

Romu: Fast Nonlinear Pseudo-Random Number Generators Providing High Quality

Mark A. Overton
Independent Researcher

Abstract

We introduce the Romu family of pseudo-random number generators (PRNGs) which combines the nonlinear operation of rotation with the linear operations of multiplication and (optionally) addition. Compared to conventional linear-only PRNGs, this mixture of linear and nonlinear operations achieves a greater degree of randomness using the same number of arithmetic operations. Or equivalently, it achieves the same randomness with fewer operations, resulting in higher speed. The statistical properties of these generators are strong, as they pass **BigCrush** and **PractRand** – the most stringent test suites available. In addition, Romu generators take maximum advantage of instruction-level parallelism in modern superscalar processors, giving them an output latency of zero clock-cycles when inlined, thus adding *no* delay to an application. Scaled-down versions of these generators can be created and tested, enabling one to estimate the maximum number of values the full-size generators can supply before their randomness declines, ensuring the success of large jobs. Such capacity-estimates are rare for conventional PRNGs.

A linear PRNG has a single cycle of states of known length comprising almost all possible states. However, a Romu generator computes pseudo-random permutations of those states, creating multiple cycles with pseudo-random lengths which cannot be determined by theory. But the ease of creating state-sizes of 128 or more bits allows (1) short cycles to be constrained to vanishingly low probabilities, and (2) thousands of parallel streams to be created having infinitesimal probabilities of overlap.

Keywords: PRNG, random number generator, nonlinear.

1. Introduction

There are many noncryptological uses of random numbers in software, such as Monte Carlo simulations, statistical analysis, testing software and hardware, random actions in games, producing textures and grain in photos, and others. But computers lack an auditable source of true randomness in them, so random numbers are often obtained from external sources such as the times at which unpredictable events occur, including packet-arrivals, key-presses, and mouse-movements. But such sources produce random numbers slowly, so when they are needed quickly, they must be synthesized with arithmetic. Such an arithmetic algorithm is referred to as a *pseudo-random number generator* (PRNG). This paper uses the term *generator* when referring to a PRNG.

Mascagni and Srinivasan (2004) wrote, “While the quality of a PRNG sequence is extremely important, significant aspects of a generator’s quality are hard to prove mathematically. Therefore, the ultimate test of a PRNG’s quality is empirical”. Accordingly, several test-

suites for generators have been developed. Some well known ones include **Diehard** (Marsaglia 1995), **Dieharder** (Brown, Eddelbuettel, and Bauer 2011), **TestU01** (L’Ecuyer and Simard 2007), and **PractRand** (Doty-Humphrey 2018). **TestU01** sports three batteries of tests, among which **BigCrush** is the most thorough, consuming about 2^{35} 32-bit values from the generator under test and taking four hours. **PractRand** tests 2^{45} bytes by default, making it even more thorough than **BigCrush**, taking over three days. **BigCrush** and **PractRand** can identify subtle flaws in a generator, so they were used to test the generators presented in this paper. We relied most heavily on **PractRand** (version 0.94) because it provides a measure of goodness of a generator, and not merely a pass/fail indication.

This paper presents a new kind of pseudo-random number generator named *Romu*, which stands for the words *rotate-multiply*, the two primary operations it performs. The key strengths of generators in the Romu family are:

- Randomness: They have excellent statistical properties, verified by passing the most stringent tests of randomness (**BigCrush** and **PractRand**).
- Speed: They require unusually few instructions and take advantage of instruction-level parallelism (ILP) offered by modern superscalar processors, making them fast. Their tiny code can be inlined, making them even faster.
- Streams: They can generate thousands of independent streams for large parallel jobs.
- Assurance: Scaled-down versions can be tested until they fail, providing estimates of the capacities of the full-size generators.

Our naming convention for these generators appends a suffix to *Romu* indicating how many internal state-variables a generator has. *Mono*, *Duo*, *Trio*, and *Quad* signify 1, 2, 3, and 4, respectively. For example, the state of RomuTrio is held in three 64-bit variables (192 bits). The name of a Romu generator employing 32-bit arithmetic is also suffixed with *32*.

To illustrate the high degree of randomness produced by the Romu family, we will compare the simplest kind of Romu generator with the popular linear congruential generator (LCG) using dot-plots. The two images in Figure 1 show all successive pairs of outputs for an entire period of (1) an LCG with a multiplier of 477, and (2) a two-operation Romu generator with a rotation of 4 and multiplier of 715. The arithmetic of both generators was modulo 2^{10} . The left plot in Figure 1 is the usual lattice produced by an LCG. The right plot produced by the simple Romu generator is clearly more random.

The C source-code of a 64-bit version of the same type of two-operation Romu generator is presented below. ROTL is a macro or template function that rotates bits leftward.

```
uint32_t result = state; // state is uint64_t.
state = ROTL(state,32) * 15241094284759029579u;
return result; // return low 32-bits of 64-bit state.
```

This generator, named RomuMono, consists of only *two* assembler-language instructions, making it very fast. Yet it passes both **BigCrush** and **PractRand** (up to 2^{46} bytes). This simple generator reveals the great randomizing power available in a rotation combined with a multiplication. RomuMono is not discussed further due to its small state-size of 64 bits.

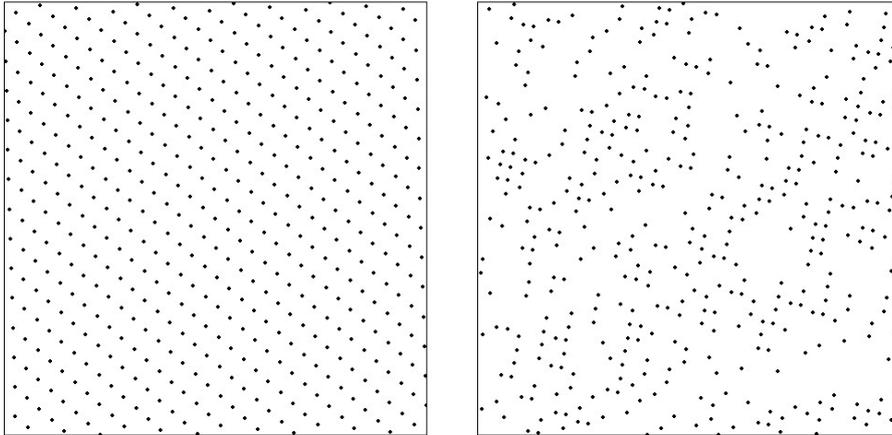


Figure 1: Dot-plots of LCG (left) and Romu (right).

Unlike popular generators such as the LCG, Romu generators do *not* have a single cycle of states encompassing all possible states. Instead, they create a pseudo-random permutation of all possible states, and consequently possess multiple cycles of pseudo-random lengths, including some that are too short. The mathematics of these cycles is intractable, so it is not possible to compute the length of the cycle containing the seed.

Fortunately, the probability distribution of cycle lengths is known, and we can rely on such theory to design generators with infinitesimal probabilities of short cycles. Furthermore, using a convolution of probability distributions, we can compute the probabilities of overlap of streams, and thus be assured that the possibility of either a short cycle or stream-overlap can be safely ignored.

Fog (2001) proposed the RANROT family of generators which are similar to Romu in that they are nonlinear with probabilistic cycle lengths. But his generators employ rotation and *addition*, whereas Romu generators employ rotation and *multiplication*. Fog rejected multiplication because it was too slow at that time, but we appreciate his insights.

Note that every Romu generator presented in this paper is *not* cryptographically secure because it is invertible. If an attacker knows its internal state, its prior outputs can be computed by running the generator backwards, allowing the attacker to decrypt prior messages.

2. Requirements of Generators

In addition to producing good random numbers, we impose several requirements on modern generators. A generator must provide thousands of *parallel streams* that do not overlap. To achieve high speed, a generator must make good use of *ILP* and add a minimum of *register pressure*. Finally, a generator must have known and sufficient *capacity*.

2.1. Parallel streams

Modern processors contain multiple cores, each of which can run a program or thread independently of the other cores. High-end computers contain hundreds of processors and several

thousand cores. To take advantage of this available parallelism, large simulations can be divided into sub-jobs to be run in parallel, one per core. Such parallel computation requires that a generator supply an independent stream of values for each of the many cores. There are two popular methods of creating streams: cycle-splitting and parameterizing.

Cycle-splitting

Mascagni and Srinivasan (2004, p. 2) mention and dismiss the “naive” method of creating streams in which “users randomly select seeds on each processor and hope that the seeds will take them to widely separated portions of the sequence, so that there will be no overlap between the sub-sequences used by different processors.” However, for a Romu generator with at least 128 bits of state, we claim that this naive approach is excellent. Table 3 tells us that such a generator can provide tens of thousands (and usually far more) sub-sequences (streams) with an infinitesimal chance of any overlap occurring among them. Cycle-splitting is well-suited for Romu generators.

Some generators offer a jump-function allowing one to jump ahead in the cycle by some large number of values. This method of cycle-splitting results in *no* chance of overlap.

Parameterizing

A second method of providing multiple, independent streams is to use a different multiplier or rotation for each stream produced by a generator. A severe disadvantage of creating streams by changing multipliers or other constants is that it is not practical to test every combination of constants. **PractRand** requires at least three days to test a generator up to 2^{45} bytes, making it impractical to test thousands of generators. Thus, when employing parameterization, the generator will be *untested*, creating an unknown risk to the job. With at least 128 bits of state, cycle-splitting virtually guarantees separation of streams in a *tested* generator, suggesting that cycle-splitting is the safer approach.

2.2. Instruction-level parallelism

Good generators rely on instruction-level parallelism (ILP) in the processor to attain high speed. Modern PC processors can typically execute up to four instructions in one clock cycle, so they are known as *four-issue superscalar* designs. A clock cycle can be regarded as having four slots in which an instruction can issued (executed). In each clock cycle, while instructions are being executed, the processor looks ahead in the instruction stream, analyzes data-dependencies among upcoming instructions, and fills as many slots as possible with instructions whose input-values will be known in the next clock cycle.

ILP is illustrated with the following RomuTrio generator which is presented in detail later.

```
uint64_t romuTrio_random () {
    uint64_t xp = xState, yp = yState, zp = zState;
    xState = 15241094284759029579u * zp;
    yState = yp - xp;  yState = ROTL(yState,12);
    zState = zp - yp;  zState = ROTL(zState,44);
    return xp;
}
```

This generator consists of five arithmetic operations: a multiplication, two rotations, and two subtractions. Thanks to ILP, these five instructions consume only three clock cycles as shown in the following ILP table.

| Cycle | Slot 1 | Slot 2 | Slot 3 | Slot 4 | Multiplier |
|-------|------------|------------|----------|--------|------------|
| 1 | subtract x | subtract y | multiply | | * |
| 2 | rotate y | rotate z | | | * |
| 3 | | | | | * |

The *Multiplier* column of an ILP table shows when the processor’s multiplier is busy. A single multiplication consumes three cycles in today’s Intel processors, which is why most Romu generators presented in this paper consume three cycles. If multiplication drops to two cycles in future processors, the speed of some Romu generators will increase proportionately. A core in an Intel PC CPU has only one multiplier, but it supports overlapping (termed *pipelining*), so the core can start exactly one new multiplication in each clock cycle. Thus, up to three multiplications can be active in a given cycle. A constraint in Intel processors is that at most two shifts and/or rotations can be performed in one clock cycle.

