

# Batched Ranged Random Integer Generation

Nevin Brackett-Rozinsky<sup>1</sup> | Daniel Lemire<sup>2</sup>

<sup>1</sup>Limerick, Maine, USA  
nevin.brackettrozinsky@gmail.com

<sup>2</sup>Data Science Research Center, Université  
du Québec (TELUQ), Montreal, Quebec,  
H2S 3L5, Canada

## Correspondence

Daniel Lemire, Data Science Research  
Center, Université du Québec (TELUQ),  
Montreal, Quebec, H2S 3L5, Canada  
Email: daniel.lemire@teluq.ca

## Funding information

Natural Sciences and Engineering Research  
Council of Canada, Grant Number:  
RGPIN-2024-03787

Pseudorandom values are often generated as 64-bit binary words. These random words need to be converted into ranged values without statistical bias. We present an efficient algorithm to generate multiple independent uniformly-random bounded integers from a single uniformly-random binary word, without any bias. In the common case, our method uses one multiplication and no division operations per value produced. In practice, our algorithm can more than double the speed of unbiased random shuffling for small to moderately large arrays.

## KEYWORDS

Random number generation, Rejection method, Randomized algorithms

## 1 | INTRODUCTION

Random numbers are useful in many computer programs. Most programming languages provide a method to generate uniformly random or pseudorandom integers in the range  $[0, 2^L)$  for some  $L$ , commonly 64. We refer to such numbers as  $L$ -bit random words, and the functions which produce them as *random number generators*. Pseudorandom generators execute algorithms to produce a sequence of numbers that approximate the properties of random numbers, starting from a given *seed* (typically a few bytes). There are a variety of efficient techniques to generate high-quality pseudorandom binary words, such as the Mersenne Twister [1], linear congruential generators [2, 3, 4, 5, 6], PCG [7], and so forth [8, 9]. Although they do not produce truly random outputs, many pseudorandom generators can pass rigorous statistical tests, and are considered to be random in practice [10].

Random binary words are readily available, yet applications often require uniformly random integers from other bounded ranges, such as  $[0, b)$  for some  $b$ . We refer to these numbers as bounded random integers, or *dice rolls*, and we focus on the case where  $0 < b \leq 2^L$ . We are particularly interested in applications which require multiple dice rolls, and where  $b$  may be chosen dynamically rather than known at compile time, such as when shuffling arrays or selecting random samples [11]. For example, we consider the Fisher-Yates random shuffle described by Knuth [12]

and by Durstenfeld [13], which we restate here in [Algorithm 1](#).

---

**Algorithm 1** – Fisher-Yates random shuffle

---

**Require:** Source of uniformly random integers in bounded ranges

**Require:** Array  $A$  made of  $n$  elements indexed from 0 to  $n-1$

**Ensure:** All  $n!$  permutations of  $A$  are equiprobable

```

1: for  $i = n-1, \dots, 1$  do
2:    $j \leftarrow$  random integer in  $[0, i]$ 
3:   exchange  $A[i]$  and  $A[j]$ 

```

---

Algorithms such as random shuffling are commonplace in simulations and other important applications [14, 15, 16, 17, 18, 19, 20]. Accordingly, shuffling algorithms are provided in the standard libraries for most programming languages, and there has also been much work done on parallelizing random permutations [21, 22, 23, 24, 25]. In the present work we demonstrate an efficient method to roll dice in batches, which can improve the performance of shuffles and similar algorithms by reducing the number of calls to a random number generator. Previous batched dice rolling algorithms were based on division, whereas ours is based on multiplication which can be faster.

Perhaps surprisingly, the computational cost of converting binary words into ranged integers can be critical to good performance. For example, Lemire [26] showed that by simplifying the conversion to avoid most division operations, the practical performance of a random shuffle could be made up to eight times faster. We aim to show that the performance can be further multiplied in some cases, through the use of batched dice rolls.

Our main theoretical result is in § 4, where we present and prove the correctness of an algorithm to generate multiple independent bounded random integers from a single random binary word, using in the common case only one multiplication and zero division operations per die roll. The method is based on an existing algorithm due to Lemire that generates one such number at a time, which we summarize in § 2. Our proof uses mixed-radix notation as described in § 3.

In § 5 we demonstrate how our theoretical result can be implemented as a practical algorithm, in § 6 we apply it to the task of shuffling an array, and in § 7 we show the results of our experiments, with timing measurements illustrating the speed of array shuffling with and without batched dice rolls. Finally in § 8 we summarize our results.

## 1.1 | Mathematical notation

We only consider non-negative integers. We use “ $\otimes$ ” to denote full-width multiplication:  $a \otimes b = (x, y)$  means  $x$  and  $y$  are integers such that  $2^L x + y = ab$ , and  $0 \leq y < 2^L$ . On x64 systems, a full-width multiplication requires only a single instruction, while ARM systems provide full-width multiplication with two instructions.

We use “ $\div$ ” to denote integer division:  $a \div b = \lfloor a/b \rfloor$  is the greatest integer less than or equal to  $a/b$ . We use “ $\bmod$ ” to denote the Euclidean remainder:  $(a \bmod b) = a - b \cdot (a \div b)$ . Because we use only non-negative integers, we have  $0 \leq (a \bmod b) < b$ . Note that division and remainder instructions are often slow in practice, and may have a latency of 18 cycles compared to merely 3 cycles for a multiplication on a recent Intel processor (e.g. Ice Lake) [27].

We use pi notation “ $\prod$ ” to denote products, and we omit the bounds when they can be inferred from context. Thus if  $b_i$  is defined for each  $i$  from 1 through  $k$ , then  $\prod b_i = b_1 b_2 \cdots b_k$ . We also use sigma notation “ $\sum$ ” to denote sums, and an underlined superscript to denote the falling factorial:  $n^{\underline{k}} = n! / (n-k)! = n(n-1) \cdots (n-(k-1))$ .

## 2 | EXISTING ALGORITHMS

There are many ways to generate bounded random integers from random bits, and it is nontrivial to do so efficiently. Several widely-used algorithms are described by Lemire [26], including a then-novel strategy to avoid expensive division operations. This method has now been adopted by several major systems: GNU libstdc++, Microsoft standard C++ library, the Linux kernel, and the standard libraries of the Go, Swift, Julia, C# and Zig languages.

