SPECIAL ISSUE PAPER

A Sparse Matrix-Vector Multiplication Method with Low Preprocessing Cost

Barış Aktemur*

¹Department of Computer Science, Ozyegin University, Turkey

Correspondence

*Barış Aktemur, Ozyegin University, Cekmekoy, 34794 Istanbul, Turkey. Email: baris.aktemur@ozyegin.edu.tr

Abstract

Sparse matrix-vector multiplication (SpMV) is a crucial operation used for solving many engineering and scientific problems. In general there is no single SpMV method that gives high performance for all sparse matrices. Even though there exist sparse matrix storage formats and SpMV implementations that yield high efficiency for certain matrix structures, using these methods may entail high preprocessing or format conversion costs. In this work, we present a new SpMV implementation, named CSRLenGoto, that can be utilized by preprocessing the Compressed Sparse Row (CSR) format of a matrix. This preprocessing phase is inexpensive enough for the associated cost to be compensated in just a few repetitions of the SpMV operation. CSRLen-Goto is based on complete loop unrolling, and gives performance improvements in particular for matrices whose mean row length is low. We parallelized our method by integrating it into a stateof-the-art matrix partitioning approach as the kernel operation. We observed up to 2.46x and on the average 1.29x speedup with respect to Intel MKL's SpMV function for matrices with short or medium-length rows.

KEYWORDS:

sparse matrix-vector multiplication; spmv; compressed sparse row

1 | INTRODUCTION

Sparse matrices are those matrices that contain a high ratio of zeros. Sparse matrix-vector multiplication (SpMV) is a fundamental operation used frequently in the solutions to many engineering and science problems. In domains such as Krylov subspace problems and iterative solvers, a sparse matrix goes into an SpMV operation tens or hundreds of times. For this reason, SpMV plays a key role in high performance computing. However, it is known that SpMV's performance falls well behind the capacity of modern computers¹. Hence, optimization of SpMV has been extensively studied (see Langr and Tvrdík², and Filippone et al.³ for comprehensive surveys).

Sparse matrices are stored in formats that provide space savings. Probably the most popular and the *de facto* standard of these formats is the *Compressed Sparse Row* (CSR) representation⁴ (explained in detail in the next section). CSR is a general-purpose sparse matrix storage format, taking no parameters such as a block size or the cache line length, and oblivious to architectural details of the target computer. However, it cannot give the best performance for all sparse matrix types. For this reason, new storage formats are developed to address performance shortcomings², sometimes targetting only a limited class of sparse matrices. When a new format is proposed, there is a crucial cost that needs to be taken into account: the cost of converting the matrix data from CSR format to the new format. If this conversion is too costly with respect to the SpMV operation, the format would find very limited use in practice. Furthermore, in addition to the time overhead of format conversion, one has to consider space overheads, too, because if the matrix is going to be used in other operations after SpMV, the CSR representation cannot be thrown away; the matrix data



FIGURE 1 A sample matrix, its CSR representation, and the SpMV function for the CSR format.

has to be kept in both the CSR format and the new format. In other words, there may be a need to keep two copies of a matrix in memory. For these reasons, there has been a recent interest to use the CSR format directly or after an inexpensive preprocessing step ^{5–10}.

In this work, we investigate improving the efficiency of SpMV based on complete loop unrolling and using a variation of CSR as the storage format. The experimental evaluation of our approach shows that significant speedups can be obtained, particularly for matrices that have short and medium-length rows, by paying a low preprocessing cost. The **contributions** made by our work are as follows:

- We present a variation of the CSR matrix storage format, named CSRLen.
- We present an SpMV implementation, named CSRLenGoto, using the CSRLen format and based on complete loop unrolling. CSRLenGoto gives substantial speedups with respect to CSR. Our current implementation is for the X64_64 CPU architecture.
- Matrix data in CSR format can be converted to the *CSRLen* format very quickly (on average, in duration equivalent to 0.11–0.25× of the baseline SpMV execution). This way, the preprocessing cost can be amortized in only a few iterations of the SpMV for most of the matrices.
- CSRLenGoto gives speedup for matrices that have low or medium mean row length, and in particular for matrices with short rows (e.g. mean row length smaller than 8).
- Our method can be parallelized straightforwardly using existing matrix partitioning methods. As an example, we integrated our method as the kernel method of the merge-based SpMV approach⁵ and obtained up to 2.46× performance with respect to Intel's Math Kernel Library (MKL) on an 8-core Intel Xeon CPU.
- Our code is publicly available at https://github.com/aktemur/CSRLenGoto.

This paper is organized as follows: Section 2 gives background information about the SpMV problem. Section 3 presents our proposed approach. Section 4 and 5 evaluate the performance in single-threaded and multi-threaded settings, respectively. Section 6 compares and contrasts our approach to the existing work. Finally, Section 7 gives our conclusions.

2 | BACKGROUND

In the CSR (*Compressed Sparse Row*) format, a matrix is represented by three arrays that we will call vals, cols, and rows. In the vals array, the nonzero elements of the matrix are stored according to the row-major ordering. The cols array stores the column indices of the nonzero elements, in the same order as the vals array. The rows array stores, for each row, the index of the first element of the row in the other two arrays. A sample matrix, its CSR representation, and the SpMV code for the CSR format are given in Figure 1. In the code, N is the number of rows of the matrix; v is the input vector; w is the output vector. SpMV calculates the following expression: $w \leftarrow w + M \cdot v$.