In the RomuTrio generator, there are three sequences of independent computations: a multiplication and two subtract-rotate pairs. Because these sequences do not depend on each other, each sequence is executed concurrently as seen in the above ILP table, greatly increasing performance.

Timing measurements confirm this understanding of ILP. Table 1 shows the execution times of combinations of one, two, and three simple RomuMono32 generators (presented later) as measured on an Intel Xeon E5-2643 v3 (3.4-3.7 GHz), which is the type of processor most commonly seen in workstations used by engineers and scientists. The table shows the number of milliseconds and clock cycles consumed by 5×10^9 invocations of the combination generator. Each base generator consists of only two instructions (rotate, multiply), and theoretically consumes 4 clock cycles. (Most generators presented in this paper are structured to consume only 3 clock cycles.) The source-code of the combination generator in row 2 of the table is shown below.

```
x = ROTL(x,r1) * c1;      // generator 1
y *= c2;  y = ROTL(y,r2); // generator 2
```

Note that the two instructions in the second generator are reversed (multiply, then rotate) to take advantage of pipelining in the multiplier.

| Row | Generator | mSec | Cycles |
|-----|------------------------|------|--------|
| 1 | One RomuMono32 | 6238 | 4.6 |
| 2 | Two RomuMono32 Combo | 6226 | 4.6 |
| 3 | Three RomuMono32 Combo | 6562 | 4.9 |

Table 1: Execution times.

Rows 1 and 2 of Table 1 show that two Romu generators in combination consumed the same time as one (within measurement error), revealing that ILP achieved 100% parallelism. Row 3

reveals that even three generators achieved nearly 100% parallelism. When used well, ILP boosts performance severalfold.

2.3. Low register pressure

The term *register pressure* refers to the number of registers the compiler wants to use at a specific point in a program. A random number generator is often called from an inner loop, and to improve performance, the generator can be inlined. Consequently, all registers devoted to the generator add to the register pressure inside the loop, and if it exceeds the number of registers in the processor, some variables in the loop will be stored in memory instead, referred to as *spilling*, reducing speed. To increase ILP, the designer of a generator might arrange for a calculation to be independent of the others, which will require more registers, causing more spills and reducing speed, defeating the purpose of increasing ILP. Thus, a compromise must be made between ILP and register pressure, as both affect the speed of the application.

We regard the usage of up to 5 registers by a generator to be good, 6 as borderline, and 7 or more as excessive. Most of the Romu generators presented in this paper are acceptable, as they require 4 to 6 registers when inlined in a loop. The `xoshiro256++` generator (Blackman and Vigna 2018) requires 7-8 registers because it keeps 4 state-variables in registers and its high level of ILP consumes 3 or 4 more. The 128-to-64 bit variant of the `pcg-xsh-rs` generator (O’Neill 2014) requires 10 registers because its 128-bit operations need twice as many registers to hold operands.

The best ways to minimize register pressure are (1) minimize the number of variables in which the generator’s state is stored, and (2) do not allow such a variable to exceed the processor’s word-size (usually 64 bits). For best performance, we suggest avoiding a generator containing over four state-variables or performing 128-bit or wider arithmetic.

2.4. Known capacity

The *capacity* of a generator is the quantity of data it can produce before statistical flaws become detectable. Our testing of multiple types of generators revealed that a generator’s capacity is almost always less than its period – often far less. Therefore, a user is wise to require that a generator’s capacity be measured or estimated.

However, one must keep in mind that a generator’s period is usually more important than its capacity because exceeding its period has worse consequences. Exceeding its capacity causes subtle correlations to become detectable only in a large quantity pseudo-random data, which will cause little to no harm to many kinds of jobs. But exceeding its period causes previously-supplied data to be supplied again, which is a severe correlation that can significantly distort results.

Exceeding capacity is not the same as *overtaxing*, a concept introduced by O’Neill (2014). Overtaxing refers to exceeding a generator’s period, and the solution is to increase its period. But increasing its period will not increase a generator’s capacity, although increasing its state-size can do so.

The existence of multiple cycles (with various periods) in a Romu generator allows us to graph its capacity as a function of period. Scaled-down versions of the Romu generators presented later in this paper were created such that their state size was 32 bits (30 bits for RomuTrio). The generators RomuQuad, RomuTrio, and RomuDuoJr were selected for this investigation.

RomuQuad is the best generator we present and RomuDuoJr is the worst, consisting of only three arithmetic instructions. The small 32-bit size made it possible to traverse all of their cycles, which is not possible with larger state-sizes, and to run **PractRand** on each of those cycles and record the number of values output just before the generator began exhibiting anomalous behavior. The results are plotted in Figure 2.

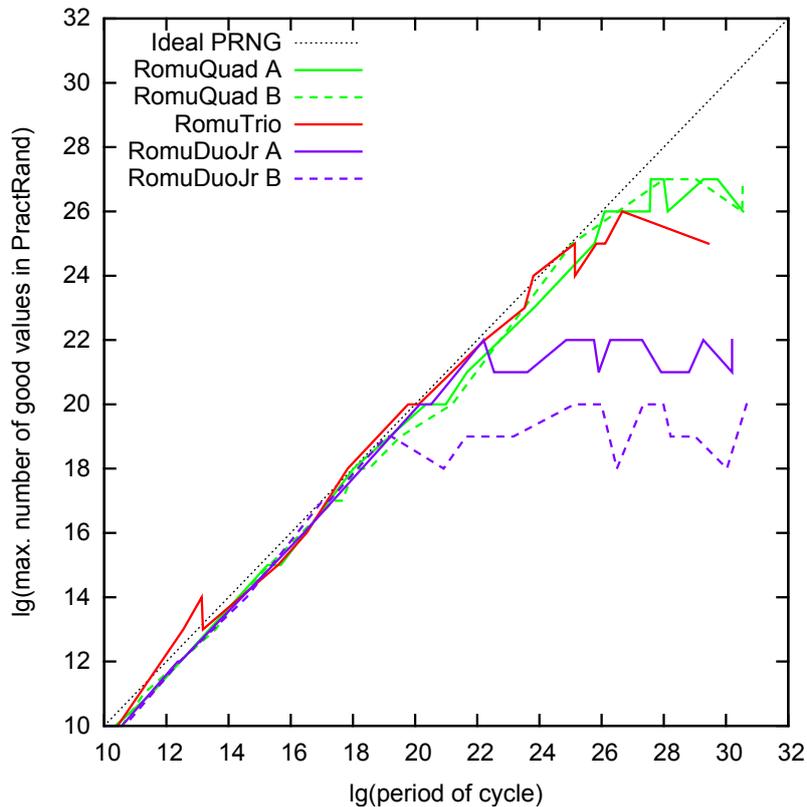


Figure 2: Good values versus cycle-period for Romu generators.

The B-versions in the graph used different multipliers than the A-versions. The horizontal axis is the base-2 logarithm of period, and the vertical axis is the base-2 logarithm of the maximum number of nonanomalous values (not bytes) supplied by the generators.

The graph reveals that each of these scaled-down generators worked well to a point, and then leveled onto a bumpy plateau in which continuing to boost period (by selecting longer cycles) failed to boost the number of good values. The plateau is the generator’s capacity. The arithmetic in RomuQuad is more complex (6 operations) than that in RomuDuoJr (3 operations). The graph shows that RomuQuad has greater capacity, as would be expected.

An important observation is that before exceeding their capacities, all Romu generators stayed close to the theoretical ideal generator identified by the dashed 45-degree line. That is, up to a period of 2^p (which must not exceed capacity), each generator could output close to 2^p values before **PractRand** could detect statistical flaws. This behavior tells us several things:

- Up to its capacity, arithmetic employing a multiplication with rotation(s) produces excellent randomness.

- With Romu generators, one need not follow the rule of thumb that the number of values used should not exceed the square-root of the generator’s period. Figure 2 tells us that most of the period is usable (up to the generator’s capacity).
- If a Romu generator happens to be in a small cycle, it will perform nearly ideally. Thus, if a job (or stream) requires 2^k values, the generator will perform well if the period is not under 2^{k+1} . The 1 was added as a safety margin, and 2^{k+1} was assumed not to exceed capacity.

As an example using numbers for an impossibly huge job, if a generator’s state size s is 192 bits, and $k+1$ is 59, then equation (3) tells us that the maximum probability that the period will be below 2^{59} is $2^{59-192+7} = 2^{-126}$, which is the probability of randomly selecting a given grain of sand in all the beaches of the world *twice in succession*. This is incomprehensibly tiny, assuring us that such a generator will always perform well.

Capacities of Romu generators cannot be determined by testing, as doing so would require centuries. Therefore, they were estimated based on tests of half-scale, quarter-scale, and eighth-scale generators. For generators using 64-bit arithmetic, the state-variables in these scaled-down generators were 32, 16, and 8 bits each. Some of these mini-generators were not the same as those used to produce Figure 2, as they were not restricted to 32 bits of state. Due to their smaller state-sizes, it was possible to run these mini-generators until they failed. The number of good values a mini-generator produced in a test was its capacity.

Estimated capacities for the corresponding full-size generators were extrapolations of the mini-generator capacities, requiring one to three doublings of the sizes of the mini-generators. For each such doubling, the base-2 logarithm of the tested capacity was multiplied by 1.4. To be conservative, 1.4 was used in lieu of 2.0 because it was observed that doubling the number of state bits in a generator (by doubling its word-size) multiplies that logarithm by a factor of 1.4 to 1.7, which is less than the 2.0 suggested by intuition.

3. Theory of Romu Generators

The values produced by linear generators consist of one cycle of known length. These generators are mathematically tractable, and thus much is known about them. The LCG and MCG are the most widely known linear generators. Xorshift (Marsaglia 2003) and the generators from Blackman and Vigna (2018) are more recent examples. But such generators create linear artifacts (such as a lattice structure) in their output values, a serious problem that greatly reduces the usable portion of their periods (L’Ecuyer and Simard 2001). Gentle (2003, p. 21) recommends not exceeding the square-root of a generator’s period, representing a severe reduction. However, it is important to note that the Blackman-Vigna generators, the LCG-based PCG (O’Neill 2014), and enhancements to xorshift perform a post facto scrambling step on the output value (not affecting state), eliminating their linear artifacts and greatly increasing their capacities.