It works as follows. Let  $r$  be a uniformly random integer in  $[0, 2^L)$ , and  $[0, b)$  the desired target range with  $0 < b \leq 2^L$ . Perform the full-width multiplication  $b \otimes r = (x, y)$ . Now if  $y \geq (2^L \bmod b)$ —which we call *Lemire's criterion*—then  $x$  is a uniformly random integer in the range  $[0, b)$ . Otherwise, try again with a new  $r$ . That is, apply the rejection method [28]. A detailed proof of the correctness of Lemire's method can be found in [26]. We provide a simplified sketch of it here.

**Lemma 1** *Lemire's method produces uniformly random integers in the range  $[0, b)$ .*

**Proof** Every list of  $2^L - (2^L \bmod b)$  consecutive integers contains the same number of multiples of  $b$ , because the length of that list is divisible by  $b$ . Thus  $br$ , which is a uniformly random multiple of  $b$  in  $[0, 2^L b)$ , is equally likely to land in each of the intervals  $[2^L x + (2^L \bmod b), 2^L(x + 1))$ , where  $x$  is an integer in  $[0, b)$ , because they all have that same length. Lemire's criterion accepts the value  $x$  when  $br$  is in such an interval, and rejects it otherwise, which ensures that the results are uniformly random.

In the common case, Lemire's method uses one  $L$ -bit random word and one multiplication per die roll. Sometimes it requires computing the remainder  $(2^L \bmod b)$ , which may involve a division instruction. However, the division can often be avoided. Because  $(2^L \bmod b) < b$ , it is only necessary to compute  $(2^L \bmod b)$  when  $y < b$ . If  $b$  is much smaller than  $2^L$  the division is often avoided, and when required it is computed at most once per die roll. As a practical matter, we note that when using  $L$ -bit unsigned integers, the value  $(2^L \bmod b)$  can be computed in C and similar languages via the expression  $-b \% b$ . This works because  $-b$  evaluates to  $2^L - b$ , which is congruent to  $2^L \bmod b$ .

### 2.1 | Batched dice rolls

One widely-known way to generate bounded random numbers in batches is that, to simulate rolling dice with  $b_1$  and  $b_2$  sides, a single die with  $b_1 b_2$  sides is rolled instead. The resulting number is uniformly random in  $[0, b_1 b_2)$ , and the desired dice rolls can be obtained by taking its remainder and quotient upon dividing by  $b_1$ . This is equivalent to choosing a random square on a  $b_1 \times b_2$  rectangular grid by numbering the squares sequentially, then taking its row and column as the results.

These values are independent and uniformly random integers in  $[0, b_1)$  and  $[0, b_2)$  respectively, and the method generalizes to more dice by taking successive remainders of quotients upon dividing by each  $b_j$ . This corresponds to choosing a random cell in a higher-dimensional rectangular lattice by numbering the cells sequentially, and using its coordinates as the results. Or in a context familiar to many programmers, it is equivalent to computing the indices of an element in a multi-dimensional array, at a random position in the underlying linear storage. This strategy uses fewer random bits than rolling each die separately, however its performance is limited by division operations which are slower in practice than multiplications on commodity processors.

When applied to the task of shuffling an array, this approach gives rise to the factorial number system described by Laisant [29] in 1888, and to the Lehmer code [30] for numbering permutations. Specifically, dice of sizes  $n, n-1, \dots, 2$  can be rolled in a batch by starting with a single  $n!$ -sided die roll, and extracting the required values through division

and remainder. The resulting dice rolls may then be used for a Fisher-Yates shuffle. If  $n!$  is too large for practical purposes, several smaller batches may be rolled instead. We are not aware of any widespread use of this batched shuffling method, which we attribute to its reliance on relatively slow division operations.

Our approach to rolling a batch of dice works somewhat differently. Rather than extract the values via division, we instead build them up through multiplication. Specifically, we extend Lemire's method in order to generate multiple bounded random integers from a single random word, with zero division operations in the common case.

### 3 | MIXED-RADIX NUMBERS

We use mixed-radix notation in the proof of our main result. Mixed-radix notation is a positional number system in which each digit position has its own base. This contrasts with decimal notation where every digit is in base 10, or binary where every digit is in base 2. Mixed-radix numbers are allowed, but not required, to have a different base for each digit. Each base is a positive integer, and each digit is a non-negative integer smaller than its base.

We denote both the bases and digits of a mixed-radix number with ordered tuples beginning with the most-significant digit, and use context to distinguish them. For example, a 2-digit mixed-radix number in base  $(b_1, b_2)$  whose digits are  $(a_1, a_2)$  represents the value  $a = a_1 b_2 + a_2$ . A  $k$ -digit mixed-radix number in base  $(b_1, b_2, \dots, b_k)$  can represent all the integers from 0 through  $b - 1$ , where  $b = b_1 b_2 \cdots b_k$ , and this representation is unique. If its digits are  $(a_1, a_2, \dots, a_k)$  then it represents the value:

$$a = \sum_{i=1}^k \left( a_i \prod_{j>i}^k b_j \right) = a_1 (b_2 b_3 \cdots b_k) + a_2 (b_3 b_4 \cdots b_k) + \cdots + a_{k-1} b_k + a_k$$

Given a non-negative integer  $a < \prod b_i$ , its mixed-radix representation can be obtained by taking the remainders of successive quotients when dividing by  $b_k, b_{k-1}, \dots, b_1$ , meaning all the  $b_i$  in reverse order.

Mixed-radix notation gives us a concise way to describe the ideas we discussed in § 2.1. If the digits of a mixed-radix number are interpreted as coordinates in a  $k$ -dimensional rectangular lattice, then its value matches the label of that cell when they are numbered sequentially. Or if its digits are used as indices in a  $k$ -dimensional C-style array, then its value gives the corresponding position in the underlying linear storage. We know that a uniformly random cell in a rectangular lattice has independent and uniformly random coordinates, and this translates directly to mixed-radix numbers as seen in Lemma 2.

**Lemma 2** *If a uniformly random integer  $a$  in  $[0, b)$  is written as a mixed-radix number in base  $(b_1, b_2, \dots, b_k)$ , where  $b = \prod b_i$ , then its digits  $(a_1, a_2, \dots, a_k)$  are independent and uniformly random integers in the ranges  $0 \leq a_i < b_i$ .*

**Proof** Every possible sequence of digits is equally likely, so each  $a_i$  takes all integer values in  $[0, b_i)$  with uniform probability regardless of the values of the other digits.

### 4 | MAIN RESULT

We are now ready to present our main result, which is an efficient method for rolling a batch of dice. We start with a random word, and through repeated multiplications build up the digits of a mixed-radix number. When appropriate conditions are met for that number to be uniformly random, then its digits are uniform and independent.

