

Efficient Computation of Positional Population Counts Using SIMD Instructions

Marcus D. R. Klarqvist^{1,2} | Wojciech Muła³ | Daniel Lemire^{*4}

¹Department of Genetics, University of Cambridge, Downing Street, Cambridge CB2 3EH, United Kingdom

²Wellcome Sanger Institute, Wellcome Genome Campus, Hinxton, Cambridge, CB10 1SA, United Kingdom

³0x80.pl, Wrocław, Poland

⁴TELUQ, Université du Québec, Québec, Canada

Correspondence

*D. Lemire, Université du Québec (TELUQ), 5800, Saint-Denis street, Montreal (Quebec) H2S 3L5, Canada.
Email: daniel.lemire@teluq.ca

Summary

In several fields such as statistics, machine learning, and bioinformatics, categorical variables are frequently represented as one-hot encoded vectors. For example, given 8 distinct values, we map each value to a byte where only a single bit has been set. We are motivated to quickly compute statistics over such encodings.

Given a stream of k -bit words, we seek to compute k distinct sums corresponding to bit values at indexes $0, 1, 2, \dots, k - 1$. If the k -bit words are one-hot encoded then the sums correspond to a frequency histogram.

This multiple-sum problem is a generalization of the population-count problem where we seek the sum of all bit values. Accordingly, we refer to the multiple-sum problem as a *positional population-count*.

Using SIMD (Single Instruction, Multiple Data) instructions from recent Intel processors, we describe algorithms for computing the 16-bit position population count using less than half of a CPU cycle per 16-bit word. Our best approach uses up to 400 times fewer instructions and is up to 50 times faster than baseline code using only regular (non-SIMD) instructions, for sufficiently large inputs.

KEYWORDS:

Vectorization, Population Counts, SIMD Instructions, Software Performance, Bioinformatics, Sequencing, Genomics


1 | INTRODUCTION

In many applications such as deep learning^{1,2,3,4}, indexing⁵, chemistry⁶, cryptography⁷, and bioinformatics^{8,9,10,11} it is desirable to compute the number of set bits in a computer word. This operation is referred to as the population count (popcnt), Hamming weight, or the sideways sum of the word. For example, the machine word `10010010` has a population count of three since there are three set bits. We have previously described efficient subroutines for computing the population count¹² of large arrays that takes advantage of SIMD (Single Instruction, Multiple Data) instructions available on most commodity processors. These instructions operate on wide data registers and reduce the number of instructions required by processing multiple machine words simultaneously.

It is common to represent categorical variables using one-hot (1-of- k) encoding^{13,14} where each categorical value maps to a corresponding bit within a k -bit word and each word may only have a single bit set. As an illustrative example, consider a categorical variable country that has the eight distinct values Australia, Canada, China, France, Japan, Portugal, Spain, and USA. These categorical string values can be unambiguously dictionary-compressed into one-hot encodings (Table 1a) resulting in lower memory usage and more efficient queries. One of the motivations for such categorical encodings is that many

TABLE 1 Example of a one-hot encoding.

(a) mapping		(b) Stream of values
Value	One-hot	
		000 <u>1</u> 0000 (France)
		000 <u>1</u> 0000 (France)
Australia	<u>1</u> 0000000	00000 <u>1</u> 00 (Portugal)
Canada	0 <u>1</u> 000000	000 <u>1</u> 0000 (France)
China	00 <u>1</u> 00000	0000000 <u>1</u> (USA)
France	000 <u>1</u> 0000	00000 <u>1</u> 00 (Portugal)
Japan	0000 <u>1</u> 000	0000000 <u>1</u> (USA)
Portugal	00000 <u>1</u> 00	0000000 <u>1</u> (USA)
Spain	000000 <u>1</u> 0	0000000 <u>1</u> (USA)
USA	0000000 <u>1</u>	00 <u>1</u> 00000 (China)

**FIGURE 1** Comparing per-word population count (popcnt) with positional population count (pospopcnt).

machine-learning algorithms require numerical variables and cannot operate directly on categorical data. These encodings are closely related to the concept of dummy variables in statistics where categorical attributes are represented as zeros or ones for computing purposes. In this context, we want to compute the frequency histogram, that is the number of occurrences of each value. For example, given a stream (Table 1b) of one-hot encoded values (Table 1a), we want to compute the number of times each country has been observed such that France has been observed three times, Portugal two times, USA four times, and China one time. To compute such cumulative frequency histograms from encoded words, we need to count the number of set bits at each position (first, second, ..., last).