Romu is nonlinear in the field of integers modulo 2^w ($\mathbf{Z}/2^w\mathbf{Z}$), where w is the processor’s word-size. Romu’s multiplication (and any additions) is linear in this field, but its rotation is not, resulting in a nonlinear function which is mathematically intractable. Mathematical analysis of a *linear* generator provides several benefits:

1. The generator's period can be computed.
2. A jump formula can be derived, from which a jump-ahead function can be coded.
3. In some cases, lattice spacings can be computed, providing measures of randomness.

On the other hand, the benefit of *nonlinear* generators is their freedom from linear artifacts, thus improving randomness and eliminating the need for a scrambling step, increasing speed. But without the above benefits of analysis, their periods are unknown, jump-ahead is impossible, and their randomness must be gauged solely by testing. Also, as mentioned above, Romu generators have multiple cycles of differing lengths. For this reason, the period obtained by an arbitrary seed is unpredictable, but is probabilistic as will be seen next.

3.1. Multiple cycles with probabilistic periods

The integers in set $S = \{1, \dots, 2^s - 1\}$ are all of the allowable states of a Romu generator having s bits of state. Note that the all-zeros state is excluded. Every Romu generator is one-to-one and onto (i.e., a bijection). Thus, it maps every member of S to every member of S . Consequently, a Romu generator is invertible: Given its state, its prior state can be computed. Also, the mapping is a permutation of the members of S , meaning the generator rearranges the integers in $\{1, \dots, 2^s - 1\}$. Crucially, this permutation by a Romu generator is pseudo-random. All permutations consist of multiple cycles of various lengths (Knuth 1973, pp 160-161), and every state is in some cycle. Fog (2001, eq. 8) and Knuth (1973, p. 518 ex. 17) have disclosed that, for random permutations of the integers $\{1, \dots, n\}$, the expected probability that a value x is in a cycle of length m is $1/n$. Curiously, this probability is independent of m . We can verify this counterintuitive statement with the following derivation:

$$\begin{aligned} \text{E}(\text{P}(|\text{cycle containing } x| = m)) &= \frac{n-1}{n} \cdot \frac{n-2}{n-1} \cdot \frac{n-3}{n-2} \cdots \frac{n-m+2}{n-m+3} \cdot \frac{n-m+1}{n-m+2} \cdot \frac{1}{n-m+1} \\ &= \frac{1}{n}. \end{aligned}$$

Thus, the expected probability that the length of the cycle is less than or equal to m is simply

$$\begin{aligned} \text{E}(\text{P}(|\text{cycle containing } x| \leq m)) &= \sum_{i=1}^m \frac{1}{n} \\ &= \frac{m}{n}. \end{aligned} \tag{1}$$

This equation is also presented in Fog (2001, eq. 15).

For any state x (i.e., the generator's seed) that was randomly selected from the $2^s - 1$ possible states in S , we wish to know the expected probability that the length of the cycle containing x is less than or equal to 2^k . We substitute $m = 2^k$ and $n = 2^s - 1$ into (1), yielding

$$\begin{aligned} \text{E}(\text{P}(|\text{cycle containing } x| \leq 2^k)) &= 2^k / (2^s - 1) \\ &\approx 2^{k-s}. \end{aligned} \tag{2}$$

In practice, an upper bound on this probability would be more useful than its expected value. A Romu generator employs a single, fixed set of cycles based on only one permutation because

its multiplier and rotation(s) are constants. The single cycle-set causes the actual probability for a given k and s to differ from the expectation given in (2). We propose that (3) is an appropriate upper bound on the actual probability. Note that (3) is identical to (2) except for the addition of 7 to the exponent.

$$P_{\text{ub}}(|\text{cycle containing } x| \leq 2^k) = 2^{k-s+7}. \quad (3)$$

Because the mathematics underlying Romu generators is intractable, it is not possible to ascertain the distribution of (2) or its upper bound based on theory. Therefore, the upper bound was determined empirically by creating scaled-down versions of the RomuQuad and RomuDuo generators having only 32 bits of state, which is small enough that all of their cycles could be discovered and traversed. 800 generators of each type (1600 total) were produced using random multipliers and rotations that were subject to heuristics (listed in Section 4) that ensure the resulting generators were good. For each $k \in [0, 31]$, the 800 actual probabilities were computed from the cycle lengths, and their maximum values are plotted as the violet lines in Figure 3.

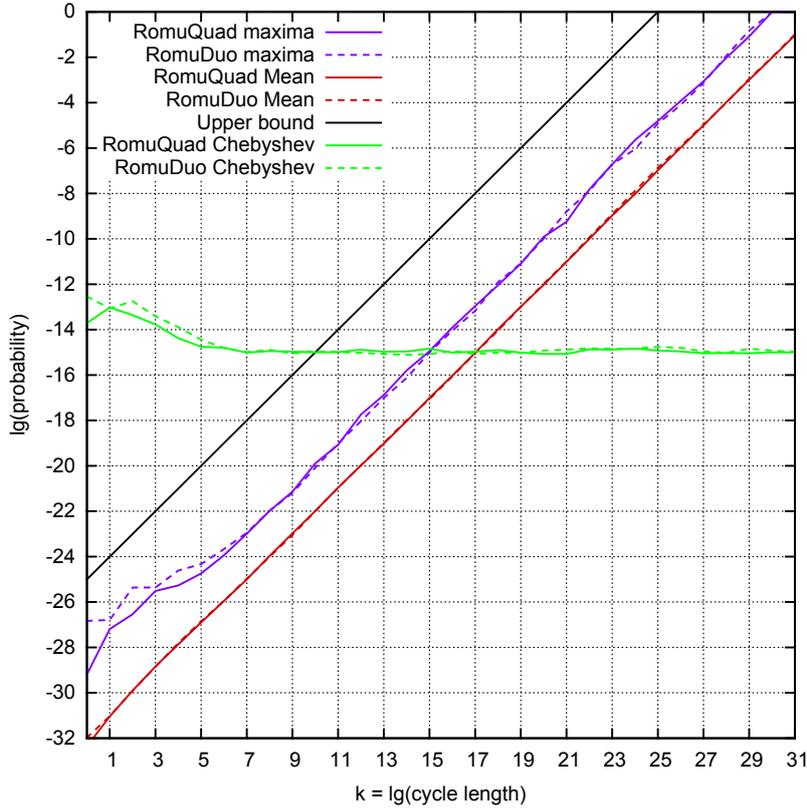


Figure 3: Probabilities vs. cycle-length

Note that this graph uses base-2 logarithms for both axes. The red lines are the mean probabilities, and are present to show that these generators follow the expected means given by (2) very closely. The black line represents (3), denoted in the graph’s legend as “Upper bound”. The two violet maxima lines are nearly linear except at their left ends where their overall slope is lower, especially for RomuDuo. But this nonlinear region corresponds to

$k \in [0, 5]$, which are cycle-lengths of 32 and below. Virtually all jobs consume more than 32 numbers, and the probabilities in this region are at their lowest, so this nonlinear anomaly in the tiniest cycles may be disregarded as harmless.

The green lines in Figure 3 are the bounds resulting from Chebyshev’s inequality:

$$\mathbb{P}(|X - \mu| \geq \delta) \leq \sigma^2/\delta^2. \quad (4)$$

This inequality is independent of distribution, and represents an upper bound on the probability that (3) will be breached. This is, it is the worst-case probability that the probability of seeing a too-short cycle exceeds (3). In (4), σ is the standard deviation in the observed values of $\mathbb{P}(|\text{cycle}| \leq 2^k)$ for each k , and $\delta = 2^{k-32+7} - 2^{k-32}$, which is the difference of (3) and (2) for a state-size of 32 bits (i.e., $s = 32$). In Figure 3, the value of the two Chebyshev lines is nearly constant at approximately 2^{-15} (1 in 30000) except in the nonlinear zone at the left, where they rise to around 2^{-12} (1 in 4000).

It should be mentioned that, outside the anomalous zone on the left, the upper bound given by (3) is around 180 standard deviations away from the mean sample probabilities, strongly suggesting that the true probabilities of breaching (3) are far lower than the worst-case probabilities given by Chebyshev’s inequality.

Based on the probabilities of 2^{-12} to 2^{-15} given by Chebyshev’s inequality, and the fact that the true probabilities are almost certainly far lower, we conclude that (3) is a reliable upper bound on the probability of encountering a too-short cycle.

| Row | State size (bits) | Min. cycle len | U.B. of prob. of shorter cycle |
|-----|-------------------|----------------|--------------------------------|
| 1 | 256 | 2^{55} | 2^{-194} |
| 2 | 192 | 2^{55} | 2^{-130} |
| 3 | 128 | 2^{55} | 2^{-66} |
| 4 | 96 | 2^{55} | 2^{-34} |
| 5 | 64 | 2^{55} | 2^{-2} |

Table 2: Upper bounds of probabilities of shorter cycles.

Table 2 provides upper bounds on shorter-cycle probabilities for several state sizes based on (3). This table assumes a minimum cycle length of 2^{55} , which exceeds the longest feasible stream as it would take over a year to consume 2^{55} values at the optimistic rate of one per nanosecond. The bounds of probabilities of encountering cycles shorter than 2^{55} are inconceivably low for state sizes of 256, 192, and 128 bits (rows 1-3). With 96 bits of state (row 4), the bounded probability for that huge job is one in 16 billion. But with only 64 bits (row 5), the bound is 0.25, showing us that 64 bits of state are not adequate for general purpose work when using an invertible nonlinear generator such as Romu.

Based on (3), we conclude that when the number of bits s in a generator’s state is at least 128, the probability of encountering a cycle shorter than the job or stream is infinitesimal.

Fog (2001) suggested adding a *self-test* to his RANROT generators which would detect the completion of a cycle by comparing the state with the seed, and take an unspecified remedial action. But with a state size of at least 128 bits, the probability of seeing a too-short cycle is so low that such a test would not be beneficial.

3.2. Overlaps among sequences

We must ask ourselves, “What is the purpose of a seed from the user’s perspective?” The user selects a different seed in order to obtain a different pseudo-random sequence. Also, in Section 2.1.1, it was recommended that multiple streams be implemented with cycle-splitting, by selecting streams with seeds. Thus, the user regards a seed as a sequence-selector, and will be angered if the sequences or streams resulting from two seeds overlap. Hence, we need to know the probability of overlap.

Generators with known periods

Let us define l as the length of each sequence (or stream) used by an application, n as the number of such sequences needed, and p as the (known) period of a generator or that of a specific cycle in a Romu generator.

Let us regard the n sequences as having been generated sequentially (i.e., not in parallel). The probability that sequence i overlaps a prior sequence is $2il/p$, $i \in [0, n - 1]$. Why the factor of 2? If we think of starting a new sequence B at a random place in the period, then if its random starting position is within l of the beginning of some prior sequence A , then B will overlap A . Hence, each sequence is preceded by a vulnerable zone of length l . A sequence itself is also vulnerable to overlap, so the combined vulnerable block has length $2l$.

