

# Diagonally-Addressed Matrix Nicknack: How to improve SpMV performance

Jens Saak <sup>†</sup>    Jonas Schulze <sup>†\*</sup>

<sup>†</sup> *Max Planck Institute for Dynamics of Complex Technical Systems,  
Sandtorstr. 1, 39106 Magdeburg, Germany.*

<sup>\*</sup> *Corresponding author. Email: [jschulze@mpi-magdeburg.mpg.de](mailto:jschulze@mpi-magdeburg.mpg.de)*

**Abstract:** We suggest a technique to reduce the storage size of sparse matrices at no loss of information. We call this technique Diagonally-Adressed (DA) storage. It exploits the typically low matrix bandwidth of matrices arising in applications. For memory-bound algorithms, this traffic reduction has direct benefits for both uni-precision and multi-precision algorithms.

In particular, we demonstrate how to apply DA storage to the Compressed Sparse Rows (CSR) format and compare the performance in computing the Sparse Matrix Vector (SpMV) product, which is a basic building block of many iterative algorithms. We investigate 1367 matrices from the SuiteSparse Matrix Collection [4] fitting into the CSR format using signed 32 bit indices. More than 95 % of these matrices fit into the DA-CSR format using 16 bit column indices, potentially after Reverse Cuthill-McKee (RCM) reordering [2]. Using IEEE 754 double precision scalars, we observe a performance uplift of 11 % (single-threaded) or 17.5 % (multithreaded) on average when the traffic exceeds the size of the last-level CPU cache. The predicted uplift in this scenario is 20 %. For traffic within the CPU's combined level 2 and level 3 caches, the multithreaded performance uplift is over 40 % for a few test matrices.

**Keywords:** sparse linear algebra, sparse matrices, traffic reduction, DAMN

**Novelty statement:** We introduce Diagonally-Adressed (DA) storage as a general technique to reduce sparse matrix traffic. Applying the DA variant of the CSR format to the Sparse Matrix Vector (SpMV) product, we demonstrate that our practical implementation achieves 87.5 % of the theoretically predicted performance uplift.

## 1 Introduction

An operation is called memory-bound, if its performance is limited by the memory bandwidth [Byte/s] of the executing hardware. In that context, it holds that

$$\frac{P_1}{P_2} = \frac{\text{traffic}_2}{\text{traffic}_1}, \quad (1)$$

where  $P_i$  denotes performance [FLOP/s], and  $\text{traffic}_i$  [Byte] accounts for all the memory involved. Hence, reducing the traffic should directly lead to a performance improvement. It is often assumed that computing the Sparse Matrix Vector (SpMV) product is memory-bound; see, e.g., [5, 7, 9]. The biggest contributor to the overall SpMV traffic is the matrix. Therefore, in the following, we present a technique that allows for the reduction of the storage size of a sparse matrix at no loss of information.

The key ingredient of the new storage technique is the observation that many matrices arising in, e.g., finite element simulations have a very low matrix bandwidth (under certain permutations).




Figure 1: Sample matrix (left) in Compressed Sparse Rows (CSR) storage (middle) and Diagonally-Adressed CSR (DA-CSR) storage (right). Colorful dots represent non-zero entries, gray dots are zero. Whiskers represent (column) indices with respect to a reference line (dashed).

Listing 1: (DA-) CSR storage of a matrix.

```

1 struct {
2     oindex_t rowptr [nrows + 1];
3     iindex_t colids [nnz];
4     scalar_t values [nnz];
5 };

```

That means, potentially after permutation, all non-zero matrix entries are located close to the matrix diagonal. This motivates storing the indices of these entries relative to the matrix diagonal rather than as an absolute position, which we call *Diagonally-Adressed (DA)* storage. Due to the small matrix bandwidth, the relative indices may be stored in a smaller (integer) data type. This technique is easily applicable to many sparse formats, e.g., Compressed Sparse Rows (CSR), Block CSR (BSR), or (one of the vectors of) Coordinate (COO) storage. In this paper we apply DA storage to the CSR format; obtaining the Diagonally-Adressed CSR (DA-CSR) format; and compare the SpMV performance against our implementation of CSR as well as Intel Math Kernel Library (MKL) [6].