We name this novel generalization of the population count operation to individual bits spanning multiple words as the *positional population count* (pospopcnt). If each word is made of k -bits then we want to compute k counts representing the total number of set bits at each of the k positions. As an illustrative example, consider three given input words A , B , and C (Fig. 1), the conventional per-word population count computes the number of set bits for each word independently (Fig. 1a). In contrast, the positional population count operation computes the *vertical* population count over these words for independent bit positions (Fig. 1b).

Our main contribution is the formalized description of the novel positional population count operation and the introduction of an efficient algorithmic family for computing it using SIMD instructions leveraging the AVX-512 instruction set architecture (ISA) available on recent Intel processors. The AVX-512 family of ISAs can operate on up to thirty-two 512-bit registers. Notably, we show that our proposed algorithm can be in excess of 50-fold faster compared to a baseline scalar approach, as long as the input is sufficiently large (several kilobytes).

2 | RELATED WORK

To our knowledge, there is no prior work in the literature describing the positional population count. In contrast, computation of the population count has enjoyed renewed interest¹² following recent advancements in hardware with wider registers and

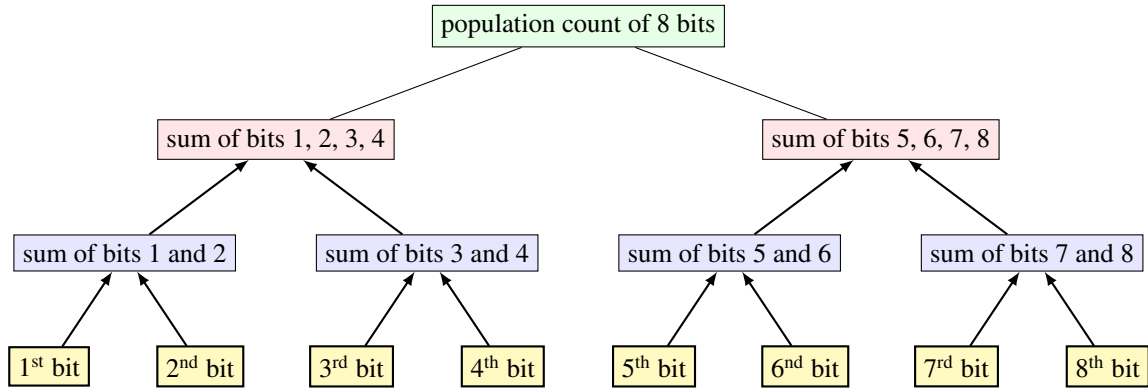


FIGURE 2 A tree of adders to compute the population count of eight bits in three steps.

```

void CSA(uint64_t* h, uint64_t* l,
         uint64_t a, uint64_t b, uint64_t c) {
    uint64_t u = a xor b;
    *h = (a bitand b) bitor (u bitand c);
    *l = u xor c;
}

```

FIGURE 3 A C++ function implementing a bitwise parallel carry-save adder (CSA). Given three input words a, b, c , we compute h, l representing the most significant and the least significant bits of the sum of the bits from a, b , and c , respectively.

more diverse and mature instruction-set architectures. Computing the population count is of such general importance that most commodity processors have dedicated instructions (e.g. `popcnt` on x64 processors). These hardware instructions are efficient when operating on single words. When large volumes of machine words are available in contiguous memory, other software-based approaches have demonstrated superior performance¹².

One of the earliest efficient algorithm to compute the population count is the Wilkes-Wheeler-Gill population-count¹⁵ (Fig. 2). This algorithm begins by adding adjacent bits into two-bit subwords, then adds these two-bit subwords into four-bit subwords, and finally adds adjacent four-bit subwords into byte values. These byte values are then summed using a multiplication and a bit-shift. It has a natural tree structure and can be described as a "tree of adders"¹⁶.

Another efficient approach¹⁷ is the Harley-Seal algorithm that is based on carry-save adder (CSA) networks frequently used in hardware microarchitectures. The CSA subroutine takes three input words and produce two output words: (1) a sequence of partial sum bits, and (2) a sequence of carry bits. Given three bit values ($a, b, c \in \{0, 1\}$), the sum $a + b + c$ is computed as a 2-bit word where the least significant bit is given by $(a \text{ xor } b) \text{ xor } c$ and the most significant bit is given by $(a \text{ bitand } b) \text{ bitor } ((a \text{ xor } b) \text{ bitand } c)$. We reuse the $a \text{ xor } b$ expression from the computation of the least significant bit when computing the most significant bit. Therefore, three values can be summed using 5 logical operations. This CSA network is particularly effective when generalized to operate on 64-bits in parallel (Fig. 3). Let us illustrate this approach by computing the population count of four inputs words using a simplified circuit (Fig. 4). Starting with two 64-bit words initialized to zero: one for the least significant bits (B_0) and one for the second least significant bits (B_1).