To compute the probability that any sequence overlaps another, we must analyze the two cases of n being even and n being odd.

$$P(\text{overlap} \mid n \text{ is even}) = 1 - \prod_{i=0}^{n-1} \left(1 - \frac{2il}{p}\right) \quad (5)$$

$$= 1 - \prod_{i=0}^{(n/2)-1} \left(1 - \frac{2il}{p}\right) \left(1 - \frac{2(n-1-i)l}{p}\right) \quad (6)$$

$$= 1 - \prod_{i=0}^{(n/2)-1} \left(1 - \frac{2l(n-1)}{p} + \underbrace{4(n-1-i)i \left(\frac{l}{p}\right)^2}_{\text{insignificant}}\right) \quad (7)$$

$$\approx 1 - \left(1 - \frac{2l(n-1)}{p}\right)^{n/2} \quad (8)$$

$$\approx \frac{(n-1)nl}{p} \quad (9)$$

$$P(\text{overlap} \mid n \text{ is odd}) = 1 - \prod_{i=1}^{n-1} \left(1 - \frac{2il}{p}\right) \quad (10)$$

$$= 1 - \prod_{i=1}^{(n-1)/2} \left(1 - \frac{2il}{p}\right) \left(1 - \frac{2(n-i)l}{p}\right) \quad (11)$$

$$= 1 - \prod_{i=1}^{(n-1)/2} \left(1 - \frac{2nl}{p} + \underbrace{4(n-i)i \left(\frac{l}{p}\right)^2}_{\text{insignificant}}\right) \quad (12)$$

$$\approx 1 - \left(1 - \frac{2nl}{p}\right)^{(n-1)/2} \quad (13)$$

$$\approx \frac{(n-1)nl}{p} \quad (14)$$

Equations (6) and (11) group the terms in the products into pairs. Equations (7) and (12) reveal that these pairs are approximately equal because the expression marked “insignificant” is $O((nl/p)^2)$, making it much smaller than the other terms in the sum when $nl/p \ll 1$. Equations (8) and (13) were simplified using the binomial approximation:

$$(1+x)^n \approx 1+nx, \quad |x| \ll 1.$$

Like the first approximation above, the binomial approximation adds an insignificant error of $O((nl/p)^2)$. Both of these approximations increase the final result, making it an upper bound. And both make small changes in the result when $nl/p \ll 1$, so (14) is fairly accurate.

Note that (9) and (14) are identical.

An approximation of (14) can be obtained from the work of Knuth (1981, p. 69). He derived (15), which is an upper bound on the average total number of collisions c of n balls thrown into m urns. Let us divide the period (of length p) of a generator into equal-sized adjacent blocks of length $2l$, where the factor of 2 accounts for vulnerable zones. Substituting the number of such blocks $m = p/2l$ into (15) yields (16), which is close to (14). When c is small, it is an upper bound on the probability of any overlaps occurring.

$$c < n^2/2m \quad (15)$$

$$P(\text{overlap}) < ln^2/p. \quad (16)$$

Fog (2015, p. 312), L’Ecuyer, Munger, Oreshkin, and Simard (2016, p. 15), and Vigna (2018, “Remarks” section) have presented (16) in their work. Though (16) is more elegant, (14) is preferable because it is significantly more accurate when n is small due to running a job on a single multi-core computer.

Generators with probabilistic periods

The periods of the multiple cycles in a Romu generator are not known, but the upper bound of their probability distribution is given in (3). To find an upper bound on the probability of overlap of its sequences, we must convolve the distributions given in (3) and (14). First, we list the definitions they need:

- $s \equiv$ number of bits in the generator’s state,
- $l \equiv$ lg(number of values in one sequence),
- $n \equiv$ number of such sequences,
- $r \equiv$ lg(period of the cycle containing a randomly chosen x),
- $k \equiv$ lg(period of a cycle).

In this paper, $\lg(x) \equiv \log_2(x)$. Also note that the definition of l differs from that used in the previous section.

For a Romu generator, consider the cycle containing a randomly selected value x . The upper bound on the probability that this cycle's period r will be in the interval $[2^{k-1}, 2^k]$ is given by

$$P_{\text{ub}}(r \in [k-1, k]) = 2^{k-s+7} - 2^{k-s-1+7} = 2^{k-s+6}. \quad (17)$$

The convolution can then be expressed as follows:

$$P_{\text{ub}}(\text{overlap}) = \sum_{k=l}^s (1/a) \cdot P_{\text{ub}}(r \in [k-1, k]) \cdot P(\text{overlap} \mid \text{period } p = 2^{k-0.5}). \quad (18)$$

The expression $2^{k-0.5}$ at the end of (18) is the geometric mean of the interval of summation, $(2^{k-1}, 2^k]$, and was employed to improve accuracy. The summation starts at $k = l$ because we are presuming that $r \geq l$, meaning that the cycle is not shorter than the sequence to be output. In addition, if we had used the expected probabilities in (17) and (18), then starting the summation at $k = l$ would cause the sum of those probabilities to be less than 1.0, necessitating a correction factor of $1/a$, where a is the sum of expected probabilities:

$$\begin{aligned} a &= \sum_{k=l}^s E(P(r \in [k-1, k])) = \sum_{k=l}^s (2^{k-s} - 2^{k-s-1}) = \sum_{k=l}^s 2^{k-s-1} \\ &= 2^{-(s+1)} \sum_{k=l}^s 2^k = 2^{-(s+1)} (2^{s+1} - 2^l) \\ &\approx 2^{-(s+1)} 2^{s+1}, \quad \text{when } 2^l \ll 2^{s+1} \\ &= 1 \end{aligned}$$

The same correction factor should be applied to the upper bounds used in (18). However, for practical Romu generators, $s \geq 96$, and for an impossibly large job, $l \leq 60$. Therefore, $2^l \ll 2^{s+1}$, so the correction factor of $1/a$ in (18) is approximately 1.0 and thus may be discarded. The resulting convolution is:

$$P_{\text{ub}}(\text{overlap}) = \sum_{k=l}^s 2^{k-s+6} \cdot 2^{l-k+0.5} (n-1)n \quad (19)$$

$$= 2^{6.5+l-s} (s-l+1)(n-1)n \quad (20)$$

$$\approx 2^{6.5+l-s} (s-l)n^2. \quad (21)$$

The loss of accuracy caused by changing $(n-1)n$ to n^2 in (21) is insignificant in comparison to the factor of 2^7 that was brought into (21) from (3) via the convolution.

Table 3 provides the probabilities of overlap for various generators and jobs. The *Known period* column contains probabilities for generators with known periods of about 2^s , and is based on (14). The probabilities in the *Romu* column are based on (21), and are upper bounds. For Romu generators:

- 256 bits of state (row 1): This large state size far exceeds what could ever be needed. Absurdly large values of l and n result in an inconceivably low probability of overlap.

| Row | s | l | n | Probability of overlap | |
|-----|-----|-----|----------|------------------------|--------------------|
| | | | | Known period | Romu (upper bound) |
| 1 | 256 | 64 | 2^{30} | 2^{-132} | $2^{-117.9}$ |
| 2 | 192 | 55 | 2^{17} | 2^{-103} | $2^{-89.4}$ |
| 3 | 128 | 55 | 2^{15} | 2^{-50} | $2^{-30.3}$ |
| 4 | 96 | 45 | 2^{11} | 2^{-29} | $2^{-16.8}$ |
| 5 | 64 | 42 | 2^5 | 2^{-12} | $2^{-1.0}$ |

Table 3: Probabilities of sequence-overlap in generators.

- 192 bits of state (row 2): The massive job in this row consists of 131 thousand streams of 2^{55} values each. Recall that consuming 2^{55} values at the optimistic rate of one per nanosecond would take over a year. The resulting upper bound of probability of overlap is lower than that of selecting a specific snowflake among all snowflakes that have fallen in human history. The human mind cannot conceive of such a low probability. Thus, in reality, streams will never overlap.
- 128 bits of state (row 3): For a massive job consisting of 33 thousand streams of 2^{55} values each, the upper bound on probability of overlap is one in a billion.
- 96 bits of state (row 4): 96 state bits is barely adequate to supply a moderately large job with an upper bound on overlap-probability about one in 100,000.
- 64 bits of state (row 5): Such a small state size is not sufficient for large applications, as a Romu generator cannot provide a large number of large streams with a low probability of overlap.

With at least 128 bits of state, we can safely ignore the possibility of overlap.

Some people might insist that a generator have exactly one period (of known length), which eliminates nonlinear generators from consideration. But with 128 or more bits of state, such a person will be unwisely losing the substantial benefit of improved performance in order to avoid a will-never-happen risk of overlap (or a too-short cycle). For 192 bits of state, randomly selecting one specific snowflake in human history will never happen. Losing benefits to avoid such a nil risk is unwise.

The right two columns show that the convolution has raised the probabilities of overlap for Romu generators (in relation to generators with known periods) by factors of $2^{11.0}$ to $2^{14.1}$. For state sizes of at least 128 bits, the probabilities are so low that this boost is harmless. But for state sizes of 64 or fewer bits, probabilities of overlap are high enough that Romu generators should be avoided except for specialized applications.

3.3. Equidistribution

Some generators have the property of *d-dimensional equidistribution*, which means that when the generator is run over its entire period, all *d*-dimensional hypercubes will contain the same number of points. A Romu generator is *not* equidistributed because every cycle contains only a subset of the 2^s possible states, causing hypercubes to contain an uneven number of points,

as would occur with truly random numbers. In fact, equidistribution can be regarded as a *defect* in a generator because truly random numbers are not equidistributed.

L’Ecuyer (1996, p. 203) claims the following benefits of equidistribution: “Such generators have strong theoretical support and lend themselves to very fast software implementations.” Unfortunately, the strong theoretical support applies only to linear generators, and not to nonlinear generators such as Romu. And the claim of “very fast” is suspect because our testing shows that the nonequidistributed Romu generators are faster than equidistributed generators with the same degree of randomness. Brent (2010) stated his conclusion clearly:

When comparing modern long-period pseudo-random number generators, (d, w) -equidistribution is irrelevant, because it is neither necessary nor sufficient for a good generator.

4. Creating and Using Romu Generators

4.1. Determination of constants

The Romu generators presented in this paper contain some constant scalars, specifically, rotation-counts and multipliers. These could not be determined mathematically because the nonlinear arithmetic in Romu generators is mathematically intractable, as mentioned above. And they could not be determined empirically with the actual generators because doing so would take centuries.