DA storage differs from Diagonal (DIA) storage in that the new technique still requires one index per non-zero, depending on the underlying technique, instead of one index per diagonal. Also, it does not impose a diagonal-major order of the entries, or require a potentially padded and full/dense storage for each of the diagonals. The Recursive Sparse Blocks (RSB) format [10] is another data structure for sparse matrices motivated by a cache-efficient and parallel implementation of the SpMV product. It divides a sparse matrix into a tree structure of sparse blocks, whose leaves are iterated in a Z- or Morton-order. The leaf blocks are stored in the COO, CSR, or Compressed Sparse Columns (CSC) format.

RSB was designed for arbitrary sparse matrices, in particular, matrices without an inherent low bandwidth (under certain permutations). RSB allows 16 bit indices as well, but only for its leaf matrices. Meanwhile, DA-CSR has a conceptually simpler non-recursive design, allowing 16 bit indices throughout, which leads to a much lower overhead in terms of Byte per non-zero. Therefore, DA storage does not directly compete with the RSB format, but could be used in the leaf blocks within the RSB format, to allow for an even smaller index type.

The remainder of this paper is structured as follows. Section 2 applies DA storage to the CSR format. Section 3 describes the selection of test matrices. Section 4 describes how to compute the SpMV product using that new DA-CSR format, and measures the performance of SpMV. We conclude the paper in Section 5.

## 2 Diagonally-Adressed Storage

The CSR storage of a matrix  $A \in \mathbb{F}^{\text{nrows} \times \text{ncols}}$  comprises three vectors, cf. Listing 1, where `nnz` denotes the number of non-zero entries, `oindex_t` and `iindex_t` are integer data types, and

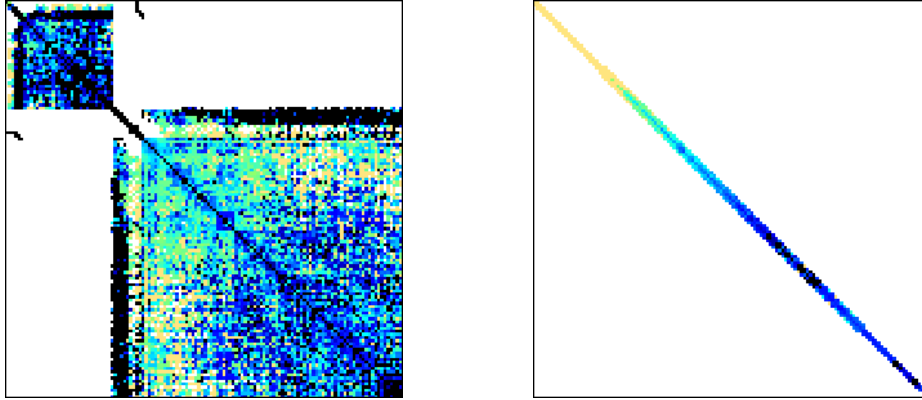


Figure 2: Sparsity patterns of the matrices GHS\_psdef/lldoor (left) as well as Janna/Bump\_2911 (right) from the SuiteSparse Matrix Collection [4].

`scalar_t` is an approximation of  $\mathbb{F}$ , e.g., IEEE 754 double or float for  $\mathbb{F} = \mathbb{R}$ . The “row pointers” stored in `rowptr` and “column indices” stored in `colids` do the bookkeeping imposed by only storing non-zero matrix entries.

The  $r$ th entry  $0 \leq \text{rowptr}[r] < \text{nnz}$  is the index into `colids` and `values` corresponding to the first non-zero of row  $r$ ,  $0 \leq r \leq \text{nrows}$ .<sup>1</sup> The  $i$ th entry `colids`[ $i$ ] is the column index and `values`[ $i$ ] is the value of the  $i$ th non-zero,  $0 \leq i < \text{nnz}$ . Let  $w$  denote the matrix bandwidth of  $A = (a_{rc})$ , i.e. the farthest distance of a non-zero matrix entry from the matrix diagonal,

$$w := \max\{|c - r| : a_{rc} \neq 0\} \ll \text{ncols}. \quad (2)$$

For CSR storage, `colids`[ $i$ ] =  $c_i$ , which lies in the range  $[0, \text{ncols}]$ . For DA-CSR storage, `colids`[ $i$ ] =  $c_i - r_i$ , which instead lies in the range  $[-w, w]$ . We illustrate this transformation in the following example.