SpMV CSR code is notorious for having poor performance. There are several reasons behind this. SpMV is a memory-bound computation ¹¹. The accesses to the input vector v are indirect and irregular; this causes underutilization of the cache and poor instruction-level parallelism ^{1,12}. Because sparse matrices usually have short rows, the trip count of the inner loop is low. For this reason, loop-related costs stay relatively high, and the branch-predictor of the CPU does not help much ¹³. As an attempt to remedy these performance problems, we can try **loop unrolling** – one of the standard optimization transformations performed by compilers. For instance, if we unroll the inner loop in Figure 1 for 4 times, we obtain the code in Figure 2. Here, we need a second inner loop to handle the left-over elements in case the number of elements in the row is not an exact multiple of

```
for (int i = 0; i < N; i++) {
   double sum = 0.0;
   int j = rows[i];
   for (; j < rows[i+1] - 3; j += 4) {
      sum += vals[j] * v[cols[j]];
      sum += vals[j+1] * v[cols[j+1]];
      sum += vals[j+2] * v[cols[j+2]];
      sum += vals[j+3] * v[cols[j+3]];
   }
   for (; j < rows[i+1]; j++)
      sum += vals[j] * v[cols[j]];
   w[i] += sum;
}</pre>
```



```
// Expected matrix format: CSR
      int j = 0, i = 0, length;
      double sum;
      goto init;
L_5: sum += vals[j] * v[cols[j]]; j++;
L_4: sum += vals[j] * v[cols[j]]; j++;
L_3: sum += vals[j] * v[cols[j]]; j++;
L_2: sum += vals[j] * v[cols[j]]; j++;
L_1: sum += vals[j] * v[cols[j]]; j++;
L_0: w[i] += sum;
      i++;
init: if (i >= N) goto end;
      length = rows[i + 1] - rows[i];
      sum = 0.0;
      goto L_length; // conceptual
end:
     :
```

FIGURE 3 CSRGoto code when the maximum row length of the matrix is 5.

the unrolling factor 4. We will refer to the code obtained by unrolling the inner loop of the CSR code k times as CSR_k . So, the original code given in Figure 1 is CSR_1 . The necessity for a second inner loop in CSR_k brings runtime overhead especially for short rows seen frequently in sparse matrices, and CSR_k does not yield the desired performance increase.

3 | OUR PROPOSED APPROACH

In this section we present *CSRLenGoto* – the new SpMV method we are proposing; but before we do that, we will discuss an intermediary method that we name *CSRGoto*.

3.1 | CSRGoto

Let us investigate unrolling the inner loop of *CSR*₁ completely instead of k times, so that we can get rid of the need to have a second inner loop to handle the left-over elements. Because the maximum iteration count of the inner loop of *CSR*₁ will be equal to the maximum row length of the matrix, we just need to unroll the loop as many times as the number of elements in the longest row. After unrolling, for each row, we have to be able to jump to the appropriate point in the code according to the length of the row. We can do that by putting labels in the code and using a goto statement. Let us call this SpMV method *CSRGoto*. We show a *conceptual* code for this method in Figure 3. Note that for this method to work correctly, it is not required that the rows of the matrix have a fixed length; row lengths may vary, but the original inner loop should be unrolled for at least the maximum row length of the matrix.

```
xor %eax, %eax
1
                                           ; j \leftarrow 0
                                           ; i \leftarrow 0
            xor %edx, %edx
2
3
            movsxd (%r11), %rcx
                                           ; rcx \leftarrow rows[0]
            jmp init
4
5
     L_5: ; sum += vals[j] * v[cols[j]]; j++;
           movslq (%r9,%rax,4), %rbx
6
7
            movsd (%r8,%rax,8), %xmm1
            incq %rax
8
9
            mulsd (%rdi,%rbx,8), %xmm1
            addsd %xmm1, %xmm0
10
11
     L_4: \ldots; same as L_5
     L_3: ... ; same as L_5
12
13
     L_2: ... ; same as L_5
14
      L 1: ... : same as L 5
15
      L_0: addsd (%rsi,%rdx,8), %xmm0 ; sum \leftarrow sum + w[i]
            movsd %xmm0, (%rsi,%rdx,8) ; w[i] \leftarrow sum
16
                                             ; i \leftarrow i + 1
17
            incq %rdx
      init:
18
19
            cmp %rdx, N
                                            ; exit if i \geq N
            jge end
20
            movslq 4(%r11,%rdx,4),%rbx ; rbx \leftarrow rows[i+1]
21
                                            ; \mathit{rcx} \leftarrow \mathit{-length}
22
            suba %rbx, %rcx
23
            imul 22, %rcx
                                            ; rcx \leftarrow -length * 22
24
            leaq -45(%rip), %r10
                                            ; r10 \leftarrow L_0
25
            addq %rcx, %r10
                                            ; r10 \leftarrow L_0 + rcx
                                            ; rcx \leftarrow rbx
            leag (%rbx), %rcx
26
                                             ; sum \leftarrow 0
27
            xorps %xmm0, %xmm0
28
            jmp *%r10
                                             ; goto L_length
29
      end:
```

FIGURE 4 X86_64 assembly code for CSRGoto when the maximum row length in the matrix is 5.

We said the code in Figure 3 is "conceptual" because the last line is not a legal C statement. Here, we have a programming challenge. How can we express in C the address that we want to jump to? We can use the label addressing operator (&&) and the *computed goto* statement¹ that do not exist in the C standard but are supported by compilers like GCC and Clang:

long delta = (&&L_0 - &&L_5) / 5; //distance between each consecutive label
goto *(void*)(&&L_0 - length * delta);

However, for this code to be correct, we have to trust the compiler to position the labels equidistantly. In our experiments, we have seen that this is not a valid assumption. Therefore, we decided to implement the SpMV function at the assembly code level, instead of source code. To determine which machine instructions to use, we compiled source codes similar to the one given in Figure 3 using the Clang, GCC, and icc compilers to produce X86_64 assembly code (with the -03 optimization level). Upon manual examination of the output of the compilers, we have seen that the statement

sum += vals[j] * v[cols[j]]; j++;

is transformed to native code similar to

movslq (%r9,%rax,4), %rbx	;	$rbx \leftarrow cols[j]$	(4 bytes)
movsd (%r8,%rax,8), %xmm1	;	x mm1 \leftarrow vals[j]	(6 bytes)
incq %rax	;	j++	(3 bytes)
mulsd (%rdi,%rbx,8), %xmm1	;	$xmm1 \leftarrow xmm1 * v[rbx]$	(5 bytes)
addsd %xmm1, %xmm0	;	$\mathit{sum} \leftarrow \mathit{sum} + \mathit{xmm1}$	(4 bytes)

Based on this observation, we write for CSRGoto the assembly code shown in Figure 4. A straightforward manual examination of the instructions above reveals that their total length is 22 bytes (4+6+3+5+4). Therefore, the distance between two consecutive labels L_j and $L_(j + 1)$ will be 22 bytes. Because this is a constant value, no runtime calculation is needed. We use this fact on line 23 in Figure 4.