**Theorem 1** Let  $r_0$  be an  $L$ -bit random word, meaning  $r_0$  is a uniformly random integer in  $[0, 2^L)$ . Let  $(b_1, b_2, \dots, b_k)$  be positive integers with  $b = \prod b_i \leq 2^L$ . Starting with  $r_0$  and  $b_1$ , for each  $i$  from 1 to  $k$ , perform the full-width multiplication  $b_i \otimes r_{i-1}$  and set  $a_i$  to the most significant  $L$  bits and  $r_i$  to the least significant  $L$  bits of the  $2L$ -bit result:

$$\begin{aligned} (a_1, r_1) &\leftarrow b_1 \otimes r_0 \\ (a_2, r_2) &\leftarrow b_2 \otimes r_1 \\ &\vdots \\ (a_k, r_k) &\leftarrow b_k \otimes r_{k-1} \end{aligned}$$

At the end of this process, if  $r_k \geq (2^L \bmod b)$  is satisfied, then the  $a_i$  are independent and uniformly random integers in the ranges  $0 \leq a_i < b_i$ .

**Proof** We will first show that when the conditions of the theorem are met, if the resulting values  $(a_1, a_2, \dots, a_k)$  are interpreted as the digits of a mixed-radix number  $a$  in base  $(b_1, b_2, \dots, b_k)$ , then  $a$  is a uniformly random integer in  $[0, b)$ . Once this is established, we will invoke [Lemma 2](#).

For each  $i$  from 1 to  $k$ , the full-width product  $b_i \otimes r_{i-1} = (a_i, r_i)$  means  $b_i r_{i-1} = 2^L a_i + r_i$ . Since both  $r_{i-1}$  and  $r_i$  are in  $[0, 2^L)$ , this implies  $0 \leq a_i < b_i$ . So each  $a_i$  is a valid mixed-radix digit for base  $b_i$ , and  $a$  is well-defined. Let  $c_i$  be the value obtained by truncating  $a$  to its first  $i$  digits. In other words,  $c_i$  is the value represented by the mixed-radix number  $(a_1, a_2, \dots, a_i)$  in base  $(b_1, b_2, \dots, b_i)$ . Thus  $c_1 = a_1$ , and  $c_i = b_i c_{i-1} + a_i$  for  $1 < i \leq k$ .

Let  $b_1 b_2 \cdots b_i$  be the product of  $b_1$  through  $b_i$ . We claim that  $(b_1 b_2 \cdots b_i) \otimes r_0 = (c_i, r_i)$ , and we prove it by finite induction on  $i$ . As a base case, the claim is true for  $i = 1$  because  $b_1 \otimes r_0 = (a_1, r_1) = (c_1, r_1)$ . The claim is equivalent to  $r_0(b_1 b_2 \cdots b_i) = 2^L c_i + r_i$ , and when  $1 < i \leq k$  we use the inductive hypothesis that this equation holds for  $i-1$  in order to prove it for  $i$ :

$$\begin{aligned} r_0(b_1 b_2 \cdots b_i) &= r_0(b_1 b_2 \cdots b_{i-1}) b_i \\ &= (2^L c_{i-1} + r_{i-1}) b_i && \text{(inductive hypothesis)} \\ &= 2^L b_i c_{i-1} + b_i r_{i-1} \\ &= 2^L b_i c_{i-1} + 2^L a_i + r_i && \text{(definition of } a_i \text{ and } r_i) \\ &= 2^L (b_i c_{i-1} + a_i) + r_i \\ &= 2^L c_i + r_i && \text{(formula for } c_i) \end{aligned}$$

This completes the induction and proves the claim for each  $i$  from 1 to  $k$ . We know that  $c_k = a$  and  $\prod b_i = b$ , hence substituting  $i \rightarrow k$  in the claim gives  $b \otimes r_0 = (a, r_k)$ . This is the full-width product of  $b$  with a random word, and  $b \leq 2^L$ , so we can apply [Lemire's criterion](#) [26]: if  $r_0$  is an  $L$ -bit random word such that  $r_k \geq (2^L \bmod b)$ , then  $a$  is a uniformly random integer in  $[0, b)$ .

By [Lemma 2](#), since  $a$  is uniformly random in  $[0, b)$ , its mixed-radix digits in base  $(b_1, b_2, \dots, b_k)$  are independent and uniformly random in the ranges  $[0, b_i)$ . But those digits are  $(a_1, a_2, \dots, a_k)$ , so the theorem is proved. Each  $a_i$  produced this way is a uniformly random integer in  $[0, b_i)$ , and the  $a_i$  are independent, provided that  $r_k \geq (2^L \bmod b)$  and  $r_0$  was picked uniformly at random.

## 4.1 | Worked example

[Theorem 1](#) provides an efficient method to roll a batch of dice, however its description is rather abstract. To develop some familiarity with the process, let us consider a small example using  $L = 4$ -bit integers. Suppose we want to flip

a coin (which is equivalent to rolling a 2-sided die) and also roll a standard 6-sided die. Then  $b_1 = 2$ ,  $b_2 = 6$ , and  $b = 2 \cdot 6 = 12$ . Since  $12 \leq 16 = 2^4$ , this batch of dice can indeed be rolled using 4-bit random numbers. There are 16 possible values for  $r_0$ , and we list all of them in [Table 1](#). The columns of the table show the steps in our method.

The first step is to multiply  $r_0$  by  $b_1$ , which yields  $2r_0$  in this example. Then  $a_1$  is the number of times 16 goes into that value, and  $r_1$  is the remainder. The next step is to multiply  $r_1$  by  $b_2$ , which yields  $6r_1$ , and find  $a_2$  and  $r_2$  in the same way. Finally, Lemire's criterion tells us that if  $r_2 < (16 \bmod 12) = 4$ , then the batch must be rejected and rolled again. We indicate this with an X in the right-most column of the table.