The solution was to employ the mini-generators described above. However, in some cases, mini-generators with 40 and 48 bits of state were created when 32 bits was found to be too small and 64 bits was too large. Using **PractRand**, each trial rotation-count was scored based on the amount of data required to produce a failure. Generators containing two rotations required evaluating many trial rotation-pairs. Preference was given to high-scoring rotation-count(s) whose neighbors (with counts differing by 1 or 2) also had high scores, indicating that the prospective rotation-count(s) was not on a narrow peak, but was at the top of a gentle hill, providing assurance that the scaled-up count would also be at or near the peak. Each prospective rotation-count(s) was also tested with multiple multipliers to ascertain its sensitivity to multipliers. The winning rotation-count(s) was then scaled up for the full-size generator.

The above procedure cannot determine multipliers because they cannot be scaled up. Instead, multipliers were created randomly and then hand-modified using the following heuristics:

1. The 4 least significant bits (LSBs) were set to 1011 to (1) make the multiplier odd so that multiplications would be invertible, and (2) mix the LSBs of multiplicands well.
2. The multiplier’s Hamming weight (i.e., its number of 1-bits) was changed to be near $w/2$, where w is the word-size of the generator’s arithmetic (64 or 32).
3. Any repeated patterns in the multiplier’s bits were broken up.
4. Long runs of 0- or 1-bits were avoided.

Several prospective 64-bit multipliers were then tested using a RomuMono generator that was modified to output 56 bits instead of the usual 32, making it poor enough that **PractRand** would fail it in a few minutes. All of the multipliers performed similarly, even with different rotations. The ratio of the best number of bytes processed to the worst was about 8.0, which is insignificant at large magnitudes. The best multiplier was selected for use in all of the 64-bit Romu generators presented in this paper.

4.2. Seeding Romu generators

Every Romu generator presented in this paper must be seeded such that its initial state is not all zeros. Any nonzero seed will suffice, including one that is almost all zeros. A seed with all zeros except for one 1-bit is acceptable. Due to its multiplication and rotation(s), a Romu generator quickly escapes from an initial mostly-zeros state. However, to avoid having the first few output values be highly correlated to the initial state, one should call the generator a few times (discarding its output) in order to mix the bits in its state.

Equations (2) and (3), representing the probability of encountering a too-short cycle, rely on random selection of cycles. Based on empirical results and the mixture of linear and nonlinear arithmetic in Romu generators, it can be argued that nonrandom seeds, such as the output of a counter, will be in random locations in random cycles. But in the absence of proof of this hypothesis, it is prudent to seed Romu generators with pseudo-random values.

As will be seen in the discussion below, the SplitMix64 generator (Vigna 2015) is well suited for seeding generators, and Vigna recommends it for seeding the Blackman-Vigna generators. Vigna has kindly placed this generator into the public domain. In addition, we have created a SplitMix32 generator to seed the Romu generators that use 32-bit arithmetic. The source-code of SplitMix64 and SplitMix32 are provided in Appendix A and B. For both, the state variable `sState` must first be set to the seed for the job or stream. That seed may be any value, including zero. Seeds for streams may be consecutive integers $[0, 1, 2, 3, \dots]$. A Romu generator can be seeded and mixed as illustrated by the following example for RomuQuad:

```
void romuQuad_init (uint64_t seed) {
    sState = seed;           // sState is the state-variable for splitMix64.
    wState = splitMix64(); // w/x/y/zState are the state-vars for RomuQuad.
    xState = splitMix64();
    yState = splitMix64();
    zState = splitMix64();
    int i;
    for (i=0; i<10; i++) romuQuad_random(); // Ensure state is mixed well.
}
```

This code is using SplitMix to produce a micro-stream of four values, one micro-stream per seed. In order to assess the independence of Romu's streams, we must gauge the randomness of the concatenation of all micro-streams. We denote each micro-stream as S_{lt} , where l is the number of values in the micro-stream and t is the user's seed, assumed to be produced by a counter. Let $S_l = \{S_{l0}, S_{l1}, S_{l2}, \dots\}$, representing the concatenation of the micro-streams. S_l is all the values that will be output by SplitMix when seeding all Romu streams. S_l was generated for $l \in [1, 2, 3, 4, \infty]$, and passed into **PractRand** for both SplitMix64 and

SplitMix32. The value $l = \infty$ applies when a single sequence from SplitMix is used to seed all Romu streams, instead of employing micro-streams.

The capacities reported by **PractRand** ranged from 2^{26} to at least 2^{35} values, which is more than sufficient for seeding all streams in the largest jobs. Thus, these SplitMix generators will seed all streams in Romu generators with pseudo-random numbers having undetectable correlations, assuring us that (1) the results of equations (2) and (3) are valid because cycles will be selected at random, and (2) Romu streams will not be correlated because their cycles and locations within cycles will be selected randomly.

4.3. Rotation operation

Because Romu generators perform rotations, we define the following rotate-left macro in C/C++:

```
#define ROTL(d,lrot) ((d<<(lrot)) | (d>>(8*sizeof(d)-(lrot))))
```

Modern compilers such as gcc and Visual Studio™ recognize the expression in this macro as a left-rotation, and emit one instruction for it. Also, the `sizeof` operator in this macro enables it to properly rotate operands of any size. All C-code examples provided in this paper use this ROTL macro. The following is the equivalent template function in C++:

```
template <class uDataT>
inline uDataT rotl (uDataT d, unsigned lrot) {
    return (d<<lrot) | (d>>(8*sizeof(d)-lrot));
}
```

Instead of the template function, we prefer the ROTL macro as it will serve both C and C++ equally well.

5. Romu Generators using 64-bit Arithmetic

5.1. Introduction to 64-bit generators

The Romu generators employing 64-bit arithmetic are presented below. Their traits are summarized in Table 4, along with the traits of some of their competitors.

In Table 4:

Bits of state is 64 multiplied by the number of state variables in the generator (2, 3, or 4).

Output latency is an important concept. It is the delay (in clock cycles) the generator imposes on the application when the generator is inlined in a superscalar processor. In a well designed generator, the output value becomes available before the generator's processing is done, allowing the application to execute in parallel with the generator. The 64-bit Romu generators presented in this paper have no output latency, and thus cause no delay in the application. From the application's point of view, these generators consume no time.

As an aside, any generator that computes its output-value (and thus has a nonzero output latency) can be modified to have no output latency by computing and saving the next output-value in a variable to be returned in the next call or invocation. But this trick comes with a

| Generator | Bits of state | Output latency | Arith. ops | Proc. cycles | Register pressure | Tested bytes | Capacity (bytes) |
|---------------------|---------------|----------------|------------|--------------|-------------------|--------------|----------------------|
| RomuQuad | 256 | 0 | 6 | 3 | 8 | 2^{48} | $\geq 2^{90}$ (est.) |
| RomuTrio | 192 | 0 | 5 | 3 | 6 | 2^{48} | 2^{75} (est.) |
| RomuDuo | 128 | 0 | 5 | 3 | 5 | 2^{48} | 2^{61} (est.) |
| RomuDuoJr | 128 | 0 | 3 | 3 | 4 | 2^{48} | 2^{51} (est.) |
| xoshiro256++ | 256 | 3 | 10 | 3 | 7-8 | 2^{45} | 2^{60} (est.) |
| xoroshiro128++ | 128 | 3 | 9 | 3 | 5 | 2^{45} | 2^{51} (est.) |
| pcg-xsh-rs (128→64) | 128 | 7 | 17 | 9 | 10 | | |

Table 4: Summary of 64-bit generators.

cost: It boosts register pressure. The additional memory-accesses due to a spill will reduce or erase the time-savings.

Arith. ops is the number of arithmetic operations performed by the generator, excluding data-movement. Note that the Romu generators have around half or fewer operations than their (linear) competitors, revealing the power of a rotation in conjunction with a multiplication. The pcg-xsh-rs generator has a large number of operations because each of its 128-bit operations creates multiple 64-bit operations; these were counted in the compiler’s assembler-language output. *Arith. ops* is also a rough indication of how fast the generator will run on a nonsuperscalar processor.

Proc. cycles is the number of clock cycles of processing time needed by the generator when inlined, assuming that its variables are held in registers. When this value is greater than *Output latency*, ILP is causing the application and generator to execute concurrently after the random number has been output. When the generator is not inlined, it is being called as a function, so the time spent reading/writing variables and the call/return overhead must be added to this processing time.

Register pressure is the number of registers devoted to the generator when inlined in a loop. It is important to count registers in a loop and not in a function, as a function usually uses fewer registers because writes to memory also free registers. These numbers were determined by compiling with g++ version 4.8.4 with the `-S -O3` options, and examining the resultant assembler-language file.

Tested bytes is the number of bytes (not number of values) supplied to **PractRand** in a long test. Each generator passed with no signs of stress.

Capacity is the maximum number of bytes each generator can supply before statistical anomalies become apparent in **PractRand**. When this value exceeds *Tested bytes*, it is an estimate (marked “(est.)” in the table) based on testing a scaled-down version of the generator as described above. A value of the form “ $\geq 2^n$ ” indicates that the scaled-down generator passed its test without stress or failure. These estimates are conservative, so the true maxima should be higher.

David Blackman, who is one of the inventors of the good Blackman-Vigna generators, kindly provided the capacity estimates for xoroshiro128++ and xoshiro256++ given in Table 4.

5.2. RomuQuad generator

The state of RomuQuad comprises four 64-bit variables (256 bits). Its processing time is 3 cycles, and it has no output latency, allowing an application to execute as soon as possible. This is the best 64-bit Romu generator presented in this paper, but it also has the highest register pressure, which is 8. RomuTrio is recommended instead due to its lower register pressure of 6. But RomuQuad can be employed for massive jobs run by users who choose to be extremely cautious about the probability of overlap.

```
uint64_t wState, xState, yState, zState; // set to nonzero seed

uint64_t romuQuad_random () {
    uint64_t wp = wState, xp = xState, yp = yState, zp = zState;
    wState = 15241094284759029579u * zp; // a-mult
    xState = zp + ROTL(wp,52);          // b-rotl, c-add
    yState = yp - xp;                    // d-sub
    zState = yp + wp;                    // e-add
    zState = ROTL(zState,19);           // f-rotl
    return xp;
}
```

The ILP table (below) for RomuQuad reveals a subtle reason this generator will slightly slow down the application. All four slots are occupied in cycle 1, causing the generator and the application to compete for one of those slots, delaying one of them by a cycle. On average, the application will be delayed by half a clock cycle, which is insignificant.

| Cycle | Slot 1 | Slot 2 | Slot 3 | Slot 4 | Multiplier |
|-------|--------|--------|--------|--------|------------|
| 1 | b-rotl | d-sub | e-add | a-mult | * |
| 2 | c-add | f-rotl | | | * |
| 3 | | | | | * |

One might regard RomuQuad as too large, but its half-size version (RomuQuad32) presented later requires all four 32-bit state variables in order to reach 128 bits of state.

| Output variable | Input var. 1 | Input var. 2 | Transform |
|-----------------|--------------|--------------|-------------|
| wState | zp | | multiply |
| xState | wp | zp | rotate, add |
| yState | xp | yp | subtract |
| zState | yp | wp | add, rotate |

Table 5: Dataflows in RomuQuad.