3.2 | CSRLenGoto

In the CSRGoto method, substantial amount of work is performed to calculate the address to jump to (Figure 4, between the init and end labels). This calculation may be a burden especially for short rows. To remedy, we can pre-compute the work that does not depend on the program counter

```
int *newRows = new int[N + 1];
for (int i = 0; i < N; i++) {
    int length = rows[i + 1] - rows[i];
    newRows[i] = -length * 22;
}
newRows[N] = 33; // distance between 'L_0' and 'end' in Figure 6</pre>
```

FIGURE 5 Converting the CSR format to CSRLen.

```
; j \leftarrow 0
     xor %eax, %eax
     xor %edx, %edx
                                     : i \leftarrow 0
     jmp init
L_5: ; sum += vals[j] * v[cols[j]]; j++;
     movslq (%r9,%rax,4), %rbx
     movsd (%r8,%rax,8), %xmm1
     incq %rax
     mulsd (%rdi,%rbx,8), %xmm1
     addsd %xmm1, %xmm0
L_4: \ldots; same as L_5
L 3: ... : same as L 5
L_2: ... ; same as L_5
L 1: ... : same as L 5
L_0: addsd (%rsi,%rdx,8), %xmm0 ; sum \leftarrow sum + w[i]
                                                             (5 bytes)
     movsd %xmm0, (%rsi,%rdx,8) ; w[i] \leftarrow sum
                                                             (5 bytes)
     incq %rdx
                                     ; i \leftarrow i + 1
                                                             (3 bytes)
init:
     xorps %xmm0, %xmm0
                                    ; sum \leftarrow 0
                                                             (3 bytes)
     movslq (%r11,%rdx,4), %rbx ; rbx \leftarrow rows[i]
                                                             (4 bytes)
     leaq -27(%rip), %r10
                                    ; r10 \leftarrow L_0
                                                             (7 bytes)
     addq %rbx, %r10
                                    ; r10 \leftarrow r10 + rbx
                                                             (3 bytes)
     jmp *%r10
                                     ; goto L_length
                                                             (3 bytes)
end
```

FIGURE 6 X86_64 assembly code for CSRLenGoto when the maximum row length in the matrix is 5.

%rip, such as the calculation of length values, and keep those values as part of the matrix data. For this, we need to go over the rows array of the matrix in a preprocessing phase. We name the matrix data obtained in this way, CSRLen, and the SpMV method using this matrix data, CSRLenGoto. Converting the data from the CSR format to CSRLen involves processing the rows array only, and is given in Figure 5. Here, we calculate for each row the distance to land on the label appropriate for the row length. Now that we are preparing a new rows array, we can as well get rid of the check for the exit condition (the cmp and jge instructions). For this, we make an addendum to the very end of the new rows array, and store the distance between the labels L_0 and end, which is 33 bytes. This value is the sum of the lengths of instructions that appear between these two labels (individual instruction lengths are shown in Figure 6).

In Figure 5 we create a new array named newRows. If the user does not need to keep around the original CSR data, the original rows array can be overwritten instead of creating a separate array. In our experiments, to be fair in the measurements, we created a new array. The runtime complexity of the conversion is O(N); the vals and cols arrays are not modified or processed. Because the size of the newRows array is the same as the original one, the size of the matrix data transferred from the memory to the SpMV function will be the same as CSR.

CSRLenGoto code is given in Figure 6. This code is quite similar to the *CSRGoto* code in Figure 4. The only major difference is that the distance to jump is read from the rows array instead of being computed on the fly according to the row length. Remember that the source-level unrolling code in Figure 2 contains one outer loop and two inner loops. The associated overheads prevent performance benefits, especially for short rows. The overheads are minimized in our proposed approach where there exists only one loop, obtained via a jump instruction, whose target is dynamically computed for each row. Runtime computation is further eliminated through the preprocessing phase. When code is unrolled at the source-level, the compiler is provided with opportunities for re-ordering the instructions. We lose this feature because we implemented our approach in assembly; however, that is a compromise we had to make to guarantee that the code sections are equally-spaced.

4 | SINGLE-THREADED PERFORMANCE EVALUATION

In this section we measure and discuss the performance of CSRLenGoto in a single-threaded execution environment.



FIGURE 7 Information about the 1878 matrices in our data set that we used for performance evaluation.

CDI I (Micro-architecture)	Cache size (Bytes)			Memory	Compiler
CFO (Micro-architecture)	L1 (I/D)	L2	L3	(GB)	Complier
Intel Xeon E5-1660v4 @ 3.20 Ghz, 8-core, 16 threads (Broadwell)	8×(32K/32K)	8×256K	20M	32	icc 17.0.0
AMD FX 8350 @ 4.00 Ghz, 8-core, 8 threads (Piledriver)	4×64K/8×16K	4×2M	8M	8	gcc 5.4.0

TABLE 1 Machines used in our experiments.

4.1 | Setup

6

We prepared a data set comprising of real-world matrices obtained from the SuiteSparse (formerly known as the University of Florida) matrix collection ¹⁴. Our set contains non-complex-valued matrices that have 10K–200M nonzero elements in the SuiteSparse collection. There are 1878 such matrices. Information about the number of rows, nonzero elements, and row lengths of this set are given in Figure 7.