**TABLE 1** Rolling a 2-sided die and a 6-sided die with 4 bits

$r_0$	$2r_0$	$a_1$	$r_1$	$6r_1$	$a_2$	$r_2$	$r_2 < 4$
0	0	0	0	0	0	0	X
1	2	0	2	12	0	12	
2	4	0	4	24	1	8	
3	6	0	6	36	2	4	
4	8	0	8	48	3	0	X
5	10	0	10	60	3	12	
6	12	0	12	72	4	8	
7	14	0	14	84	5	4	
8	16	1	0	0	0	0	X
9	18	1	2	12	0	12	
10	20	1	4	24	1	8	
11	22	1	6	36	2	4	
12	24	1	8	48	3	0	X
13	26	1	10	60	3	12	
14	28	1	12	72	4	8	
15	30	1	14	84	5	4	

The rows with an X in the last column will be rejected, so let us look at the other rows, which produce  $a_1$  and  $a_2$  as output. We see there are 12 such rows, and we know they are all equally likely because  $r_0$  is chosen uniformly at random. There is one such row for each possible pair of values  $(a_1, a_2)$  with  $a_1$  in  $\{0, 1\}$  and  $a_2$  in  $\{0, 1, 2, 3, 4, 5\}$ , so we conclude that in this example the coin flip and die roll are indeed fair and independent. [Theorem 1](#) proves that this will always be the case.

## 5 | IMPLEMENTATION

In some applications the values of  $b_i$  are known ahead of time, possibly even at compile time. In that case the value of  $t = (2^L \bmod b)$  can be precomputed, and [Theorem 1](#) can be implemented succinctly as shown in [Algorithm 2](#). In other applications the values of  $b_i$  are not known ahead of time. In that case the threshold  $t$  must be computed when

needed, which involves a division operation. It can be avoided when  $r_k \geq b$ , as shown in [Algorithm 3](#).

---

**Algorithm 2** – Batched dice rolls (known threshold)

---

**Require:** Source of uniformly random integers in  $[0, 2^L)$

**Require:** Target intervals  $[0, b_i)$  for  $i$  in  $1 \dots k$ , with  $1 \leq \prod b_i \leq 2^L$

**Require:** The value  $t = (2^L \bmod \prod b_i)$

**Ensure:** The  $a_i$  are independent and uniformly random in  $[0, b_i)$

```

1: repeat
2:    $r \leftarrow$  random integer in  $[0, 2^L)$ 
3:   for  $i$  in  $1 \dots k$  do
4:      $(a_i, r) \leftarrow b_i \otimes r$  ▷ Full-width multiply
5: until  $r \geq t$ 
6: return  $(a_1, a_2, \dots, a_k)$ 

```

---



---

**Algorithm 3** – Batched dice rolls (unknown threshold)

---

**Require:** Source of uniformly random integers in  $[0, 2^L)$

**Require:** Target intervals  $[0, b_i)$  for  $i$  in  $1 \dots k$ , with  $1 \leq \prod b_i \leq 2^L$

**Ensure:** The  $a_i$  are independent and uniformly random in  $[0, b_i)$

```

1:  $r \leftarrow$  random integer in  $[0, 2^L)$ 
2: for  $i$  in  $1 \dots k$  do
3:    $(a_i, r) \leftarrow b_i \otimes r$  ▷ Full-width multiply
4:  $b \leftarrow \prod b_i$ 
5: if  $r < b$  then
6:    $t \leftarrow (2^L \bmod b)$ 
7:   while  $r < t$  do
8:      $r \leftarrow$  random integer in  $[0, 2^L)$ 
9:     for  $i$  in  $1 \dots k$  do
10:       $(a_i, r) \leftarrow b_i \otimes r$ 
11: return  $(a_1, a_2, \dots, a_k)$ 

```

---

We must have  $b \leq 2^L$  in order to use [Theorem 1](#), and for [Algorithm 3](#) we would prefer to have  $b$  at least an order of magnitude smaller than  $2^L$ . If  $b$  is too close to  $2^L$  then there is a high probability of taking the slow path that needs to calculate  $t$ , and possibly having to reroll the whole batch of dice. In many applications it is possible to bound the  $b_i$  in such a way that a value  $u$  satisfying  $b \leq u \ll 2^L$  is known ahead of time. This allows for a faster implementation that avoids computing  $b$  most of the time, by enclosing lines 4–10 of [Algorithm 3](#) within an “if  $r < u$ ” block.

Whichever version of the algorithm is used, at its core is a loop containing a single full-width multiplication “ $b_i \otimes r$ ”. The value of  $b_i$  is known, but each pass through the loop computes the value of  $r$  that will be used for the next iteration. This constitutes a loop-carried dependency, and it means that each iteration must complete before the next can begin. Modern commodity processors generally have the ability to carry out more than one operation at a time, a feature

called instruction-level parallelism or superscalarity [31]. Thus we expect that the computation of **Algorithm 3** can be executed while other operations are completed (e.g. memory loads and stores), or the processor might speculatively use generated values of  $a_j$  for upcoming operations. More than one batch of dice rolls may also be interleaved.

## 6 | SHUFFLING ARRAYS

The Fisher-Yates shuffle of **Algorithm 1** is widely used for permuting the elements of an array. It requires  $n-1$  dice rolls to shuffle  $n$  elements. With a traditional implementation this involves  $n-1$  calls to a random number generator, but by rolling dice in batches of  $k$  we can reduce that by a factor of  $k$ . As written, **Algorithm 3** requires the computation of the product  $b = \prod b_j$ . In a batched shuffle, that is  $b = n^k$ . One key insight is that we can replace an exact computation by an upper bound  $u \geq b$ . As long as  $r_k \geq u$ , there is no need to compute  $b$  exactly.

We denote by  $n_k$  the largest array length  $n$  at which we will use batches of  $k$  dice. We want an upper bound  $u \geq n^k$ , ideally with  $u \ll 2^L$  so the fast path succeeds with high probability. One possible choice is  $u_k = n_k^k$ , which is the largest product we will ever see for a batch of  $k$  dice. We could use  $u_k$  as the upper bound for all batches of size  $k$ , however to improve efficiency we would like to lower the value of  $u$  as  $n$  decreases. A convenient approach is, whenever the current upper bound fails and we need to compute the true product, we assign that product to  $u$  and use it for subsequent batches of size  $k$  in the shuffle.

We illustrate this approach in **Algorithm 4**, which carries out the dice rolls and swaps for a single batch of size  $k$ , when there are  $n$  elements to shuffle. It takes an upper bound  $u$  as input, and at the end returns an upper bound for the next iteration. Usually the return value equals the input, however if the batch needed to calculate its true product then the return value equals that product.