The dataflows from input variables to output variables are instructive. Column *Input var. 1* of Table 5 shows that each output variable is computed from the prior input variable. For example, `yState` is computed from `xp`. Drawing from the prior input variable ensures that data flows circularly among the four state-variables. If those were the only dataflows, then after every four calls, no variable would have affected any other, so in essence there would be

four independent 64-bit generators operating in tandem, creating poor randomness. Therefore, every output variable (except `wState`) is also computed from a second input variable (in column *Input var. 2*), causing the variables to mix with each other. Except for `wState`, every output variable is computed from two input variables. Furthermore, every input variable (except `yp`) supplies two output variables. This two-out/two-in design mixes the state variables well. A rotation contributes to every other output variable (`xState` and `zState`), creating an alternation between rotations and additions, improving the mixing of bits.

An alternate design is to have every output variable use the prior two input variables, except for the product in `wState`. For example, `yState` would be computed from `wp` and `xp`. Tests on a 32-bit generator consisting of four 8-bit state variables indicated that this approach matches the randomness of RomuQuad. This design was not investigated further as it improves neither performance nor randomness, and it was feared that the repeating prior-two-into-one pattern might produce subtle patterns in output values.

5.3. RomuTrio generator

RomuTrio has 192 (3×64) bits of state, and as seen in Table 3, such a generator can create many huge sequences with essentially no chance of overlap (one snowflake in human history). Furthermore, RomuTrio easily passes all statistical tests. RomuQuad is fast, and RomuTrio is faster due to its lower register pressure and sparser ILP table. For these reasons, this generator is recommended most highly.

```
uint64_t xState, yState, zState; // set to nonzero seed

uint64_t romuTrio_random () {
    uint64_t xp = xState, yp = yState, zp = zState;
    xState = 15241094284759029579u * zp;
    yState = yp - xp;  yState = ROTL(yState,12);
    zState = zp - yp;  zState = ROTL(zState,44);
    return xp;
}
```

The ILP Table (below) for RomuTrio reveals that it makes excellent use of ILP. Note that slot 4 is empty in all three cycles, allowing the application to execute in parallel with this generator, maximizing its speed.

| Cycle | Slot 1 | Slot 2 | Slot 3 | Slot 4 | Multiplier |
|-------|------------|------------|----------|--------|------------|
| 1 | subtract x | subtract y | multiply | | * |
| 2 | rotate y | rotate z | | | * |
| 3 | | | | | * |

As seen in Table 4, the processing time of `xoshiro256++` (by Blackman and Vigna) matches that of RomuQuad and RomuTrio, but its output latency is 3 cycles, a noticeable speed loss in time-critical applications. The ILP table (below) for `xoshiro256++` reveals the genius behind this good generator, as all four slots have been thoughtfully filled. But this full ILP table could cause some speed loss if the application inserts an instruction or two in these slots,

pushing the generator’s instruction(s) into additional clock cycle(s). Overall, xoshiro256++ is slower than RomuTrio due to its 3-cycle output latency, fragile ILP-table, and higher register pressure.

| Cycle | Slot 1 | Slot 2 | Slot 3 | Slot 4 | Multiplier |
|-------|--------|--------|--------|--------|------------|
| 1 | add | shift | xor | xor | |
| 2 | rotate | xor | xor | xor | |
| 3 | add | rotate | | | |

The xoshiro256++ generator has the advantage of offering a jump-ahead function, which is supplied with the generator’s source-code. No Romu generator can have this feature due to its intractable nonlinear arithmetic. But does this feature help an application? Jump-ahead is used to create nonoverlapping streams. The 192 bits of state in RomuTrio ensure that its streams will not overlap, eliminating the need for jump-ahead.

5.4. RomuDuo generator

RomuDuo will be faster than RomuTrio when its lower register pressure causes fewer memory-spills. Its state size of 128 bits is large enough to support many streams with essentially no chance of overlap. The estimated capacity of RomuDuo is 2^{58} values (2^{61} bytes), which is significantly lower than that of RomuTrio or RomuQuad. But this lower capacity is still larger than any job, and it would take 9 years to consume it at the rate of one per nanosecond. This estimate of capacity should be too low because it is conservative. But because it is an estimate and not a measurement, a prudent worker will maintain some margin between job/stream-size and capacity for large jobs requiring a high degree of randomness.

```
uint64_t xState, yState; // set to nonzero seed

uint64_t romuDuo_random () {
    uint64_t xp = xState;
    xState = 15241094284759029579u * yState;
    yState = ROTL(yState,36) + ROTL(yState,15) - xp;
    return xp;
}
```

The ILP table (below) for RomuDuo is sparse, leaving plenty of slots for the application to use, which can only improve performance.

| Cycle | Slot 1 | Slot 2 | Slot 3 | Slot 4 | Multiplier |
|-------|----------|--------|----------|--------|------------|
| 1 | rotate | rotate | multiply | | * |
| 2 | add | | | | * |
| 3 | subtract | | | | * |

5.5. RomuDuoJr generator

RomuDuoJr is a simplification of RomuDuo, removing a rotation and addition. Hence, we appended Jr to its name, which stands for “Junior”. The reduced number of operations and

register pressure make it the fastest generator presented in this paper using 64-bit arithmetic. RomuDuoJr is suitable for most applications, and it should be preferred when speed is paramount. However, a large job can exceed its reduced capacity, so one must be careful.

```
uint64_t xState, yState; // set to nonzero seed

uint64_t romuDuoJr_random () {
    uint64_t xp = xState;
    xState = 15241094284759029579u * yState;
    yState = yState - xp; yState = ROTL(yState,27);
    return xp;
}
```

The ILP table (below) for RomuDuoJr is even sparser than that of RomuDuo. In fact, as pointed out in Section 2.4, this generator consists of only three arithmetic instructions.

| Cycle | Slot 1 | Slot 2 | Slot 3 | Slot 4 | Multiplier |
|-------|----------|----------|--------|--------|------------|
| 1 | subtract | multiply | | | * |
| 2 | rotate | | | | * |
| 3 | | | | | * |

We can compare O’Neill’s PCG (O’Neill 2014) to RomuDuo/RomuDuoJr. The PCG is an LCG with a randomized scrambler added as a post-processing step. The `pcg-xsh-rs` variant uses 128-bit-wide arithmetic, and returns 64 bits of output. Its state-size is 128 bits, matching RomuDuo/RomuDuoJr. But its 128-bit-wide arithmetic makes it much slower than all Romu generators, as Table 4 shows. Furthermore, its large register pressure of 10 will cause memory-spills, slowing the application even more. Like the Blackman-Vigna generators, the PCG offers a jump-ahead function. And as mentioned above, jump-ahead is unnecessary for generators with 128 or more bits of state due to the extremely low probabilities of sequence-overlap.

6. Romu Generators using 32-bit Arithmetic

32-bit generators are needed when running on specialized hardware lacking 64-bit arithmetic, such as the Intel Nios II ® soft core and DSP chips such as the Texas Instruments C66x family. Their small word-size of 32-bits makes it difficult for their generators to have a large number of state-bits without resorting to an internal array, which would reduce performance due to reading/writing memory and managing indices or pointers.

RomuQuad32 is suitable for large jobs, and RomuTrio32 and RomuMono32 are suitable for medium-to-large and small jobs, respectively. Among these choices, RomuTrio32 is the fastest, so we recommend preferring it for jobs that will not exceed its capacity.

6.1. RomuQuad32 generator

This 128-bit generator is suitable for large jobs, is very fast (with no output latency), and has excellent statistical quality. It is the only Romu generator using 32-bit arithmetic this

| Generator | Bits of state | Output latency | Arith. ops | Proc. cycles | Register pressure | Tested bytes | Capacity (bytes) |
|--------------------|---------------|----------------|------------|--------------|-------------------|--------------|----------------------|
| RomuQuad32 | 128 | 0 | 6 | 3 | 7 | 2^{45} | $\geq 2^{62}$ (est.) |
| RomuTrio32 | 96 | 0 | 5 | 3 | 5 | 2^{45} | 2^{53} (est.) |
| RomuMono32 (32→16) | 32 | 1 | 2 | 4 | 2 | 2^{27} | 2^{27} |
| xoshiro128** | 128 | 5 | 10 | 5 | 6-7 | 2^{45} | |
| pcg-xsh-rr (32→16) | 32 | 4 | 7 | 4 | 3 | 2^{27} | 2^{27} |
| LCG (32→upper 16) | 32 | 0 | 2 | 4 | 2 | 2^{17} | 2^{17} |

Table 6: Summary of 32-bit generators.

paper recommends for general purpose use. However, as mentioned above, RomuTrio32 is faster due to its lower register pressure, so we recommend employing it when feasible.

Earlier, it was stated that a register pressure of 7 or more is regarded as excessive. But a state size of 128 bits requires that four 32-bit variables be kept in registers, and after adding the necessary temporary registers, high register pressure became unavoidable in RomuQuad32. The register pressure of RomuQuad32 is lower than that of RomuQuad because a quirk in the Intel/AMD instruction set forces a register to be dedicated to holding a 64-bit multiplier, an aberration not present in its 32-bit instructions.

```
uint32_t wState, xState, yState, zState; // set to nonzero seed

uint32_t romuQuad32_random () {
    uint32_t wp = wState, xp = xState, yp = yState, zp = zState;
    wState = 3323815723u * zp; // a-mult
    xState = zp + ROTL(wp,26); // b-rotl, c-add
    yState = yp - xp; // d-sub
    zState = yp + wp; // e-add
    zState = ROTL(zState,9); // f-rotl
    return xp;
}
```

The ILP table (below) is identical to that of RomuQuad because this is a half-size version of RomuQuad.

| Cycle | Slot 1 | Slot 2 | Slot 3 | Slot 4 | Multiplier |
|-------|--------|--------|--------|--------|------------|
| 1 | b-rotl | d-sub | e-add | a-mult | * |
| 2 | c-add | f-rotl | | | * |
| 3 | | | | | * |

Blackman and Vigna’s xoshiro128** generator competes against RomuQuad32. Both have 128 bits of state using 32-bit arithmetic with about the same register pressure, producing top quality values. But Table 6 shows that RomuQuad32 is 5 cycles faster. Though xoshiro128** is slower, it has the advantage of performing *no* multiplications (despite the stars in its name) because it multiplies by the constants 5 and 9 which the compiler translates into `lea` instructions or shifts-and-adds. If multiplication is slow or absent in the processor,

xoshiro128** is probably the best choice. Otherwise, the higher speed of RomuQuad32 makes it preferable.

