastest on PC, though wasn't able to pass all tests.
4 Plus, it occupies vast memory. As its original form (used here), it can only be
5 recommended when randomness requirements are loose but time limitations are
6 dominant. If it is possible to make it fail-free hereafter, then it might be a good alternative
7 to xorshiftR+. Such improvements on the LFSR could definitely be a worthy future work.

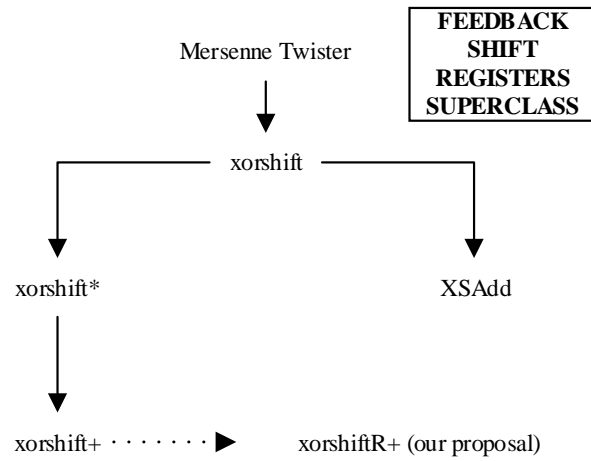
8 **Acknowledgement**

9 We sincerely thank to Hakan Altaş for his valuable efforts during our experiments.

10 **References**

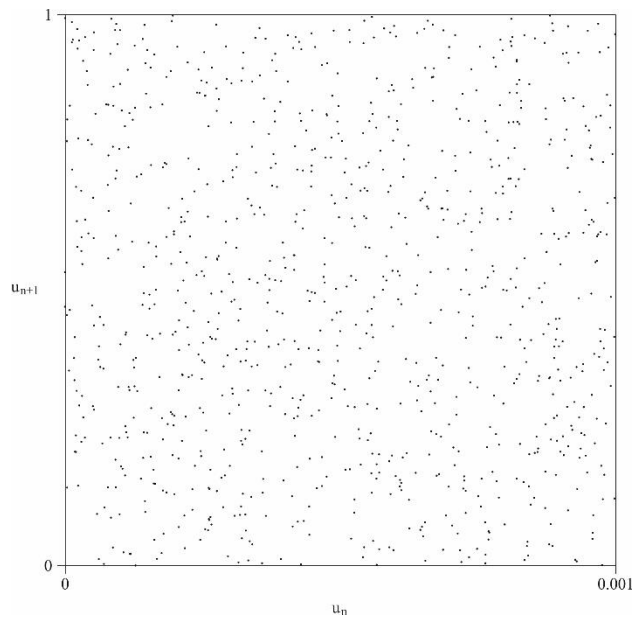
- 11 [1] Hellekalek P. Good random number generators are (not so) easy to find. *Math Comput*
12 *Simulat* 1998; 46: 485-505.
- 13 [2] Özcanhan MH, Dalkılıç G. Mersenne twister-based RFID authentication protocol.
14 *Turk J Elec Eng & Comp Sci* 2015; 23: 231-254.
- 15 [3] Eastlake D, Crocker S, Schiller J. Randomness recommendations for security. RFC
16 1750, 1994.
- 17 [4] Marsaglia G. The DIEHARD battery of tests of randomness. Technical Report,
18 Florida State University, 1995.
- 19 [5] L'Ecuyer P, Simard R. TestU01: A C library for empirical testing of random number
20 generators. *ACM T Math Software* 2007; 33: Article No. 22.
- 21 [6] Bassham LE, et al. A statistical test suite for random and pseudorandom number
22 generators for cryptographic applications. NIST special publication 800-22rev1a, 2010.
- 23 [7] Barker E, Kelsey J. Recommendation for random bit generator (RBG) constructions,
24 NIST special publication 800-90C second draft, 2016.