1. Two input words are summed with B_0 using a CSA subroutine (Fig. 3) resulting in two outputs: one corresponding to the least significant bits stored in B_0 and another corresponding to the second most significant bit: let us call it x .
2. Two new input words are summed with B_0 using the same CSA subroutine. Again the output corresponding to the least significant bit is stored in B_0 . We have a second output corresponding to the second most significant bits, let us call it y .
3. In the final step, x, y , and B_1 are added together. The output corresponding to the least significant bits is stored in B_1 . Similarly, the output corresponding to the most significant bit is stored in B_2 .

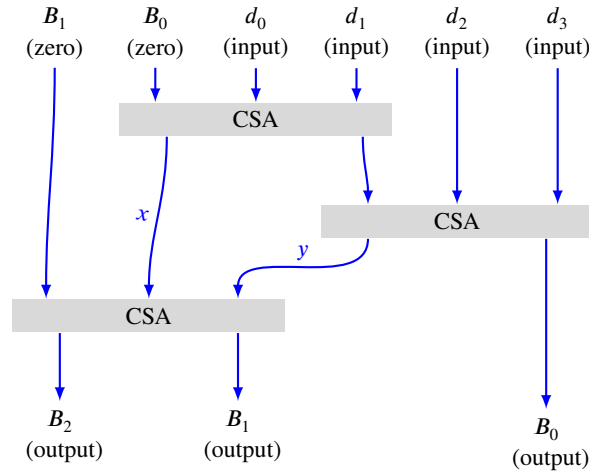


FIGURE 4 CSA circuit algorithm aggregating four inputs (d_0, d_1, d_2, d_3), producing three outputs B_0, B_1, B_2 corresponding to the least significant bit, second most significant bit and most significant bit of the sum. Observe that B_0 and B_1 are both inputs and outputs.

If c_i is the population count of B_i , then the population count of the four words is $1c_0 + 2c_1 + 4c_2$. In general, given 2^N input words, we can produce $N + 1$ words corresponding to the sum of the bits using a CSA circuit. Then the per-word population count can be computed from the resulting N words. More complex circuits are difficult to illustrate but can be built by recursively applying the same circuit. Assume that we have a circuit that takes 2^N inputs as well as B_0, \dots, B_{N-1} and produces the updated B_0, \dots, B_N corresponding to the sum. We can represent this circuit as a function $F_N(d_0, \dots, d_{2^N-1}; B_0, \dots, B_{N-1})$ that outputs B_0, \dots, B_N . We can then construct a larger circuit that takes twice as many inputs by two calls to the circuit together with an additional CSA routine. Indeed, to compute $F_{N+1}(d_0, \dots, d_{2^{N+1}-1}; B_0, \dots, B_N)$, we can start by computing $F_N(d_0, \dots, d_{2^N-1}; B_0, \dots, B_{N-1})$. We provide the first N outputs (B_0, \dots, B_N) from this first call to F_N to a second call to the circuit F_N , this time using the 2^N following inputs $d_N, \dots, d_{2^{N+1}-1}$. We are left with three values B_N , one was provided as an input to F_{N+1} , and two have been produced by a call to the smaller circuit F_N . We can combine these with a CSA routine. By such a recursive argument, we can show that $2^N - 1$ CSA routines is sufficient to process 2^N inputs.

In addition to being scaled to more inputs, the CSA circuit can be parallelized using wide SIMD registers (e.g. 512-bit words using the AVX-512 ISA). Instead of processing 64-bit words one-by-one, we can process multiple 64-bit words in parallel resulting in greater speed. Given sufficiently large input arrays (e.g., 1 kB), such a SIMD-accelerated CSA network is superior to all known population count alternatives on current Intel processors¹².

Furthermore, our general strategy and many of our constructions are familiar to cryptographers¹⁸. Some of these researchers have produced related optimizations using advanced SIMD instructions^{19,20}.