6.2. RomuTrio32 generator

With 96 bits of state, this generator is not suitable for the largest jobs, as Table 3 reveals. But it can supply the needs of most applications. Because its register pressure is lower than RomuQuad32, there will be fewer spills, causing an application to run faster.

```
uint32_t xState, yState, zState; // set to nonzero seed

uint32_t romuTrio32_random () {
    uint32_t xp = xState, yp = yState, zp = zState;
    xState = 3323815723u * zp;
    yState = yp - xp;  yState = ROTL(yState,6);
    zState = zp - yp;  zState = ROTL(zState,22);
    return xp;
}
```

Because this generator is a half-size version of RomuTrio, its ILP table (below) is identical. The empty slot 4 allows the application to run at maximum possible speed because it will never wait on the generator.

| Cycle | Slot 1 | Slot 2 | Slot 3 | Slot 4 | Multiplier |
|-------|------------|------------|----------|--------|------------|
| 1 | subtract x | subtract y | multiply | | * |
| 2 | rotate y | rotate z | | | * |
| 3 | | | | | * |

6.3. RomuMono32 generator

We first published this elegant generator in Overton (2011), but named it “CMR”, which stands for “constant, multiply, rotate”, because it multiplies by a constant and rotates the result. RomuMono32 is tiny, consisting of only those two arithmetic operations plus one more to extract the returned result. Because RomuMono32 has only 32 bits of state, the period of its longest cycle can be traversed in a few seconds, allowing one to search for multipliers and rotation-counts that yield cycles with periods very close to 2^{32} . If p is the period resulting from such a multiplier-rotation pair, let $d = 2^{32} - p$, which is the amount the period falls short of the ideal of 2^{32} . For several months, all cores of several computers randomly searched for multiplier-rotation pairs that yield small values of d .

If d is small (i.e., under about 10000), and if one were to sort all numbers in the longest cycle, there would be large blocks of numbers that are consecutive on the number line. These can be used to seed the generator, guaranteeing that every seed from the block will be in the longest cycle. When seeded in this manner, the generator’s period is known and is not probabilistic. The length of the largest such *seed-block*, as we term them, determines the number of bits of seed the generator can accept. The mean length of a seed-block is $2^{32}/d$. For example, if d is 47, seed-blocks will average $2^{26.4}$ values. The largest seed-block will be larger due to random

variation in their sizes, so such a generator can accept more than 26 bits of seed, as seen in function `romuMono32_init` below.

```
uint32_t state;

void romuMono32_init (uint32_t seed) {
    state = (seed & 0x1fffffffu) + 1156979152u; // Accepts 29 seed-bits.
}

uint16_t romuMono32_random () {
    uint16_t result = state >> 16;
    state *= 3611795771u; state = ROTL(state,12);
    return result;
}
```

The d -value of this generator is 47, so its period is $2^{32}-47$. The unusual method of seeding in `romuMono32_init` adds the 29 LSBs of the seed to the beginning of the largest seed-block for the given multiplier and rotation.

Appendix C lists all multiplier-rotation pairs having d -values under 2000 that the long search found. Rotations up to only 16 are included because testing revealed that randomness suddenly drops with larger rotations. However, despite its reduced randomness, the second row in the table with a rotation-count over 16 was included due to its extraordinarily low d -value. The period of `RomuMono32` is about 2^{32} , which is small by today's standards, so it should be used only for small applications. This small period creates the temptation to combine two or three such generators, perhaps by mixing their outputs with an add or xor, yielding periods of 2^{64} or 2^{96} . This experiment was done, and it was found that two `RomuMono32` generators in combination can output 32 bits each (instead of the usual 16), creating 32-bit combined outputs which will pass **BigCrush** and **PractRand** at 2^{45} bytes. When combining these generators, one must ensure that their periods are coprime in order to achieve the longest possible combined period. It is also prudent to choose rotations that differ by at least 2.

But such a combination generator has a problem: There is likely to be detectable correlation between portions of two sequences with lengths approaching 2^{31} . And for lengths exceeding 2^{31} , such correlated portions will surely exist because a base generator must have contributed the same subsequence to both output sequences, causing detectable correlation between them. This is the reason no combination `RomuMono32` generator was presented in this paper.

Though this tiny generator employs only two arithmetic operations to update its state, it matches the statistical quality of the 32-bit `pcg-xsh-rr` variant of the PCG. The capacity of both is 2^{26} 16-bit values. But the PCG requires 7 arithmetic operations to accomplish this feat, versus 3 for `RomuMono32`. The output latency of the PCG is 4 cycles versus 1 for `RomuMono32`, making it much faster. The purpose of saying this is not to denigrate the PCG, as it is a good generator; the purpose is to emphasize that a mixture of linear and nonlinear arithmetic can produce a good generator with significantly fewer arithmetic operations.

Table 6 reveals this advantage in its comparison of `RomuMono32` with an LCG that returns its high-order 16 bits. Both consist of 3 arithmetic instructions. Both execute in the same

amount of time. Both have 32 bits of state. But the LCG can output only 2^{16} values before troubles become apparent, versus 2^{26} values for RomuMono32.

That 2^{26} is large enough to be useful. When a small job running in a 32-bit processor needs a tiny generator, RomuMono32 is a good choice.

7. Conclusion

Members of the Romu family offer the following advantages over other good generators:

- They are faster. Except for RomuMono32, they consume only 3 clock cycles due to good ILP, and when inlined, they cause *no* delay in applications, maximizing their speed.
- They are more random for a given number of arithmetic operations.
- Their approximate capacities are known, so they do not impose unknown risks to jobs.

| Generator | Capacity (bytes) | Bits of state | Arithmetic |
|------------|------------------|---------------|------------|
| RomuQuad | 2^{90} | 256 | 64 |
| ✓RomuTrio | 2^{75} | 192 | 64 |
| RomuDuo | 2^{61} | 128 | 64 |
| RomuDuoJr | 2^{51} | 128 | 64 |
| RomuQuad32 | 2^{62} | 128 | 32 |
| RomuTrio32 | 2^{53} | 96 | 32 |
| RomuMono32 | 2^{27} | 32 | 32 |

Table 7: Summary of Romu generators.

Table 7 summarizes the Romu generators presented in this paper. RomuTrio has a check mark because it is most highly recommended for general purpose use, as it offers both high speed and infinitesimal probabilities of too-short cycles or stream-overlap. The second recommendation is RomuQuad, but it has higher register pressure and its improved probabilities (compared to RomuTrio) will make no difference in practice. RomuDuo and RomuDuoJr might be even faster than RomuTrio due to lower register pressure, and their probabilities of overlap are low, but the largest jobs might approach the capacity of RomuDuo, and will exceed that of RomuDuoJr. For 32-bit processors, RomuQuad32 is intended for general purpose work, RomuTrio32 can supply most jobs, and RomuMono32 is suitable only for small jobs.

It appears that theoreticians avoid nonlinear generators because their periods and structural traits cannot be determined mathematically. However, their periods are probabilistic, and the convolution presented in this paper shows that, with a state size of at least 128 bits, the probability of a too-short cycle or stream-overlap in a large job will be vanishingly low, causing any user to confidently say, “It will never happen.”

Arithmetic that combines the linear operation of multiplication with the nonlinear operation of rotation is inherently more random than employing only linear operations. For example, using only three arithmetic instructions, no 64-bit linear generator can match the randomness of RomuDuoJr. And using only two arithmetic instructions, no linear generator can match the randomness of RomuMono. Because they perform fewer operations, Romu generators are faster than the equivalent linear generators.

We conclude that the invertible nonlinear generators in the Romu family are superior to linear generators, and we hope researchers can improve upon those presented in this paper.

Acknowledgements

We would like to thank Chris Doty-Humphrey for creating the **PractRand** suite of tests for generators. It is the only test suite that provides a systematic quantitative measurement of the quality of a generator. It performed all quality measurements needed by this paper.

References

- Blackman D, Vigna S (2018). “Scrambled Linear Pseudorandom Number Generators.” *CoRR*, **abs/1805.01407**. [1805.01407](https://arxiv.org/abs/1805.01407), URL <http://arxiv.org/abs/1805.01407>.
- Brent RP (2010). “The Myth of Equidistribution for High-Dimensional Simulation.” Accessed: 2019-12-31, URL <https://arxiv.org/pdf/1005.1320.pdf>.
- Brown RG, Eddelbuettel D, Bauer D (2011). “Dieharder: A Random Number Test Suite.” Accessed: 2019-12-26, URL <http://webhome.phy.duke.edu/~rgb/General/dieharder.php>.
- Doty-Humphrey C (2018). “Practically Random: C++ Library of Statistical Tests for RNGs.” Accessed: 2019-12-26, URL <http://pracrand.sourceforge.net>.
- Fog A (2001). “Chaotic Random Number Generators with Random Cycle Lengths.” Accessed: 2020-1-5, URL <https://www.academia.edu/8247581>.
- Fog A (2015). “Pseudo-Random Number Generators for Vector Processors and Multicore Processors.” *Journal of Modern Applied Statistical Methods*, **14**(1), 308–334. [doi:10.22237/jmasm/1430454120](https://doi.org/10.22237/jmasm/1430454120). URL <http://digitalcommons.wayne.edu/jmasm/vol14/iss1/23>.
- Gentle JE (2003). *Random Number Generators and Monte Carlo Methods*. 2nd edition. Springer-Verlag, New York.
- Knuth DE (1973). *The Art of Computer Programming. Volume 1: Fundamental Algorithms*. 2nd edition. Addison-Wesley, Boston, MA.
- Knuth DE (1981). *The Art of Computer Programming. Volume 2: Seminumerical Algorithms*. 2nd edition. Addison-Wesley, Boston, MA.
- L’Ecuyer P (1996). “Maximally Equidistributed Combined Tausworthe Generators.” *Mathematics of Computation*, **65**(213). [doi:10.1090/S0025-5718-96-00696-5](https://doi.org/10.1090/S0025-5718-96-00696-5). Accessed: 2019-12-31, URL <https://pdfs.semanticscholar.org/7c7e/19630879d9981a5ac1ccbf6c2233d7c5e772.pdf>.
- L’Ecuyer P, Munger D, Oreshkin B, Simard R (2016). “Random Numbers for Parallel Computers: Requirements and Methods, With Emphasis on GPUs.” *Mathematics and Computers in Simulation*. [doi:10.1016/j.matcom.2016.05.005](https://doi.org/10.1016/j.matcom.2016.05.005). URL <https://www.iro.umontreal.ca/~lecuyer/myftp/papers/parallel-rng-imacs.pdf>.