---

### Algorithm 4 – Batched partial shuffle

---

**Require:** Source of uniformly random integers in  $[0, 2^L)$

**Require:** Array  $z$  whose first  $n$  elements need to be shuffled

**Require:** Batch size  $k \leq n$  for which  $n^k \leq 2^L$

**Require:** Upper bound  $u \geq n^k$

**Ensure:** Only the first  $(n - k)$  elements of  $z$  remain to be shuffled

```

1:  $r \leftarrow$  random integer in  $[0, 2^L)$ 
2: for  $i$  in  $1 \dots k$  do
3:    $(a_i, r) \leftarrow (n+1-i) \otimes r$  ▷ Full-width multiply
4: if  $r < u$  then
5:    $u \leftarrow n^k$  ▷ Falling factorial
6:    $t \leftarrow (2^L \bmod u)$ 
7:   while  $r < t$  do
8:      $r \leftarrow$  random integer in  $[0, 2^L)$ 
9:     for  $i$  in  $1 \dots k$  do
10:       $(a_i, r) \leftarrow (n+1-i) \otimes r$ 
11: for  $i$  in  $1 \dots k$  do
12:   exchange  $z[a_i]$  and  $z[n-i]$  ▷ Zero-based indexing
13: return  $u$  ▷ For the next batch

```

---



In [Algorithm 4](#), it would be possible to wrap lines 6–10 in another “if  $r < u$ ” block, thus avoiding the computation of  $t$  some fraction of the time. We have omitted this for simplicity. We also note that in practice it is rare to need to update  $u$  at all, so there is little to gain by further optimizing that scenario.

To shuffle a full array, we first select the largest  $k$  such that  $n_k \geq n$ , and set  $u = u_k$ . Then we shuffle in batches of  $k$ , updating  $u$  along the way, until  $n \leq n_{k+1}$ . At that point we set  $u = u_{k+1}$  and shuffle in batches of  $k+1$ , and so forth up to some predetermined maximum batch size. Finally, when the number of remaining elements becomes smaller than the previous batch size, we finish the shuffle with one last batch.

## 6.1 | Batch sizes

The batch size  $k$  presents a tradeoff. On one hand, we want to roll as many dice as we can with each random word. On the other hand, we want each batch to succeed on the first try with high probability so we do not have to reroll. These goals are in opposition, and we seek a balance between them. The question becomes, at what array length  $n$  should we start rolling dice in batches of  $k$ . In other words, what values of  $n_k$  should be used. This depends on the bit-width  $L$ , and can only truly be answered through benchmarks on the target hardware. However, a preliminary analysis can help to identify the right ballpark.

We begin by estimating the computational cost for each die roll. This will enable us to compare the relative performance of different batch sizes  $k$  at each array length  $n$ . If our analysis is successful, then we will be able to use that information to select values of  $n_k$  which minimize the expected computational cost for a Fisher-Yates shuffle that uses our batched dice rolling technique of [Algorithm 4](#).

To estimate the cost for each die roll, let us assume that calling the random number generator is as fast as  $c_{rng} = 2$  multiplications, and dividing is as slow as  $c_{div} = 16$  multiplications. These are conservative estimates, since a fast random number generator reduces the benefit of batching, and a slow division operation increases the cost of rerolling. Using those parameters, and assuming that  $u \approx n^k$ , we can find the cost per element for batches of size  $k$  at each  $n$ .

Each batch has probability  $p_1 = u/2^L \approx n^k/2^L$  of failing the first threshold, which incurs the cost of computing both  $b = n^k$  and  $(2^L \bmod b)$ . We treat this as  $k-1$  multiplications and 1 division, for a cost of  $(c_{div} + k - 1)p_1$ . The probability of rerolling the batch is  $p_2 = (2^L \bmod b)/2^L$ , so the expected number of rolls including the first is  $1/(1-p_2)$ . Every roll incurs the cost of one random number and  $k$  multiplications, for an average of  $(c_{rng} + k)/(1-p_2)$ . The batch produces  $k$  dice rolls, so the expected cost per die roll is:

$$c_{avg} = \frac{1}{k} \left( \frac{c_{rng} + k}{1 - p_2} + (c_{div} + k - 1)p_1 \right)$$

By choosing  $n_k$  so this cost is cheaper with a batch of  $k$  than  $k-1$  for all  $n$  up to  $n_k$ , but not  $n_k + 1$ , we obtain the values shown in [Table 2](#). These are likely to be overestimates for the optimal  $n_k$  because we have omitted some of the cost (e.g. misprediction cost). As a rough attempt to adjust for this, we may reduce the values of  $n_k$  from [Table 2](#) to the next-lower power of 2. We use those reduced values in our experiments, as described in [§ 7](#).

**TABLE 2** Estimated values for  $n_k$ 

$k$	$L = 64$	$L = 32$
2	1 358 187 913	20 724
3	929 104	581
4	26 573	109
5	3225	
6	815	
7	305	
8	146	

## 7 | EXPERIMENTS

We have implemented [Algorithm 4](#) in the C programming language, and used it for a batched Fisher-Yates shuffle. To help ensure that our results are reproducible, we make our source code freely available.<sup>1</sup> We use three different random number generators:

- Our fastest generator is a linear congruential generator proposed by Lehmer [32]. It has good statistical properties [2]. Using a 128-bit integer seed as the state, we multiply it by the fixed 64-bit integer `0xda942042e4dd58b5`. The most-significant 64 bits of the resulting state are returned.
- Our second generator is a 64-bit version of O'Neill's PCG [7]. This relies on a 128-bit parameter  $m$  acting as a multiplier (`0x2360ed051fc65da44385df649fccf645`). With each call, a 128-bit state variable  $s$  is updated:  $s \leftarrow ms + c$  where  $s$  and  $c$  are initialized once. The 64-bit random value is generated from the 128-bit state with a bit rotate and an exclusive or. We use O'Neill's own implementation adapted to our code base.
- Finally, we use ChaCha as a 64-bit cryptographically strong generator [33]. The chosen implementation<sup>2</sup> is written in conventional C, without advanced optimizations.

We measure the speed of our batched shuffling algorithm across a range of array lengths, with a maximum batch size of either two or six. To get accurate measurements, we shuffle the same array until the total elapsed time is 100  $\mu$ s. Each element of the array occupies 64 bits and the entire arrays fit in CPU cache. We compare the results with two other implementations. One is a standard Fisher-Yates shuffle without any batching, but using an efficient bounded-number algorithm [26]. The other is a *naive* batched approach where two dice rolls are extracted from a single random number through division and remainder. Code is provided in [Appendix A](#) for our method with batches of up to six, and for both reference methods.