3 | CARRY-SAVE-ADDED (CSA) ROUTINES IN AVX-512

The conventional approach for computing the CSA routine requires 5 logical operations. On processors with the AVX-512 instruction set architecture available, we can simplify this subroutine down to two instructions by using the three-operand instruction `vpternlogd` using the `_mm512_ternarylogic_epi32` intrinsic function (Fig. 5). These compiler-specific intrinsic functions enables access to specific processor instructions without having to directly write in assembly or machine language. The `vpternlogd` instruction operates on three input bits a, b, c and an 8-bit sequence i and returns the bit at index $1a + 2b + 4c$ of i . Given three 512-bit input registers, the instructions generates a 512-bit output, executing 512-bit operations in parallel. The `vpternlogd` instruction can compute any ternary Boolean function using a single instruction.

This instruction is useful in computing the most and least significant bits in the CSA routine. To determine the value of the 8-bit sequence i in the context of computing the least and most significant bits, it suffices to enumerate all 8 possible inputs (Table 2).

TABLE 2 CSA routine in the context of the `vpternlogd` instruction

(a) least significant bit					(b) most significant bit				
c	b	a	$a \text{ xor } b \text{ xor } c$	$1a + 2b + 4c$	c	b	a	$(a \text{ and } b) \text{ or } ((a \text{ xor } b) \text{ and } c)$	$1a + 2b + 4c$
0	0	0	0	0	0	0	0	0	0
1	0	0	1	1	1	0	0	0	1
0	1	0	1	2	0	1	0	0	2
1	1	0	0	3	1	1	0	1	3
0	0	1	1	4	0	0	1	0	4
1	0	1	0	5	1	0	1	1	5
0	1	1	0	6	0	1	1	1	6
1	1	1	1	7	1	1	1	1	7

```

void CSA_AVX512(__m512i* h, __m512i* l,
  __m512i a, __m512i b, __m512i c)
{
  *h = _mm512_ternarylogic_epi32(c, b, a, 0b11101000);
  *l = _mm512_ternarylogic_epi32(c, b, a, 0b10010110);
}

```

FIGURE 5 The carry-save adder update step using AVX-512-based instructions. The three inputs are a, b, c , the two outputs are h and l corresponding respectively to the most significant and least significant bits.

- The least significant bit can be computed as the xor of the three inputs. Therefore, the `vpternlogd` instruction computes the desired output when the bits at index $\{1, 2, 4, 7\}$ of i are set. In software, this bit sequence is represented as the integer `0b10010110` in binary or 150 in decimal.
- Computing the most significant bits involves the expression $(a \text{ and } b) \text{ or } ((a \text{ xor } b) \text{ and } c)$. We can achieve this result by setting the bits at index $\{3, 5, 6, 7\}$ of i that can be represented as the binary string `0b11101000` or 232 in decimal.

When comparing the resulting AVX-512 routine to its 64-bit equivalent (Fig. 3), we observe that the AVX-512-based CSA subroutine can process 8 times as many bits (512 versus 64) while simultaneously reducing the number of logical operations from 5 to 2. Taken together, we achieve a gain factor of $5/2 \times 8 = 20$ in terms of instructions per input bits.

4 | POSITIONAL POPULATION-COUNT ALGORITHMS

Implementations of the positional population count depends on the desired target word size. For simplicity, we focus on 16-bit words in this work. Wider words (larger universes) can be encoded into multiple distinct streams of 16-bit words. In addition, our algorithms can be generalized to wider words without difficulty.

To compute the positional population count, we start with sixteen zero-initialized bit-counters. Next, we iterate over the bits in each word in sequence while incrementing the corresponding target bit-counter as needed. A sensible baseline algorithm for computing the positional population count involves a shift-mask-add subroutine (Fig. 6):

1. Right shift a 16-bit word by the bit index $p \in \{0, 1, \dots, 15\}$. Mathematically, this is equivalent to an integer division by 2^p with the result being the integer quotient.
2. Mask out all but the least significant bit from the shifted word. This bit is set when the quotient is odd, and zero otherwise.

```

void pospopcnt_u16_scalar_basic(uint16_t* data,
    uint32_t len, uint32_t* counters) {
    for (int i = 0; i < len; ++i) {
        // Each bit in every input word.
        uint16_t w = data[i];
        counters[0] += ((w >> 0) & 1);
        counters[1] += ((w >> 1) & 1);
        ...
        counters[15] += ((w >> 15) & 1);
    }
}

```

FIGURE 6 Reference algorithm for computing the positional population count. Given some input data we compute the total number of set bits at each position by using a branchless mask-shift-add update step.