- L'Ecuyer P, Simard R (2001). "On the Performance of Birthday Spacings Tests with Certain Families of Random Number Generators." *Mathematics and Computers in Simulation*, **55**(1–3), 131–137. doi:10.1016/S0378-4754(00)00253-6.
- L'Ecuyer P, Simard R (2007). "TestU01: A C Library for Empirical Testing of Random Number Generators." *ACM Transactions on Mathematical Software (TOMS)*, **33**(4). doi:10.1145/1268776.1268777. Accessed: 2019-12-26, URL <http://simul.iro.umontreal.ca/testu01/tu01.html>.
- Marsaglia G (1995). "The Marsaglia Random Number CDROM, with The Diehard Battery of Tests of Randomness." Accessed: 2019-12-27, URL <https://web.archive.org/web/20160125103112/http://stat.fsu.edu/pub/diehard/>.
- Marsaglia G (2003). "Xorshift RNGs." *Journal of Statistical Software*, **8**(14). doi:10.18637/jss.v008.i14. URL <https://www.jstatsoft.org/article/view/v08i14>.
- Mascagni M, Srinivasan A (2004). "Parameterizing Parallel Multiplicative Lagged-Fibonacci Generators." *Parallel Computing*, **30**(7), 899–916. doi:10.1016/j.parco.2004.06.001.
- O'Neill M (2014). "PCG: A Family of Simple Fast Space-Efficient Statistically Good Algorithms for Random Number Generation." *Report HMC-CS-2014-0905*, Harvey Mudd College, CS dept. Accessed: 2019-12-26, URL <http://www.pcg-random.org>.
- Overton M (2011). "Fast, High-Quality, Parallel Random-Number Generators: Comparing Implementations." *Dr. Dobb's Journal*. Accessed: 2020-2-25, URL www.drdoobs.com/tools/fast-high-quality-parallel-random-number/231000484.
- Vigna S (2015). "splitmix64.c source-code." Accessed: 2020-3-21, URL <http://xorshift.di.unimi.it/splitmix64.c>.
- Vigna S (2018). "xoshiro/xoroshiro Generators and the PRNG Shootout." Accessed: 2019-12-26, URL <http://prng.di.unimi.it>.

A. SplitMix64

The code below is identical to Vigna's except that the names of two identifiers were changed to be consistent with the conventions of this paper.

```
/* Written in 2015 by Sebastiano Vigna (vigna@acm.org)
```

```
To the extent possible under law, the author has dedicated all copyright  
and related and neighboring rights to this software to the public domain  
worldwide. This software is distributed without any warranty.
```

```
See <http://creativecommons.org/publicdomain/zero/1.0/>. */
```

```
#include <stdint.h>
```

```
/* This is a fixed-increment version of Java 8's SplittableRandom generator  
See http://dx.doi.org/10.1145/2714064.2660195 and  
http://docs.oracle.com/javase/8/docs/api/java/util/SplittableRandom.html
```

```
It is a very fast generator passing BigCrush, and it can be useful if  
for some reason you absolutely want 64 bits of state. */
```

```
static uint64_t sState; /* The state can be seeded with any value. */
```

```
uint64_t splitMix64 () {  
    uint64_t z = (sState += 0x9e3779b97f4a7c15);  
    z = (z ^ (z >> 30)) * 0xbf58476d1ce4e5b9;  
    z = (z ^ (z >> 27)) * 0x94d049bb133111eb;  
    return z ^ (z >> 31);  
}
```

B. SplitMix32

```
// Copyright 2020 Mark A. Overton
//
// Licensed under the Apache License, Version 2.0 (the "License");
// you may not use this file except in compliance with the License.
// You may obtain a copy of the License at
//
//     http://www.apache.org/licenses/LICENSE-2.0
//
// Unless required by applicable law or agreed to in writing, software
// distributed under the License is distributed on an "AS IS" BASIS,
// WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
// See the License for the specific language governing permissions and
// limitations under the License.

#include <stdint.h>
static uint32_t sState; // State can be seeded with any value, including 0.

uint32_t splitMix32 () {
    sState = 1664525u * (sState + 314159265u); // Multiplier is from Knuth.
    uint32_t z = sState;
    z = (z ^ (z >> 15)) * 0x5ce4e5b9u;
    z = (z ^ (z >> 13)) * 0x1331c1ebu;
    return z ^ (z >> 15);
}
```

C. Good Multiplier-Rotation Pairs for RomuMono32

In the table below, the two *RM* columns apply when the rotation is done before the multiplication, and the *MR* columns apply when that order is reversed. The randomness of the generator is the same for both orders. The two *base* columns are the lowest values in the largest seed-blocks, and the *bits* columns are the corresponding numbers of bits of seed those blocks can accommodate.

| Multiplier | Lrot | <i>d</i> -value | RM-base | RM-bits | MR-base | MR-bits |
|------------|------|-----------------|------------|---------|------------|---------|
| 2540121707 | 14 | 2 | 437125826 | 31 | 1 | 31 |
| 3731015275 | 18 | 3 | 1 | 31 | 1564370705 | 30 |
| 2336447867 | 16 | 43 | 3779345575 | 28 | 2076771216 | 29 |
| 3611795771 | 12 | 47 | 342645537 | 28 | 1156979152 | 29 |
| 3952805931 | 13 | 157 | 1159051389 | 26 | 2515406761 | 26 |
| 3276993211 | 14 | 184 | 226123367 | 27 | 1335894518 | 27 |
| 4084487243 | 12 | 347 | 1339372056 | 26 | 2141099809 | 25 |
| 4127380763 | 12 | 397 | 606354474 | 26 | 1315333761 | 25 |
| 3563976171 | 16 | 420 | 2941035005 | 25 | 1377002680 | 25 |
| 3651999659 | 14 | 429 | 3209498982 | 25 | 2227458680 | 26 |
| 3365008619 | 12 | 476 | 874642096 | 25 | 2055179747 | 26 |
| 3953463755 | 16 | 514 | 4174029353 | 26 | 2817505743 | 25 |
| 3989591211 | 13 | 592 | 3721670923 | 25 | 3643450503 | 25 |
| 3332453915 | 14 | 650 | 3818020724 | 25 | 1456733700 | 25 |
| 3586487947 | 16 | 681 | 3925643197 | 25 | 3307401597 | 25 |
| 4272641883 | 16 | 709 | 3002109605 | 25 | 2625009597 | 25 |
| 3525693099 | 14 | 758 | 144357697 | 25 | 3440120341 | 25 |
| 3690361499 | 14 | 768 | 2120353896 | 24 | 2598302657 | 25 |
| 1698147467 | 13 | 775 | 1398136959 | 25 | 2483006683 | 25 |
| 3319523819 | 15 | 838 | 3085144041 | 25 | 2129925777 | 24 |
| 3256525067 | 14 | 840 | 1109021212 | 24 | 1180712204 | 25 |
| 3552236683 | 13 | 890 | 1495506820 | 25 | 2898009531 | 25 |
| 3839154475 | 15 | 913 | 1001325849 | 25 | 1262306147 | 24 |
| 3636548587 | 13 | 920 | 3509505418 | 25 | 3475628570 | 25 |
| 4193921835 | 14 | 937 | 1699403590 | 25 | 1473717306 | 24 |
| 4074675915 | 15 | 1051 | 1274742985 | 24 | 1207307984 | 25 |
| 3589580459 | 16 | 1070 | 2080133737 | 25 | 3604530546 | 24 |
| 3937305835 | 15 | 1077 | 1503300770 | 24 | 240484477 | 24 |
| 3291786395 | 15 | 1121 | 2583623346 | 24 | 652306397 | 24 |
| 4070215643 | 15 | 1171 | 4077664074 | 24 | 2173685177 | 24 |
| 3603758123 | 12 | 1144 | 3012437991 | 24 | 2526760669 | 25 |
| 3475128667 | 14 | 1188 | 2293176165 | 24 | 3061649031 | 24 |
| 3852709339 | 15 | 1256 | 2753678026 | 24 | 2610555647 | 24 |
| 3316538187 | 13 | 1271 | 2052216962 | 24 | 1133599252 | 24 |
| 2759225787 | 12 | 1296 | 566123425 | 24 | 3813660306 | 24 |
| 3939093339 | 12 | 1336 | 3269117755 | 24 | 2086898213 | 24 |
| 3953739083 | 16 | 1338 | 534924345 | 24 | 3580323618 | 24 |

| Multiplier | Lrot | d -value | RM-base | RM-bits | MR-base | MR-bits |
|------------|------|------------|------------|---------|------------|---------|
| 2662206315 | 15 | 1370 | 3876606173 | 24 | 2061627875 | 24 |
| 1422968075 | 16 | 1377 | 332616298 | 24 | 3202323436 | 24 |
| 4219099339 | 15 | 1379 | 2712314424 | 24 | 3285589456 | 24 |
| 3715414331 | 13 | 1388 | 111908111 | 24 | 4000728640 | 24 |
| 3662642315 | 14 | 1406 | 781158169 | 25 | 2775093201 | 24 |
| 3240747339 | 15 | 1443 | 2061641557 | 24 | 109294508 | 24 |
| 3505407659 | 16 | 1507 | 2349523361 | 24 | 327552291 | 24 |
| 3603832939 | 14 | 1535 | 2430830656 | 24 | 883175065 | 24 |
| 2334149515 | 15 | 1560 | 2225828117 | 24 | 3989107689 | 24 |
| 3804926571 | 12 | 1563 | 3116634076 | 24 | 3686568596 | 24 |
| 3265651915 | 13 | 1570 | 2075741463 | 24 | 3938950213 | 24 |
| 3687027755 | 15 | 1620 | 2410875726 | 24 | 1183455639 | 24 |
| 3586546027 | 16 | 1648 | 3460233636 | 24 | 2419013898 | 24 |
| 3602222507 | 16 | 1699 | 4258250079 | 24 | 3864885105 | 24 |
| 3650858107 | 14 | 1704 | 3375209167 | 24 | 3796856390 | 24 |
| 3437182379 | 12 | 1710 | 1800368255 | 23 | 2177538044 | 24 |
| 4112721003 | 14 | 1726 | 3545442555 | 24 | 3432760323 | 24 |
| 3697088363 | 15 | 1743 | 1054856781 | 24 | 3645416328 | 24 |
| 3663407467 | 13 | 1790 | 3930504896 | 24 | 2046051146 | 24 |
| 3235068267 | 16 | 1797 | 3421127945 | 24 | 1653095304 | 23 |
| 4229866859 | 12 | 1872 | 2184042351 | 23 | 3065271249 | 24 |
| 3804923435 | 14 | 1956 | 1837917379 | 23 | 2079520793 | 23 |
| 2917649707 | 16 | 1969 | 1061812134 | 23 | 2096832732 | 24 |
| 2836121387 | 14 | 1991 | 410197697 | 23 | 2228488688 | 24 |
| 2273517239 | 14 | 1995 | 2148212120 | 23 | 3778174847 | 23 |

2020-5-9

Author info:

Mark A. Overton

1026 W. El Norte Pkwy #56

Escondido, CA 92026

Email: markovertondev@gmail.comURL: <http://www.romu-random.org>