We ran the tests on two X86_64 computers, one with an Intel CPU, the other with an AMD CPU. The properties of the machines are given in Table 1. The codes were compiled by passing the -03 flag² to the compiler. Our testbeds were kept unloaded during benchmarking to minimize interference with other processes. When measuring the SpMV duration, we repeatedly invoked the function in a loop and recorded the total elapsed time to avoid possible noise. We set the number of repetitions to a sufficiently large value to keep the total elapsed time reasonably long (e.g. \sim 1s, or longer). We divided the elapsed time to the number of repetitions to find the duration taken by a single SpMV invocation. This way, we measured the SpMV time for each matrix for 3 times, and recorded the smallest (i.e. fastest execution). We used the double-precision floating point type (64-bit double) for the nonzero values, and integer type (32-bit int) for row and column indices. For the *CSR*_k method, we used the following values for k: {1, 4, 8, 16, 32}. For *CSRLenGoto*, in addition to the SpMV duration, we also measured the preprocessing time where CSR format is converted to CSRLen. We do not include the cost of preprocessing as part of SpMV time, because in a typical "inspector-executor" setting, first, an inspection phase is performed for optimizations, followed by the execution phase where the optimized SpMV function is executed many times in an iteration. Intel's Math Kernel Library (MKL) provides an inspector-executor API³ as well. We compare our library against MKL in our multi-threaded performance evaluation (Section 5). We do not evaluate against MKL in this section, because MKL is tuned for parallelism and does not perform well in sequential execution.

On each testbed machine, we set the best-performing CSR_k as the baseline method. We determined the best-performing CSR_k as follows: For every matrix, we normalized each CSR_k result with respect to the fastest of the CSR_k methods for that matrix. We observed that on Intel, CSR_1 is almost always the best CSR_k method. Hence, we picked CSR_1 as the baseline on Intel. On AMD, CSR_1 was 8%, CSR_4 was 4%, CSR_8 was 5%, CSR_1 was 8%, and CSR_{32} was 13% worse than the fastest CSR_k when averaged over all the matrices. Hence, we set CSR_4 as the baseline method on AMD.

We express the performance of the *CSRLenGoto* method with respect to the performance of the baseline method. For this, we divide the duration taken by the baseline method to the time taken by *CSRLenGoto*. Having a ratio smaller than 1 means that *CSRLenGoto* caused a slowdown; a ratio larger than 1 means that we obtained speedup.

²Vectorization is enabled by default with this flag; however, SpMV implementation for CSR does not benefit from vectorization because of the indirect access to the input vector v.

³https://software.intel.com/en-us/mkl-developer-reference-c-inspector-executor-sparse-blas-routines



FIGURE 8 CSRLenGoto's performance with respect to the baseline method in the single-threaded setting on Intel (left) and AMD (right).



FIGURE 9 Percentage of matrices that give speedup when a certain mean row length value is used as (a) an upper bound to define "matrices with short rows", (b) a lower bound to define "matrices with long rows".

4.2 | Results

When examining the results, we noticed that the performance of *CSRLenGoto* is related to the mean row lengths of the matrices. Therefore, we are presenting the performance results by associating them with mean row lengths. Figure 8 gives the results. We can see on both machines that as the mean row length increases, performance drops down. This observation suggests that we can consider the matrices to be in one of three categories based on their mean row length and the general trend in performance: For matrices with *short rows*, substantial speedups are obtained for the majority. The performance for the matrices with *medium-length rows* is mixed, with results both below and above the baseline. Finally, hardly any speedup is achieved for matrices with *long rows*. This observation naturally raises the next question: Which row length values should we use to draw the boundaries between these three categories? We will soon discuss *where* exactly we can cut the data set, but let us first discuss *why* such a categorization helps.

An SpMV library needs to make runtime decisions about which SpMV method to use among possibly many options for a particular matrix. If the matrix can be put into a category for which a particular SpMV method is known to consistently give good performance, the library can pick that method and proceed with the SpMV operation. The categorization of the matrix is to be made based on install-time information learned on the target machine and a dynamic analysis of the matrix. This is a typical *auto-tuning* approach that has been successfully applied in the case of SpMV¹⁵⁻²⁰. Ideally, the dynamic analysis of the matrix should be as inexpensive as possible. In the case of *CSRLenGoto*, one can quickly check whether the given matrix falls into the category of "short rows" or "long rows". This is just a matter of calculating the mean row length of the matrix via a single division operation (i.e. NZ/N) and comparison. If the mean row length is low, the library can decide with high confidence that *CSRLenGoto* will yield speedups and perform the SpMV operation using *CSRLenGoto*. Similarly, if the mean row length is high, the library can predict that *CSRLenGoto* is unlikely

7

to perform better than the baseline, and do not attempt using it. When the mean row length is inbetween, the library may prefer to perform further analysis based on matrix features other than NZ/N (e.g. max row length, variation of row lengths, etc.), or monitor the first few executions of *CSRLenGoto* to decide whether to fall back to the default method or continue using it.

The next question is, what threshold values should we use for the mean row length so a matrix can be categorized as having short/medium/long rows? Figure 8 suggests that these ranges should be machine-specific. Figure 9(a) depicts an analysis of what percentage of matrices give speedup when a certain mean row length value is used as an upper bound to define "matrices with short rows". For instance, 90% of the matrices whose mean row length is 8.0 or less (767 matrices) on Intel, and 27.5 or less (1360 matrices) on AMD give speedup when using *CSRLenGoto*. Similarly, Figure 9(b) gives an analysis of what percentage of matrices give *slowdown* when a certain mean row length value is used as a *lower* bound to define "matrices with long rows". For instance, among the matrices whose mean row length is 14.5 or more (738 matrices), 90% give slowdown on Intel. So, we can say with high confidence that *CSRLenGoto* will not yield any benefits if the mean row length of a matrix is bigger than 14.5. The situation is better on AMD; for example, among the matrices whose mean row length is 40.0 or more (372 matrices), 70% give slowdown. If we use these boundaries to define the range for matrices with medium-length rows (i.e. 8.0–14.5 on Intel, 27.5–40.0 on AMD), then the percentage of matrices that show speedup is 56% on Intel (208 out of 373), and 62% on AMD (90 out of 146). In general, an analysis like this one can be carried on a machine-basis to evaluate the likelihood of *CSRLenGoto* giving speedup for a matrix by just looking at the NZ/N value of the matrix. If an autotuning approach is taken, using a decision tree classifier is potentially an effective method for finding the thresholds. In the rest of the paper, we will use the threshold values 8.0 and 14.5 on Intel, and 27.5 and 40.0 on AMD, when categorizing matrices as having short/medium/long rows.