We use two different computer architectures, Apple M2 and Intel Xeon Gold 6338, as shown in [Table 3](#). We verify that the CPU frequency is stable during the benchmarks: the Intel Xeon Gold 6338's effective frequency remains at 3.2 GHz throughout, while the Apple M2 remains within 10% of its 3.49 GHz maximum frequency.

For each architecture and random number generator, we plot the resulting average time per element to shuffle an array using all four algorithms. See [Figure 1](#), [Figure 2](#) and [Figure 3](#). Normalizing per element keeps the scale of the results stable, and makes the relative speeds easy to see. As part of our benchmark, we check that the minimal time is close to the average time (within 10%).


<sup>1</sup>[https://github.com/lemire/batched\\_random](https://github.com/lemire/batched_random)

<sup>2</sup><https://github.com/nixberg/chacha-rng-c>


**TABLE 3** Systems

Processor	Intel Xeon Gold 6338	Apple M2
Frequency	3.2 GHz	3.49 GHz
Microarchitecture	Ice Lake (x64, 2019)	Avalanche (aarch64, 2022)
Memory	DDR4 (3200 MT/s)	LPDDR5 (6400 MT/s)
Compiler	LLVM 16	Apple/LLVM 14
Cache (LLC)	48 MiB	16 MiB

In these graphs, “shuffle” is the standard unbatched Fisher-Yates shuffle, and “naive shuffle\_2” uses division to roll batches of two dice. The “shuffle\_2” and “shuffle\_6” methods use our multiplication-based technique to roll batches of at most two and six dice respectively. In particular, shuffle\_6 begins using batches of two when the number of elements to shuffle is at most  $2^{30}$  (this is beyond what appears in the graphs), then switches to batches of three when  $n \leq 2^{19}$ , batches of four when  $n \leq 2^{14}$ , batches of five when  $n \leq 2^{11}$ , and batches of six when  $n \leq 2^9$ .




(a) Intel Xeon Gold 6338




(b) Apple M2

**FIGURE 1** Shuffle timings with Lehmer random number generator




(a) Intel Xeon Gold 6338



(b) Apple M2

**FIGURE 2** Shuffle timings with PCG random number generator




**FIGURE 3** Shuffle timings with ChaCha random number generator

We see in the graphs that using our method for batched dice rolls results in a significantly faster shuffle than either reference method, with larger batches having greater benefits. Across most of the graphs, there is a clear trend: unbatched shuffles are the slowest, batches using division are somewhat faster, and batches using multiplication are the fastest. In one case, using the Lehmer generator on the Intel system, the division-based batched shuffle is no faster than the unbatched shuffle, and sometimes slower. This is explained by the relatively slow division instruction on Intel processors, and the high speed of the Lehmer generator.

In another case, using the ChaCha generator on the Apple system, the division-based method is nearly as fast as the multiplication-based shuffle\_2. This reflects the better performance of division instructions on Apple systems [34, 35], and the proportionally greater amount of time taken to generate random numbers with ChaCha, which dominates the time required for the shuffle.

Relative to the unbatched approach, our shuffle\_6 method increases the speed of shuffling by a significant margin across the entire range of sizes we tested, on both platforms, with all three random number generators. This is shown in Figure 4, where the y-axis indicates the relative speed of shuffle\_6 compared to the unbatched shuffle.



**FIGURE 4** Speed ratios between shuffle\_6 and a conventional unbatched shuffle

On the Intel system, the speedups range from 1.4–1.8 $\times$  with the Lehmer generator, 1.8–2.5 $\times$  with the PCG generator, and 2.9–4.7 $\times$  with ChaCha. On the Apple system, the speedups range from 1.5–2.3 $\times$  with the Lehmer generator, 1.9–2.6 $\times$  with the PCG generator, and 2.6–3.9 $\times$  with ChaCha. Overall on both systems, with a fast gen-

erator (e.g. Lehmer or PCG) the speedups are generally 1.5–2.5 $\times$ , and with a slow generator (e.g. ChaCha) they are 2.5–4.5 $\times$ . Notably, with the ChaCha random number generator the speedups approach the theoretical ideal of 2 $\times$  for batches of two, 3 $\times$  for batches of three, and 4 $\times$  for batches of four. The faster generators yield proportionally smaller speedups.

All of the graphs in [Figure 4](#) exhibit a peak somewhere in the middle, with lower values to either side. The decrease at large array sizes is a result of using smaller batches. The decrease at small array sizes is likely due in part to fixed costs being amortized across fewer elements, and it may also represent an opportunity for additional optimization.

## 7.1 | Instruction counts


To better understand these results, we use performance counters to record the number of instructions retired by our functions during the shuffling.<sup>3</sup> The number of instructions does not, by itself, determine the performance because a variety of factors affect the number of instructions that can be executed in a given unit of time. In particular, cache and memory latency limit our maximal speed. Nevertheless, we find that our batched procedures (`shuffle_2` and `shuffle_6`) use significantly fewer instructions per element than a conventional unbatched shuffle. Furthermore, the number of instructions retired per cycle is maintained, or even increased. See [Table 4](#) where we consider the case when there are 16 384 elements in the array to be shuffled. We also present the results in [Figure 5](#).

**TABLE 4** Instructions retired per element and instructions per cycle (16 384 elements)

generator	function	Intel Xeon Gold 6338		Apple M2	
		ins./el.	ins./cycle	ins./el.	ins./cycle
Lehmer	shuffle	18	3.8	18	4.2
	naive shuffle_2	14	4.4	13	4.6
	shuffle_2	13	4.6	14	5.5
	shuffle_6	10	4.0	11	5.8
PCG64	shuffle	26	3.4	24	4.1
	naive shuffle_2	20	3.9	16	4.5
	shuffle_2	19	4.6	17	5.3
	shuffle_6	12	4.0	13	5.6
ChaCha	shuffle	139	2.8	133	5.3
	naive shuffle_2	76	3.1	71	5.1
	shuffle_2	75	3.1	72	5.3
	shuffle_6	39	3.1	39	5.2

The reduction in the number of instructions between a conventional shuffle function and our most aggressively batched one (`shuffle_6`) is about 40% in the case of the Lehmer generator, and rises up to a threefold reduction when using the ChaCha generator. This suggests that the batched shuffles run at a higher speed in large part because they

<sup>3</sup>Modern processors execute more instructions than the program flow strictly requires, a process known as *speculative execution*. Among the speculatively executed instructions, only those necessary for the actual program execution flow are considered *retired*.



**FIGURE 5** Instructions retired per element for arrays of 16384 64-bit elements

use far fewer instructions. In turn, they use fewer instructions because fewer calls to the random number generator are required.