3. Increment the counter at position p with the value from step 2.

We expect this code to be compiled to efficient assembly. On superscalar processors, branches must be predicted and a branch misprediction can cost tens of cycles or more. Fortunately, in our tests, the executed code does not trigger a significant number of mispredicted branches: the end of the short loop (16) is corrected predicted.

For sufficiently long input streams, we estimate the number of executed instructions to ~ 4 per bit for a total of ~ 64 instructions per 16-bit word. Given that most recent Intel processors can retire four instructions per CPU cycle, we expect a processing speed of around 16 cycles per 16-bit word for this scalar shift-mask-add subroutine (Fig. 6).

4.1 | AVX-512

We build CSA circuits over 512-bit registers using a CSA routine (Fig. 5). The simplest network (Fig. 4) takes four 512-bit registers (or 256 B) as inputs (Fig. 7) and uses three calls to the `CSA_AVX512` function. We can double the size of the network and process 512 B inputs using seven calls to the `CSA_AVX512` function. Again, doubling the network size, we can process 1 kB inputs using fifteen calls to the `CSA_AVX512` function (Fig. 8). Each call to the `CSA_AVX512` function involves only two instructions (two times `vpternlogd`). Excluding load and store instructions, we use (1) 6 instructions per block of 256 B or (2) 14 instructions per block of 512 B or (3) 30 instructions per block of 1 kB. Notably, the number of instructions per input volume increases with wider circuits.

At the end of each circuit, we have a set of registers $\{B_0, B_1, B_2, \dots\}$ corresponding, for each bit position in $i \in [0, 1, \dots, 512)$, to the sum of set bits. For simplicity let us assume that we are working with the smaller circuit (256 B). Before starting the procedure, the registers B_0, B_1, B_2 and the 16 output bit-counters are zero-initialized. Next, the first 512 B of inputs are processed with the circuit and receive as output the updated values of B_0, B_1, B_2 . In addition, we also receive a new register B_3 corresponding to the fourth least significant bit of the sum of the set bits at position $i \in [0, 1, \dots, 512)$. We can update our sixteen bit-counters by checking the bits in B_3 : the first counter needs to be incremented by eight times the sum of the bits at indexes $\{0, 16, 32, \dots, 496\}$, the second counter needs to be incremented by eight times the sum of the bits at indexes $\{1, 17, 33, \dots, 497\}$, and so forth. The process of loading the next 512 B of input data, call the circuit, and providing it with the updated values B_0, B_1, B_2 is repeated until no more data is available. At this point, when no more input data is available, we must increment our bit-counters with the remaining bits in B_0, B_1, B_2 . These are processed as B_3 except that the multiplier is 1, 2 and 4 respectively. The algorithms using the large circuits (512 B and 1 kB) work similarly.

The inner loop of our algorithm, where the bit-counters from the B_3 register are updated could become expensive. This loop executes 16 sums of 32 bit values and multiplies the result by eight. This final multiplication can be deferred until the end of the main loop. Further, we vectorize the sum (Fig. 9) by using sixteen vector of counters, each spanning 512 bits, or thirty-two 16-bit counters, instead of using sixteen scalar counters. Given the B_3 register, we select the least significant bit of each 16-bit word by using a mask-select operation involving the `_mm512_and_si512` intrinsic function. Next, we increment counters using a horizontal add operation involving the `_mm512_add_epi16` intrinsic. Finally, we shift each 16-bit subword of the B_3 register right by one bit with the `_mm512_srli_epi16` intrinsic function. This procedure can be written in scalar form as:

```

// inputs: B0, B1
//         data[i], ..., data[i + 3]
//
// temporary variables: B1A, B1B
//
// outputs: B0, B1, B2
//
CSA_AVX512(&B1A , &B0 , B0 , data[i] , data[i + 1] );
CSA_AVX512(&B1B , &B0 , B0 , data[i + 2] , data[i + 3] );
CSA_AVX512(&B2 , &B1 , B1 , B1A , B1B );

```

FIGURE 7 CSA circuit for our AVX-512 implementation in C++ processing 4 registers or 256 B, the CSA_AVX512 is as in Fig. 5.

```

for (pos = 0; pos < 16; ++pos) {
    for (i = 0; i < 32; ++i) {
        counter[pos] += B[i] & 1;
        B[i] >>= 1;
    }
}

```

with the convention that the subword B at index i ($B[i]$) selects a 16-bit word within B . By unrolling this vectorized code, we achieve $3 \times 16 = 48$ instructions, without considering overhead and load/store instructions. We can add these 48 instructions to our previous counts: (1) 54 instructions per block of 256 B or (2) 62 instructions per block of 512 B or (3) 78 instructions per block of 1 kB. Dividing by the number of 16-bit words processed, we get the following instruction counts per word:

1. 0.42 when using blocks of 256 B,
2. 0.24 when using blocks of 512 B, and
3. 0.15 when using blocks of 1 kB.