It is hard to say precisely why *CSRLenGoto* shows performance relative to the mean row length value, but a reasonable explanation is due to the utilization of the cache. Modern processors include a small cache, called the *micro-op cache*, with the purpose of optimizing the instruction decode time especially for loops with small bodies. For *CSRLenGoto*, as the matrix row length increases, the "loop body" is going to be long, and this diminishes the advantages of the micro-op cache. When the row length is further longer, we begin losing the benefits of the level-1 instruction cache as well. This is why, we believe, *CSRLenGoto* provides better speedup for matrices with shorter rows. The reason we are seeing different speedups on our test machines is likely related to the difference in their cache sizes. The AMD CPU has larger level-1 instruction cache; hence, better speedups were observed on that CPU for matrices with high row lengths.

4.3 | Preprocessing Cost

Recall that to use *CSRLenGoto*, a matrix has to go through a preprocessing phase to convert the matrix data from CSR format to CSRLen. We have measured the cost of conversion relative to the cost of one SpMV function execution using the baseline method. This way, we are attempting to answer the question *"How many SpMV operations could we have executed using the baseline method instead of doing preprocessing for CSRLenGoto?"* The answer to this question is given in the top row in Figure 10. The relative cost of preprocessing tends to be higher for matrices with short rows compared to the matrices with long rows. Yet, the costs are very small, considering that in an expensive SpMV method, preprocessing may take time equivalent to tens or hundreds of SpMV invocations ^{6,19,21–23}. On the average, CSRLen preprocessing time is equivalent to 0.25 SpMV executions on Intel and 0.11 on AMD. The largest cost we have measured is 1.56 on Intel and 0.59 on AMD.

The next question we ask is "How many times do we have to repeat SpMV for a matrix so that CSRLenGoto compensates its preprocessing cost and becomes quicker than the baseline method? This question makes sense especially in the context of iterative solvers, where the same matrix is multiplied with some vector many times until a convergence point is reached. In this context, we can afford to pay a preprocessing cost, and if the SpMV operation is repeated for a sufficient number of iterations, we get profits in the overall time spent. The number of iterations needed to start getting profit is called the *break-even point*. The bottom row in Figure 10 shows the break-even points for matrices for which the measured performance is at least $1.01 \times$. There are 961 such matrices on Intel, and 1418 on AMD. Most of the break-even points are in the range of just a couple SpMV iterations. The average break-even points are 4.7 and 1.2 for Intel and AMD, respectively; the maximums are 42.0 and 18.5 for Intel and AMD, respectively. This shows that *CSRLenGoto* can be used in iterative solver contexts even when the iteration count is low.

4.4 Comparison to Yzelman's Sparse Library

Yzelman provides an SpMV benchmark software, called *Sparse Library*, that contains implementation of several SpMV methods²⁴. We ran the sequential methods in Sparse Library and measured the performance in order to compare *CSRLenGoto* to a method other than the baseline CSR. We evaluated the following methods: incremental CSR (ICSR); zig-zag CSR (ZZ-CSR); zig-zag incremental CSR (ZZ-ICSR); sparse vector matrix (SVM); Hilbert-ordered triplet schema (HTS); bi-directional incremental CSR (BICSR); hilbert-ordered triplets, backed by BICSR (Hilbert); sparse matrix blocking, backed by Hilbert (Block Hilbert); sparse matrix blocking by bisection, backed by Hilbert (Bisection Hilbert); compressed bi-directional incremental CSR (CBICSR). We again used the 1878-matrix set. We performed the measurements on our Intel machine only.



FIGURE 10 Top row: The cost of preprocessing normalized with respect to the cost of one SpMV operation using the baseline method on Intel (left) and AMD (right). Bottom row: Break-even points of matrices for which the performance is at least $1.01 \times$ on Intel (left) and AMD (right).

Method	Wins	Method	Wins
ICSR	44 (2.3%)	BICSR	9 (0.5%)
ZZ-CSR	788 (42.0%)	Hilbert	48 (2.6%)
ZZ-ICSR	39 (2.1%)	Block Hilbert	1 (0.1%)
SVM	0 (0.0%)	Bisection Hilbert	39 (2.1%)
HTS	51 (2.7%)	CBICSR	12 (0.6%)
Baseline CSR	847 (45.1%)		



FIGURE 11 The SpMV methods we evaluated from Yzelman's Sparse Library and the number of matrices for which they gave the best performance (left); the performance of *CSRLenGoto* with respect to ZZ-CSR (right).

The results are given in Figure 11, where the table on the left lists the evaluated SpMV mehods and the number of matrices for which the method was the best performer among the methods listed in the table. In 847 (45.1%) of the matrices, no method was better than the baseline CSR. In other cases, ZZ-CSR dominates the table with a count of 788 (42.0%) matrices. For this reason, we picked ZZ-CSR as the method to compare *CSRLenGoto* against. The result is given in the right hand side in Figure 11. Here, we see that *CSRLenGoto* again performs very well for matrices with short rows. It should be noted that the zig-zag CSR method needs to reorder the nonzero values and the column indices of the CSR format in its preprocessing phase; *CSRLenGoto*, however, uses these two arrays verbatim.