The division-based approach (naive shuffle\_2) uses about as many instructions as our shuffle\_2 function. However, shuffle\_2 tends to retire more instructions per cycle. This aligns with the fact that multiplication is generally faster than division. It might be possible to further tune the performance of the naive approach, however we expect that division would remain a bottleneck.

## 8 | CONCLUSION

We have shown that Lemire’s nearly-divisionless method of generating bounded random integers can be extended to generate multiple such numbers from a single random word. When rolling several dice or shuffling an array, this batched approach can reduce the number of random bits used, without increasing the amount of computation required. In our benchmark tests we saw speedups by a factor of 1.5–2.5 $\times$  when shuffling an array using a fast random number generator (e.g. Lehmer [2] or PCG [7]), and by even greater factors when the random bits are more computationally expensive.

Our results are based on a system-oblivious computational model described in § 6. We expect that further tuning, especially system-specific tuning, might have some additional benefits. Our implementation is in the C language using a popular compiler (LLVM). We expect that our results should carry over to other languages such as C++, Rust, Go, Swift, Java, C# and so forth with relative ease. However, care might be needed to ensure that the generated compiled code is comparable to the result of our C code under LLVM. We expect that our approach can find broad applications. Future work should examine other applications such as sampling algorithms and simulations, as well as the shuffling of very large and very small arrays.

## A | CODE SAMPLES

Unbatched shuffle:

```
// generates a bounded random integer
uint64_t random_bounded(uint64_t range, uint64_t (*rng)(void)) {
    __uint128_t random64bit, multiresult;
    uint64_t leftover;
    uint64_t threshold;
    random64bit = rng();
    multiresult = random64bit * range;
    leftover = (uint64_t)multiresult;
    if (leftover < range) {
        threshold = -range % range;
        while (leftover < threshold) {
            random64bit = rng();
            multiresult = random64bit * range;
            leftover = (uint64_t)multiresult;
        }
    }
    return (uint64_t)(multiresult >> 64);
}

// Fisher-Yates shuffle
void shuffle(uint64_t *storage, uint64_t size,
             uint64_t (*rng)(void)) {
    uint64_t i;
    for (i = size; i > 1; i--) {
        uint64_t nextpos = random_bounded(i, rng);
        uint64_t tmp = storage[i - 1];
        uint64_t val = storage[nextpos];
        storage[i - 1] = val;
        storage[nextpos] = tmp;
    }
}
```

Batched shuffle:

```
// performs k steps of a shuffle
void batched(uint64_t *storage, uint64_t n, uint64_t k,
             uint64_t bound, uint64_t (*rng)(void)) {
    __uint128_t x;
    uint64_t r = rng();
    uint64_t pos1, pos2;
    uint64_t val1, val2;
    uint64_t indexes[k];
    for (uint64_t i = 0; i < k; i++) {
        x = (__uint128_t)(n - i) * (__uint128_t)r;
        r = (uint64_t)x;
        indexes[i] = (uint64_t)(x >> 64);
    }
    if (r < bound) {
        bound = n;
        for (uint64_t i = 1; i < k; i++) {
            bound *= n - i;
        }
        uint64_t t = -bound % bound;
        while (r < t) {
            r = rng();
            for (uint64_t i = 0; i < k; i++) {
                x = (__uint128_t)(n - i) * (__uint128_t)r;
                r = (uint64_t)x;
                indexes[i] = (uint64_t)(x >> 64);
            }
        }
    }
    for (uint64_t i = 0; i < k; i++) {
        pos1 = n - i - 1; pos2 = indexes[i];
        val1 = storage[pos1]; val2 = storage[pos2];
        storage[pos1] = val2; storage[pos2] = val1;
    }
    return bound;
}

// Fisher-Yates shuffle, rolling up to six dice at a time
void shuffle_6(uint64_t *storage, uint64_t size,
              uint64_t (*rng)(void)) {
    uint64_t i = size;
    for (; i > 1 << 30; i--) {
        batched(storage, i, 1, i, rng);
    }
}
```

```

}
// Batches of 2 for sizes up to 2^30 elements
uint64_t bound = (uint64_t)1 << 60;
for (; i > 1 << 19; i -= 2) {
    bound = batched(storage, i, 2, bound, rng);
}
// Batches of 3 for sizes up to 2^19 elements
bound = (uint64_t)1 << 57;
for (; i > 1 << 14; i -= 3) {
    bound = batched(storage, i, 3, bound, rng);
}
// Batches of 4 for sizes up to 2^14 elements
bound = (uint64_t)1 << 56;
for (; i > 1 << 11; i -= 4) {
    bound = batched(storage, i, 4, bound, rng);
}
// Batches of 5 for sizes up to 2^11 elements
bound = (uint64_t)1 << 55;
for (; i > 1 << 9; i -= 5) {
    bound = batched(storage, i, 5, bound, rng);
}
// Batches of 6 for sizes up to 2^9 elements
bound = (uint64_t)1 << 54;
for (; i > 6; i -= 6) {
    bound = batched(storage, i, 6, bound, rng);
}
if (i > 1) {
    batched(storage, i, i - 1, 720, rng);
}
}
}

```

#### Naive batched shuffle:

```

// performs k steps of a shuffle
void naive_batched(uint64_t *storage, uint64_t n, uint64_t k, uint64_t (*rng)(void)) {
    uint64_t pos1, pos2;
    uint64_t val1, val2;
    uint64_t bound = n;
    for (uint64_t i = 1; i < k; i++) {
        bound *= n - i;
    }
    // Next we generate a random integer in [0, bound)
    uint64_t r = random_bounded(bound, rng);
    for (uint64_t i = 0; i < k - 1; i++) {
        pos2 = r % (n - i);
        r /= (n - i);
        pos1 = n - i - 1;
        val1 = storage[pos1];
        val2 = storage[pos2];
        storage[pos1] = val2;
        storage[pos2] = val1;
    }
    pos2 = r;
    pos1 = n - k;
    val1 = storage[pos1];
    val2 = storage[pos2];
    storage[pos1] = val2;
    storage[pos2] = val1;
}

void naive_shuffle_2(uint64_t *storage, uint64_t size,
                    uint64_t (*rng)(void)) {
    uint64_t i = size;
    for (; i > 0x100000000; i--) {
        naive_batched(storage, i, 1, rng);
    }
    for (; i > 1; i -= 2) {
        naive_batched(storage, i, 2, rng);
    }
}