- 1 [8] Barker E, Kelsey J. Recommendation for random number generation using
2 deterministic random bit generators, NIST special publication 800-90A, 2012.
- 3 [9] Turan MS, Barker E, Kelsey J, McKay KA, Baish ML, Boyle M. Recommendation
4 for the entropy sources used for random bit generation, NIST special publication 800-
5 90B second draft, 2016.
- 6 [10] Marsaglia G. Xorshift RNGs. *J Stat Softw* 2003; 8: 1-6.
- 7 [11] O'Neill ME. PCG: A Family of simple fast space-efficient statistically good
8 algorithms for random number generation. *ACM T Math Software* V; 46 pages.
- 9 [12] Vigna S. An experimental exploration of Marsaglia's xorshift generators, scrambled.
10 ArXiv e-print 2014.
- 11 [13] Matsumoto M, Nishimura T. Mersenne twister: a 623-dimensionally equi-distributed
12 uniform pseudo-random number generator. *ACM T Model Comput S* 1998; 8: 3-30.
- 13 [14] Vigna S. Further scramblings of Marsaglia's xorshift generators, *J Comput Appl*
14 *Math* 2016; 315:175-181.
- 15 [15] Singh KP, Kumar D. Performance evaluation of low power MIPS crypto processor
16 based on cryptography algorithms. *Int J Eng Res Applic* 2012; 2: 1625-1634.
- 17 [16] Brent RP. Some long-period random number generators using shifts and xors.
18 *Anziam J* 2007; 48: 188-201.
- 19 [17] Thakur J, Kumar N. DES, AES and Blowfish: Symmetric key cryptography
20 algorithms simulation based performance analysis. *Int J Emerg Technol Adv Eng* 2011;
21 1: 2250-2459.
- 22 [18] L'Ecuyer P, Simard R. TestU01: A software library in ANSI C for empirical testing
23 of random number generators user's guide document. Canada, 2014.



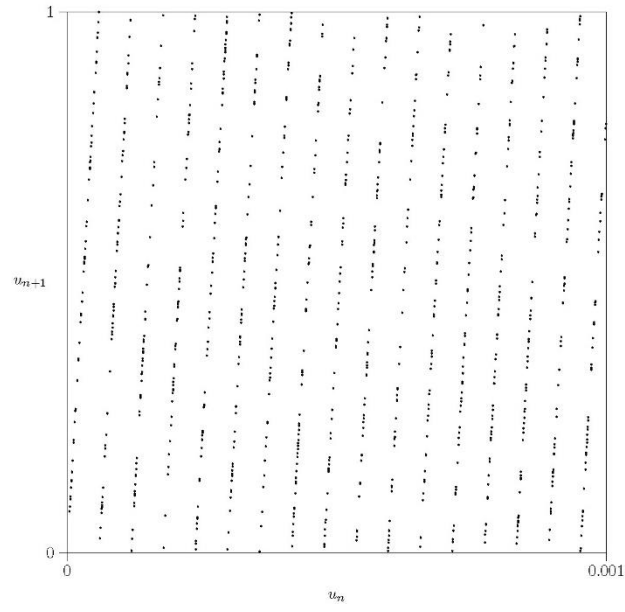
1
2

Figure 1. Development lapse of (some) xorshift based RNGs.



3
4
5

Figure 2. Bitmap of xorshiftR+; 1000 points plotted using scatter in TestU01 (u_n and u_{n+1} stand for sequential random numbers. See definitions in [18]).



1

2 **Figure 3.** Bitmap of simple LCG; 1000 points plotted using scatter in TestU01 (u_n and

3 u_{n+1} stand for sequential random numbers. See definitions in [18]).

1 **Table.** Comparative data of empirical test results from Big Crush for all subject RNGs.

Empirical Results	Big Crush score (# of runs)	Big Crush score (# of tests)	Run time of Big Crush	Time for 1 number on PC	Time for 1 number on WISP
xorshift+	159/160 99.3%	105/106 99.1%	7h:35:03	0.02 μ s	0.5515 ms
xorshift*	156/160 97.5%	102/106 96.2%	7h:28:10	0.0202 μ s	0.2703 ms
Reduced xorshift+var1	155/160 96.9%	103/106 97.2%	7h:27:02	0.0191 μ s	0.1955 ms
Reduced xorshift+var2	160/160 100%	106/106 100%	7h:32:40	0.0192 μ s	0.1954 ms
Reduced xorshift+var3 [xorshiftR+]	160/160 100%	106/106 100%	7h:30:07	0.0191 μ s	0.1953 ms
Simple LCG	91/153 59.5%	57/106 53.8%	7h:25:45	0.0177 μ s	N/A
LFSR113	154/160 96.3%	100/106 94.3%	6h:47:57	0.0175 μ s	0.1641 ms
C rand()	0/160 0%	0/106 0%	N/A	N/A	N/A

2