5 | MULTI-THREADED PERFORMANCE EVALUATION

In this section we evaluate *CSRLenGoto*'s performance in a multi-threaded environment. SpMV is a highly parallelizable computation; matrix rows can be processed concurrently without having to do synchronization among rows. A typical parallelization approach is to partition the matrix into



FIGURE 12 Performance of parallelized CSRLenGoto with respect to MKL (left), and merge-based SpMV (right).

horizontal "stripes", consisting of rows. Each stripe is then assigned to a thread. In such a row-based decomposition, the elements of a row is never split among different partitions. This decomposition may fail to divide the work equally among threads if the matrix is highly irregular; that is, if there are unusually long or many empty rows in the matrix. Merrill and Garland's recent *merge-based* method ⁵ provides strict load balancing between threads by allowing a row at a partition border to be split. Their approach is based on the "merge-path" algorithm, which is an efficient parallelization of the merge operation of two sorted lists, A and B, where each thread is assigned an equal share of the total |A| + |B| steps. The merge-path algorithm is applied to the problem of partitioning CSR matrix data by logically merging the rows array, which is sorted, with the natural numbers from 0 to NZ (i.e. indices of the cols and vals arrays). Once the matrix is partitioned, each partition can be processed sequentially by a thread. We implemented a multi-threaded version of *CSRLenGoto* by integrating it into the *merge-based SpMV* code base²⁵ as the kernel function that is executed sequentially by a thread. In principle, it should be possible to use *CSRLenGoto* with any CSR-based matrix partitioning approach for concurrent execution. We preferred merge-based SpMV because it is state-of-the-art and its code is public.

In our experiment we used the same data set that contains 1878 matrices with 10K-200M non-complex nonzero values (Figure 7). Our setup is the same as explained in the single-threaded performance evaluation section except the following: We ran the experiment only on our Intel computer (see Table 1 for properties) because the merge-based SpMV code base requires the Intel compiler (icc), and because we wanted to compare parallel *CSRLenGoto* against Intel's Math Kernel Library (MKL) as well. Concurrent execution is obtained via OpenMP pragmas. We ran three methods in this benchmarking:

- MKL's inspector-executor interface: In the inspector phase, we first create an MKL-internal CSR matrix viamkl_sparse_d_create_csr, then call mkl_sparse_set_mv_hint and mkl_sparse_optimize. As the SpMV function, we use mkl_sparse_d_mv.
- Merge-based SpMV: In the preprocessing phase, the matrix is partitioned according to the merge-path method. Plain CSR-based SpMV function is executed on each partition. This is Merrill and Garland's original code.
- CSRLenGoto: In the preprocessing phase, the matrix is partitioned according to the merge-path method, then the CSR data is converted to CSRLen. CSRLenGoto's SpMV function is executed on each partition.

For all methods, we again timed the preprocessing/inspector phases separately from the SpMV phase. We used 16 threads, which is the number of hardware threads available on the CPU (8 cores, 2 hyper-threads per core). CSR-CSRLen conversion is straightforwardly parallelized via OpenMP compiler directives.

Performance of *CSRLenGoto* with respect to MKL and merge-based methods are given in Figure 12. The overall speedup trend related to the mean row length seems to apply in the multi-threaded context as well. On the average, *CSRLenGoto*'s performance is $1.34 \times$ of MKL and $1.33 \times$ of merge-based SpMV among matrices with short rows (i.e. mean row length 8.0 or less). For matrices with medium-length rows (i.e. mean row length between 8.0 and 14.5), the average performances are $1.19 \times$ and $1.22 \times$ for MKL and merge-based SpMV, respectively. The maximum performance obtained is $2.46 \times$ wrt MKL, and $2.33 \times$ wrt merge-based SpMV. Similar to the sequential execution, slow-down cases dominate when the matrices have long rows. Preprocessing costs with respect to one MKL SpMV function call are shown in Figure 13. On the average, preprocessing takes time equivalent to 0.52, 0.52, and 0.34 SpMV calls for matrices for which the performance is less than $1.01 \times$ (744 matrices). For the remaining 1134 matrices, the average break-even point is 5.08 SpMV iterations.



FIGURE 13 Preprocessing costs (left), and break-even points (right) of parallelized CSRLenGoto with respect to one MKL SpMV cost.

6 | RELATED WORK

Since SpMV is used many scientific problems, improving its performance implies a wide impact. For this reason, optimizing SpMV has been extensively studied by many researchers (e.g. ^{1,5,15,21-23,26-28}). Two recent work provide a detailed survey ^{2,3}.

Our motivation in this work has been to use a format similar to CSR so as to avoid heavy preprocessing/inspection costs. Reordering the matrix elements and using a storage format other than CSR may give substantial speedup for some matrices ^{13,21–23,28}. However, the cost of preprocessing and storage format conversion in these approach can be equivalent to tens or even hundreds of SpMV iterations ^{6,19,21–23}. Compensating these costs can be possible only if the optimized SpMV routine will be called hundreds or thousands of times. Yet, it is common for the number SpMV iterations in iterative solvers to be less than a hundred. In such context, SpMV optimization approaches that have high preprocessing costs do not give benefits. For this reason, some recent work have focused on optimizing SpMV with little or no preprocessing cost on the GPU ^{6,7,9,10}, CPU ⁸, or both ⁵. Among those for the GPU, Ashari et al. propose to group matrix rows of the same length together to speed up SpMV ⁶. Greathouse and Daga dynamically distribute CSR matrix data to GPU threads, and efficiently use GPU's scratchpad memory to obtain improvements ⁷. Liu and Schmidt apply no preprocessing to the CSR matrix data in their *LightSpMV* library; they achieve speedup by using GPU's atomic operations, warp shuffle mechanism, and dynamic load balancing ¹⁰. In another GPU-oriented method, Liu et al. propose the LSRB-CSR format, where they decompose the matrix data into segments tuned according to the architectural properties and the warp length of the target GPU ⁹. Ohshima et al. investigate the effect of various OpenMP thread scheduling settings on the CSR-based SpMV performance on CPU ⁸. Merrill and Garland propose a row-oriented decomposition technique for the CSR format based on the *merge-path* approach ⁵. Their technique promises better load balancing for parallel execution of SpMV on both CPU and GPU, in particular for matrices with irregular shapes.