Based on these numbers, as block sizes get larger, we expect fewer instructions per input word and therefore improved performance. In comparison, we estimate about 64 instructions per 16-bit word for the scalar version. This corresponds to a >400-fold reduction in the number of executed instructions.

This analysis ignores the overhead cost of recovering the counts from both the vectorized counters and the running registers $\{B0, B1, B2, \dots\}$. For small inputs, this overhead can occupy a dominant proportion of the total cost. For example, the sixteen vector counters span 1 kB on their own. Hence, to achieve good performance, the input stream should at least exceed 1 kB.

Taken together, we describe several SIMD-accelerated CSA networks of different sizes for computing the positional population count with up to >400-fold reduction in the number of executed instructions compared to a scalar implementation. Our SIMD-based approach can compute the positional population count for inputs that are a multiple of 256 B. Any residual input values will be processed using a scalar subroutine.

5 | EXPERIMENTS

We implemented the algorithms in C99 and we make them available online at https://github.com/lemire/pospopcnt_avx512 under the Apache 2.0 license. Code was compiled with GCC 8.2 using the optimization flags `-O2 -march=native`. All tests were performed using a host machine with a Cannon Lake microarchitecture (Table 3). Performance was measured using the Linux performance counters (e.g. `PERF_COUNT_HW_CPU_CYCLES` to count processor cycles). The host processor has 32 kB of L1 cache per core, 256 kB of L2 cache per core and 4 MB of L3 cache. We verified that the processor is not subject to downclocking when running AVX-512 instructions²¹ using the `avx-turbo`²² benchmarking tool.

For all experiments, we generate random data using a Mersenne Twister pseudorandom number generator. However, we find that performance is independent of the input data as our algorithms do not branch on the content of the data.


```

// inputs: B0, B1, B2, B2
//      data[i], ..., data[i + 15]
//
// temporary variables: B1A, B1B, B2A, B2B,
//      B2A, B2B
//
// outputs: B0, B1, B2, B2, B3
//
CSA_AVX512(&B1A , &B0 , B0 , data[i] , data[i + 1] );
CSA_AVX512(&B1B , &B0 , B0 , data[i + 2] , data[i + 3] );
CSA_AVX512(&B2A , &B1 , B1 , B1A , B1B );
CSA_AVX512(&B1A , &B0 , B0 , data[i + 4] , data[i + 5] );
CSA_AVX512(&B1B , &B0 , B0 , data[i + 6] , data[i + 7] );
CSA_AVX512(&B2B , &B1 , B1 , B1A , B1B );
CSA_AVX512(&B2A , &B2 , B2 , B2A , B2B );
CSA_AVX512(&B1A , &B0 , B0 , data[i + 8] , data[i + 9] );
CSA_AVX512(&B1B , &B0 , B0 , data[i + 10] , data[i + 11] );
CSA_AVX512(&B2A , &B1 , B1 , B1A , B1B );
CSA_AVX512(&B1A , &B0 , B0 , data[i + 12] , data[i + 13] );
CSA_AVX512(&B1B , &B0 , B0 , data[i + 14] , data[i + 15] );
CSA_AVX512(&B2B , &B1 , B1 , B1A , B1B );
CSA_AVX512(&B2B , &B2 , B2 , B2A , B2B );
CSA_AVX512(&B3 , &B2 , B2 , B2A , B2B );

```

FIGURE 8 CSA circuit for our AVX-512 implementation in C++ processing 16 registers or 1 kB, the CSA_AVX512 is as in Fig. 5.

```

// input: B
// counter[0], ..., counter[15] are 512-bit registers
__m512i one = _mm512_set1_epi16(1); // 1-mask (000...1)
for (pos = 0; pos < 16; ++pos) {
    __m512i masked = _mm512_and_si512(B, one); // Select LSB (bit) with 1-mask
    counter[pos] = _mm512_add_epi16(counter[pos], masked); // Horizontal add
    B = _mm512_srli_epi16(B, 1); // Shift 16-bit words right by 1 bit
}

```

FIGURE 9 Vectorized counter increment.