```



## references

- [1] Matsumoto M, Nishimura T. Mersenne Twister: A 623-dimensionally Equidistributed Uniform Pseudo-random Number Generator. *ACM Trans Model Comput Simul* 1998 Jan;8(1):3–30.
- [2] L'Ecuyer P. Tables of linear congruential generators of different sizes and good lattice structure. *Mathematics of Computation* 1999;68(225):249–260.
- [3] L'Ecuyer P, Blouin F, Couture R. A Search for Good Multiple Recursive Random Number Generators. *ACM Trans Model Comput Simul* 1993 Apr;3(2):87–98.
- [4] De Matteis A, Pagnutti S. Parallelization of random number generators and long-range correlations. *Numerische Mathematik* 1988;53(5):595–608.
- [5] Fishman G. Monte Carlo: concepts, algorithms, and applications. corrected edition ed. Berlin: Springer Science & Business Media; 2013.
- [6] L'Ecuyer P. Random Numbers for Simulation. *Commun ACM* 1990 Oct;33(10):85–97.
- [7] O'Neill ME, PCG: A Family of Simple Fast Space-Efficient Statistically Good Algorithms for Random Number Generation; 2014. <https://www.pcg-random.org/pdf/hmc-cs-2014-0905.pdf> [last checked August 2024].
- [8] L'Ecuyer P. History of uniform random number generation. In: *Proceedings of the 2017 Winter Simulation Conference* New York, NY, USA: IEEE Press; 2015. p. 1–17.
- [9] L'Ecuyer P. Random number generation. In: *Handbook of Computational Statistics* Berlin: Springer; 2012.p. 35–71.
- [10] Sleem L, Couturier R. TestU01 and Pracrand: Tools for a randomness evaluation for famous multimedia ciphers. *Multimedia Tools and Applications* 2020;79(33):24075–24088.
- [11] Vitter JS. Random Sampling with a Reservoir. *ACM Trans Math Softw* 1985 Mar;11(1):37–57.
- [12] Knuth DE. *Seminumerical Algorithms*, vol. 2 of *The Art of Computer Programming*. Boston: Addison-Wesley; 1969.
- [13] Durstenfeld R. Algorithm 235: random permutation. *Communications of the ACM* 1964;7(7):420.
- [14] Devroye L. Random Variate Generation for Multivariate Unimodal Densities. *ACM Trans Model Comput Simul* 1997 Oct;7(4):447–477.
- [15] Calvin JM, Nakayama MK. Using Permutations in Regenerative Simulations to Reduce Variance. *ACM Trans Model Comput Simul* 1998 Apr;8(2):153–193.
- [16] Owen AB. Latin Supercube Sampling for Very High-dimensional Simulations. *ACM Trans Model Comput Simul* 1998 Jan;8(1):71–102.
- [17] Osogami T. Finding Probably Best Systems Quickly via Simulations. *ACM Trans Model Comput Simul* 2009 Aug;19(3):12:1–12:19.
- [18] Amrein M, Künsch HR. A Variant of Importance Splitting for Rare Event Estimation: Fixed Number of Successes. *ACM Trans Model Comput Simul* 2011 Feb;21(2):13:1–13:20.
- [19] Hernandez AS, Lucas TW, Carlyle M. Constructing Nearly Orthogonal Latin Hypercubes for Any Nonsaturated Run-variable Combination. *ACM Trans Model Comput Simul* 2012 Nov;22(4):20:1–20:17.
- [20] Hinrichs C, Ithapu VK, Sun Q, Johnson SC, Singh V. Speeding Up Permutation Testing in Neuroimaging. In: *Proceedings of the 26th International Conference on Neural Information Processing Systems NIPS'13, USA*: Curran Associates Inc.; 2013. p. 890–898.

- [21] Shterev ID, Jung SH, George SL, Owzar K. permGPU: Using graphics processing units in RNA microarray association studies. *BMC Bioinformatics* 2010 Jun;11(1):329. <https://doi.org/10.1186/1471-2105-11-329>.
- [22] Langr D, Tvrdík P, Dytrych T, Draayer JP. Algorithm 947: Paraperm—Parallel Generation of Random Permutations with MPI. *ACM Trans Math Softw* 2014 Oct;41(1):5:1–5:26.
- [23] Sanders P. Random permutations on distributed, external and hierarchical memory. *Information Processing Letters* 1998;67(6):305–309.
- [24] Gustedt J. In: *Engineering Parallel In-Place Random Generation of Integer Permutations* Berlin, Heidelberg: Springer Berlin Heidelberg; 2008. p. 129–141.
- [25] Waechter M, Hamacher K, Hoffgaard F, Widmer S, Goesele M. In: *Is Your Permutation Algorithm Unbiased for  $n \neq 2^m$ ?* Berlin, Heidelberg: Springer Berlin Heidelberg; 2012. p. 297–306.
- [26] Lemire D. Fast Random Integer Generation in an Interval. *ACM Transactions on Modeling and Computer Simulation* 2019 Jan;29(1):1–12. <http://dx.doi.org/10.1145/3230636>.
- [27] Abel A, Reineke J. uops.info: Characterizing latency, throughput, and port usage of instructions on intel microarchitectures. In: *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*; 2019. p. 673–686.
- [28] Von Neumann J. Various techniques used in connection with random digits. *National Bureau of Standards Series* 1951;12:36–38.
- [29] Laisant CA. Sur la numération factorielle, application aux permutations. *Bulletin de la Société Mathématique de France* 1888;16:176–183.
- [30] Lehmer DH. Teaching combinatorial tricks to a computer. In: *Proceedings of the Tenth Symposium in Applied Mathematics of the American Mathematical Society USA*: American Mathematical Society; 1960. p. 179–193.
- [31] Shen JP, Lipasti MH. *Modern processor design: fundamentals of superscalar processors*. Waveland Press; 2013.
- [32] Lehmer DH. Mathematical models in large-scale computing units. *Annals of the Computation Laboratory (Harvard University)* 1951;26:141–146.
- [33] Bernstein DJ. ChaCha, a variant of Salsa20. In: *Workshop record of SASC, vol. 8*; 2008. p. 3–5.
- [34] Johnson D, Apple Microarchitecture Research; 2023. <https://dougallj.github.io/applecpu/firestorm.html> [last checked July 2024].
- [35] Fog A. *Instruction tables: Lists of instruction latencies, throughputs and micro-operation breakdowns for Intel, AMD and VIA CPUs*. Copenhagen, Denmark: Copenhagen University College of Engineering; 2022.