SpMV is a highly parallelizable computation because the matrix rows can be processed independent from each other. The usual approach for parallelization is to decompose the matrix to as many pieces as the number of threads available on the target platform. There exist row-oriented, one-dimensional decomposition approaches ^{5,28,29} as well as two-dimensional ones ³⁰⁻³³. After the matrix data is decomposed, each portion is processed using a sequential kernel. Therefore, even though we essentially proposed an improvement of the sequential execution of SpMV, our method can be combined with existing matrix partitioning methods for parallelization. As an example, we integrated our method as the kernel SpMV implementation in Merrill and Garland's state of the art merge-based SpMV approach and showed that our method can provide substantial speed improvement in a multi-threaded setting (see Section 5).

Our method is based on the fundamental idea of loop unrolling. We had previously shown that loop unrolling provides speedup for SpMV³⁴, but we had done that for formats that re-arrange matrix data. Another paper that investigates loop unrolling for SpMV is by Mellor-Crummey and Garvin¹³. Here, a sparse matrix representation called LCSR that reorders matrix elements is used. For this reason, converting CSR format to LCSR involves restructuring both the vals and cols arrays. The conversion from CSR to our CSRLen format does not process these two arrays; we rebuild the rows array only (see Figure 5). *CSRLenGoto* method can be modified to work with a reordered matrix; however, this does not fit the "low-overhead preprocessing" motivation we had put forward. Similarly, in principle it is possible to reduce memory traffic by combining *CSRLenGoto* with data compression approaches such as value compression in the CSR-VI format²⁷ or other data compression techniques³⁵; however, compressing the values or column indices requires a preprocessing phase over the vals and cols arrays, and may incur additional costs such as frequently accessing a hashtable. Therefore, we did not prefer incorporating these compression techniques. Kumahata et al. ³⁶ employ loop unrolling to optimize the SpMV portion of the high performance conjugate gradient (HPCG) problem on the K computer. Their approach is based on generating an SpMV loop

for each possible row length value. In our approach, it is possible to have a single SpMV function. That single function can be used for any matrix, as long as the maximum row length of the matrix is smaller than the maximum row length assumed when implementing our code.

7 | CONCLUSION

We proposed a new SpMV implementation, named *CSRLenGoto*, that relies on complete loop unrolling and computed jump instructions based on row lengths. Our method gives substantial speedup for matrices whose mean row length is not high; for these matrices we measured up to $2.46 \times$ and on the average $1.29 \times$ speedup with respect to Intel's MKL. *CSRLenGoto* uses a variation of the CSR sparse matrix storage format, which can be built via a low-overhead preprocessing phase. On the average, the cost of preprocessing is equivalent to half of the time that would be spent in a single execution of the SpMV function. This way, *CSRLenGoto* is able to pay off its preprocessing cost and start bringing speed benefits in only 5.08 iterations of the SpMV function on the average.

References

- Goumas Georgios I., Kourtis Kornilios, Anastopoulos Nikos, Karakasis Vasileios, Koziris Nectarios. Understanding the Performance of Sparse Matrix-Vector Multiplication. In: :283–292; 2008.
- 2. Langr Daniel, Tvrdik Pavel. Evaluation Criteria for Sparse Matrix Storage Formats. IEEE Trans. Parallel Distrib. Syst.. 2016;27(2):428-440.
- Filippone Salvatore, Cardellini Valeria, Barbieri Davide, Fanfarillo Alessandro. Sparse Matrix-Vector Multiplication on GPGPUs. ACM Trans. Math. Softw. 2017;43(4):30:1–30:49.
- 4. Saad Y.. Iterative Methods for Sparse Linear Systems. SIAM; 2003.
- Merrill Duane, Garland Michael. Merge-based Parallel Sparse Matrix-vector Multiplication. In: SC '16:58:1–58:12IEEE Press; 2016; Piscataway, NJ, USA.
- Ashari Arash, Sedaghati Naser, Eisenlohr John, Parthasarathy Srinivasan, Sadayappan P., Fast Sparse Matrix-vector Multiplication on GPUs for Graph Applications. In: SC '14:781–792IEEE Press; 2014; Piscataway, NJ, USA.
- Greathouse Joseph L., Daga Mayank. Efficient Sparse Matrix-vector Multiplication on GPUs Using the CSR Storage Format. In: SC '14:769– 780IEEE Press; 2014; Piscataway, NJ, USA.
- Ohshima Satoshi, Katagiri Takahiro, Matsumoto Masaharu. Performance Optimization of SpMV Using CRS Format by Considering OpenMP Scheduling on CPUs and MIC. In: MCSOC '14:253–260IEEE Computer Society; 2014; Washington, DC, USA.
- 9. Liu L., Liu M., Wang C., Wang J., LSRB-CSR: A Low Overhead Storage Format for SpMV on the GPU Systems. In: :733-741; 2015.
- 10. Liu Yongchao, Schmidt Bertil. LightSpMV: Faster CUDA-Compatible Sparse Matrix-Vector Multiplication Using Compressed Sparse Rows. Journal of Signal Processing Systems. 2017;.
- Gropp W., Kaushik D., Keyes D., Smith B.. Toward realistic performance bounds for implicit CFD codes. In: Ecer A., others, eds. Proceedings of Parallel CFD'99, Elsevier; 1999.
- 12. Pichel J., Heras D., Cabaleiro J., Rivera F.. Improving the locality of the sparse matrix-vector product on shared memory multiprocessors. In: :66-71; 2004.
- Mellor-Crummey John, Garvin John. Optimizing Sparse Matrix-Vector Product Computations Using Unroll and Jam. Int. J. High Perform. Comput. Appl.. 2004;18(2):225–236.
- 14. Davis Timothy A., Hu Yifan. The University of Florida Sparse Matrix Collection. ACM Trans. Math. Softw.. 2011;38(1):1:1-1:25.
- 15. Vuduc Richard, Demmel James W, Yelick Katherine A. OSKI: A library of automatically tuned sparse matrix kernels. *Journal of Physics: Conference Series*. 2005;16(1):521.