Our proposed AVX-512-based approach has a fixed overhead, irrespective of the input stream size. This overhead is non-negligible for small inputs. For tiny inputs (512 words), we need about 30 instructions per 16-bit word, and almost all of these instructions are part of the overhead. However, once we reach a few thousands words (e.g., 4096 words or 8 kB), our AVX-512-based approach becomes beneficial. We simulated random input data for 8 and 512 kB, as well as 16, 64, 256, and 1024 MB and benchmarked our three different CSA circuits against a scalar implementation (Table 4). To guarantee reliability, we repeat each test three times and verify that the standard error is well within a 5%. We also present the speeds in GB/s for input sizes up to 64 MB (see Fig. 10). For the scalar algorithm, we observe a fixed speed over the range of inputs in our experiments: 17 cycles and 65 instructions per input word. Overall, the fastest approach is our proposed AVX-512 algorithm with large block size (1 kB), except maybe for inputs of less than 64 kB. Increasing the input size reduce the number of instructions per input word and simultaneously increase speed. Due to diminishing returns, the performance for 1 GB inputs are practically identical to that of 256 MB. In the best scenario, the AVX-512 approach uses 300 times fewer instructions compared to the scalar approach.

TABLE 3 Hardware

Processor	Base Frequency	Max. Frequency	L1 data cache per core	Microarchitecture	Compiler
Intel i3-8121U	2.2 GHz	3.2 GHz	32 kB	Cannon Lake (2018)	GCC 8.2

TABLE 4 Performance results for different input sizes: CPU cycles per 16-bit word, instructions per 16-bit word and speed in GB/s.

	(a) 8 kB input			(b) 512 kB input		
	Cycles/word	Ins./word	Speed (GB/s)	Cycles/word	Ins./word	Speed (GB/s)
Scalar	17	65	0.22	17	65	0.36
AVX-512 (256 B)	1.2	3.9	0.44	0.22	0.57	15
AVX-512 (512 B)	1.5	4.7	0.45	0.15	0.39	18
AVX-512 (1 kB)	1.5	4.7	0.45	0.14	0.29	19
	(c) 16 MB input			(d) 64 MB input		
	Cycles/word	Ins./word	Speed (GB/s)	Cycles/word	Ins./word	Speed (GB/s)
Scalar	17	65	0.36	17	65	0.36
AVX-512 (256 B)	0.65	0.52	9.7	0.65	0.52	12
AVX-512 (512 B)	0.55	0.32	11	0.44	0.32	14
AVX-512 (1 kB)	0.51	0.23	12	0.41	0.23	15
	(e) 256 MB input			(f) 1 GB input		
	Cycles/word	Ins./word	Speed (GB/s)	Cycles/word	Ins./word	Speed (GB/s)
Scalar	17	65	0.36	17	65	0.36
AVX-512 (256 B)	0.47	0.52	14	0.47	0.52	14
AVX-512 (512 B)	0.39	0.32	16	0.39	0.32	16
AVX-512 (1 kB)	0.36	0.22	18	0.36	0.22	18

As a reference point, the C `memcpy` function runs at a speed of 20 GB/s for large inputs compared to 18 GB/s for the positional `popcount` operation on large inputs.

6 | CONCLUSION

Taken together, we introduce a novel population count operation referred to as the *positional population count* (`pospopcnt`) and introduce efficient SIMD-accelerated algorithms based on carry-saver add (CSA) circuits. We demonstrate that we can process 16-bit words at nearly the speed of a memory copy when AVX-512 is available and when the input size is sufficiently large (e.g., 256 MB). Our best algorithm is over 40 times faster compared to a non-vectorized approach.




FIGURE 10 Processing speed (in GB/s) for various input size. Our Intel-based system has 4 MB of L3 cache.

Acknowledgements

This work was funded by a Wellcome Ph.D. studentship grant 109082/Z/15/A (M.D.R.K.) as well as by a grant (RGPIN-2017-03910) from the Natural Sciences and Engineering (D.L.). We are grateful to J. D. McCalpin (University of Texas at Austin) for benchmarking advice. We are grateful to members of the open-source community for reviewing and improving the implementation.