- Byun Jong-Ho, Lin Richard, Yelick Katherine A., Demmel James. Autotuning Sparse Matrix-Vector Multiplication for Multicore. UCB/EECS-2012-215: EECS Department, University of California, Berkeley; 2012.
- 17. Li Jiajia, Tan Guangming, Chen Mingyu, Sun Ninghui. SMAT: An Input Adaptive Auto-tuner for Sparse Matrix-vector Multiplication. In: PLDI '13:117–126ACM; 2013; New York, NY, USA.
- Sedaghati Naser, Mu Te, Pouchet Louis-Noel, Parthasarathy Srinivasan, Sadayappan P. Automatic Selection of Sparse Matrix Representation on GPUs. In: ICS '15:99–108ACM; 2015; New York, NY, USA.
- Yılmaz Buse, Aktemur Barış, Garzarán Maria J., Kamin Sam, Kiraç Furkan. Autotuning Runtime Specialization for Sparse Matrix-Vector Multiplication. ACM Trans. Archit. Code Optim. 2016;13(1):5:1–5:26.
- 20. Elafrou A., Goumas G., Koziris N., Performance Analysis and Optimization of Sparse Matrix-Vector Multiplication on Modern Multi- and Many-Core Processors. In: :292-301; 2017.
- 21. Belgin Mehmet, Back Godmar, Ribbens Calvin J., A Library for Pattern-based Sparse Matrix Vector Multiply. International Journal of Parallel Programming. 2011;39(1):62-87.
- 22. Karakasis Vasileios, Gkountouvas Theodoros, Kourtis Kornilios, Goumas Georgios, Koziris Nectarios. An Extended Compression Format for the Optimization of Sparse Matrix-Vector Multiplication. *IEEE Trans. Parallel Distrib. Syst.*. 2013;24(10):1930–1940.
- Liu Weifeng, Vinter Brian. CSR5: An Efficient Storage Format for Cross-Platform Sparse Matrix-Vector Multiplication. In: ICS '15:339– 350ACM; 2015; New York, NY, USA.
- 24. Yzelman A.N.. Sparse Library http://albert-jan.yzelman.net/software.php.Accessed Nov. 2017; .
- 25. Merrill Duane. Merge-based Parallel Sparse Matrix-Vector Multiplication https://github.com/dumerrill/merge-spmv.Accessed Nov. 2017; .
- 26. Buluç Aydin, Fineman Jeremy T., Frigo Matteo, Gilbert John R., Leiserson Charles E., Parallel Sparse Matrix-vector and Matrix-transpose-vector Multiplication Using Compressed Sparse Blocks. In: SPAA '09:233–244ACM; 2009; New York, NY, USA.
- Kourtis Kornilios, Goumas Georgios, Koziris Nectarios. Exploiting Compression Opportunities to Improve SpMxV Performance on Shared Memory Systems. ACM Trans. Archit. Code Optim.. 2010;7(3):16:1–16:31.
- Williams Samuel, Oliker Leonid, Vuduc Richard, Shalf John, Yelick Katherine, Demmel James. Optimization of Sparse Matrix-vector Multiplication on Emerging Multicore Platforms. *Parallel Comput.* 2009;35(3):178–194.
- 29. Yang W., Li K., Mo Z., Li K.. Performance Optimization Using Partitioned SpMV on GPUs and Multicore CPUs. *IEEE Transactions on Computers*. 2015;64(9):2623-2636.
- Catalyurek Umit, Aykanat Cevdet. Hypergraph-Partitioning-Based Decomposition for Parallel Sparse-Matrix Vector Multiplication. IEEE Trans. Parallel Distrib. Syst. 1999;10(7):673–693.
- Martone Michele. Efficient multithreaded untransposed, transposed or symmetric sparse matrix-vector multiplication with the Recursive Sparse Blocks format. *Parallel Computing*. 2014;40(7):251 - 270.
- Yzelman A. N., Bisseling Rob H., Cache-Oblivious Sparse Matrix-Vector Multiplication by Using Sparse Matrix Partitioning Methods. SIAM Journal on Scientific Computing. 2009;31(4):3128-3154.
- Yzelman A. N., Roose D., High-Level Strategies for Parallel Shared-Memory Sparse Matrix-Vector Multiplication. IEEE Transactions on Parallel and Distributed Systems. 2014;25(1):116-125.
- Kamin Sam, Garzarán Maria J., Aktemur Barış, Xu Danqing, Yılmaz Buse, Chen Zhongbo. Optimization by Runtime Specialization for Sparse Matrix-vector Multiplication. In: GPCE '14:93–102; 2014.
- 35. Willcock Jeremiah, Lumsdaine Andrew. Accelerating Sparse Matrix Computations via Data Compression. In: ICS '06:307–316ACM; 2006; New York, NY, USA.
- Kumahata Kiyoshi, Minami Kazuo, Maruyama Naoya. High-performance Conjugate Gradient Performance Improvement on the K Computer. Int. J. High Perform. Comput. Appl.. 2016;30(1):55–70.

AUTHOR BIOGRAPHY



Barış Aktemur is an assistant professor of Computer Science at Ozyegin University. He received his PhD and MSc degrees from the University of Illinois at Urbana-Champaign, and his BSc degree from Bilkent University, all in computer science.