References

1. Dai B, Guo R, Kumar S, He N, Song L. Stochastic Generative Hashing. In: Precup D, Teh YW., eds. *Proceedings of the 34th International Conference on Machine Learning*. 70 of *Proceedings of Machine Learning Research*. PMLR. ; 2017: 913–922.
2. Courbariaux M, Hubara I, Soudry D, El-Yaniv R, Bengio Y. Binarized Neural Networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. arXiv preprint arXiv:1602.02830; 2016.
3. Zhang D, Yang J, Ye D, Hua G. LQ-Nets: Learned Quantization for Highly Accurate and Compact Deep Neural Networks. In: Ferrari V, Hebert M, Sminchisescu C, Weiss Y., eds. *Computer Vision – ECCV 2018* Springer International Publishing. ; 2018; Cham: 373–390.
4. Hubara I, Courbariaux M, Soudry D, El-Yaniv R, Bengio Y. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research* 2017; 18(1): 6869–6898.
5. Lemire D, Ssi-Yan-Kai G, Kaser O. Consistently faster and smaller compressed bitmaps with Roaring. *Software: Practice and Experience* 2016; 46(11): 1547–1569.
6. Haque IS, Pande VS, Walters WP. Anatomy of High-Performance 2D Similarity Calculations. *Journal of Chemical Information and Modeling* 2011; 51(9): 2345–2351. doi: 10.1021/ci200235e
7. Sanyal A, Kusner M, Gascon A, Kanade V. TAPAS: Tricks to Accelerate (encrypted) Prediction As a Service. In: PMLR. ; 2018: 4497–4506.
8. Layer RM, Kindlon N, Karczewski KJ, Quinlan AR, Quinlan AR. Efficient genotype compression and analysis of large genetic-variation data sets. *Nature Methods* 2016; 13(1): 63–65. doi: 10.1038/nmeth.3654

9. Danek A, Deorowicz S. GTC: how to maintain huge genotype collections in a compressed form. *Bioinformatics* 2018; 34(11): 1834–1840. doi: 10.1093/bioinformatics/bty023
10. Wu TD, Nacu S. Fast and SNP-tolerant detection of complex variants and splicing in short reads. *Bioinformatics* 2010; 26(7): 873–881. doi: 10.1093/bioinformatics/btq057
11. Purcell S, Neale B, Todd-Brown K, et al. PLINK: A Tool Set for Whole-Genome Association and Population-Based Linkage Analyses. *The American Journal of Human Genetics* 2007; 81(3): 559–575. doi: 10.1086/519795
12. Muła W, Kurz N, Lemire D. Faster population counts using AVX2 instructions. *The Computer Journal* 2017; 61(1): 111–120.
13. Lippert C, Listgarten J, Davidson RI, et al. An exhaustive epistatic SNP association analysis on expanded Wellcome Trust data. *Scientific reports* 2013; 3: 1099.
14. Mittag F, Römer M, Zell A. Influence of feature encoding and choice of classifier on disease risk prediction in genome-wide association studies. *PloS one* 2015; 10(8): e0135832.
15. Wilkes MV, Wheeler DJ, Gill S. *The Preparation of Programs for an Electronic Digital Computer*. Boston, USA: Addison-Wesley Publishing. second ed. 1957.
16. Knuth DE. *The Art of Computer Programming: volume 4A: combinatorial algorithms, part 1*. Boston, Massachusetts: Addison-Wesley . 2011.
17. Warren HS. The quest for an accelerated population count. In: Wilson G, Oram A. , eds. *Beautiful Code: Leading Programmers Explain How They Think*O'Reilly Media, Sebastopol, California: chapter 10. 2007 (pp. 147-160).
18. Chou T. QcBits: constant-time small-key code-based cryptography. In: Springer. ; 2016: 280–300.
19. Guimarães A, Aranha DF, Borin E. Optimized implementation of QC-MDPC code-based cryptography. *Concurrency and Computation: Practice and Experience* 2019; 31(18): e5089.
20. Drucker N, Gueron S, Kostic D. QC-MDPC decoders with several shades of gray. In: Springer. ; 2020: 35–50.
21. Intel 64 and IA-32 Architectures Optimization Reference Manual. Tech. Rep. 248966-033, Intel Corporation; Santa Clara, California, U.S.: 2016.
22. Downs T. avx-turbo: Test the non-AVX, AVX2 and AVX-512 speeds across various active core counts. <https://github.com/travisdowns/avx-turbo>; 2019.