**Example 1.** The sample matrix shown in Figure 1 has `nrows` = `ncols` = 4, `nnz` = 5, and  $w = 1$ . Its CSR representation is given by

$$\begin{cases} \text{rowptr} = (0, 1, 2, 4, 5) \\ \text{colids} = (0, 2, 1, 3, 3) = (\mathbf{1}, \mathbf{2}, \mathbf{1}, \mathbf{3}, \mathbf{3}) \\ \text{values} = (\mathbf{\bullet}, \mathbf{\bullet}, \mathbf{\bullet}, \mathbf{\bullet}, \mathbf{\bullet}) \end{cases} \quad (3)$$

while the DA storage replaces `colids` to become

$$\text{colids} = (0, 1, -1, 1, 0) = (\mathbf{1}, \mathbf{2}, \mathbf{-1}, \mathbf{3}, \mathbf{0}). \quad (4)$$

Observe that `colids` covers its full range in either storage scheme:  $[0, \text{ncols} - 1]$  for CSR and  $[-w, w]$  for DA-CSR.

Sometimes it is necessary to reduce the bandwidth of a matrix before DA storage can be applied effectively, which we observe in the next example.

**Example 2.** The matrix GHS\_posdef/lldoor from the SuiteSparse Matrix Collection [4] has dimension 952 203 and 46 522 475 pattern entries,<sup>2</sup> distributed over a bandwidth of 686 979. On average, this matrix has 49 pattern entries per row. See Figure 2 (left) for its sparsity pattern. Due to its block structure, the original matrix bandwidth is fairly large. Still, a Reverse Cuthill-McKee (RCM) reordering [2] reduces the bandwidth to about 9100,<sup>3</sup> which is only about 1% of the matrix dimension. Ignoring the colors, the corresponding sparsity pattern would look almost identical to Figure 2 (right). The reduced bandwidth is less than  $2^{15} = 32\,768$  and therefore allows for the usage of 16 bit column indices in DA storage, while both the matrix dimension (for plain CSR storage) and the original bandwidth (for naive DA-CSR storage) would require 32 bit indices.

<sup>1</sup>The final entry `rowptr`[`nrows`] is set to `nnz` for ease of use.

<sup>2</sup>Technically, `nnz` refers to the number of pattern entries, which contains non-zeros as well as explicitly stored zeros.

<sup>3</sup>Our implementation yields a bandwidth of 9120, while the SuiteSparse Matrix Collection [4] reports 9134.

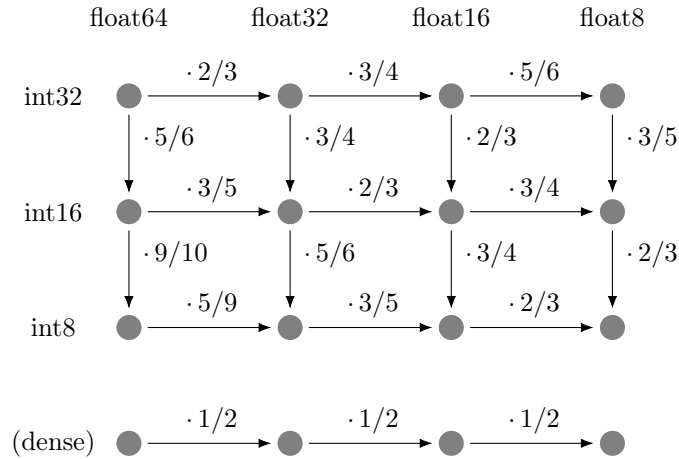


Figure 3: Approximate matrix-related traffic reduction when exchanging the data types used to store the matrix scalars (horizontally) or column indices (vertically) of (DA-) CSR storage, as well as dense storage (no indices required).

Following Listing 1, the matrix-related traffic amounts to

$$(\text{nrows} + 1) \cdot \text{sizeof}(\text{oindex\_t}) + \text{nnz} \cdot (\text{sizeof}(\text{iindex\_t}) + \text{sizeof}(\text{scalar\_t})). \quad (5)$$

If  $w$  may be stored in a smaller (integer) data type than  $\text{ncols}$ , this allows for a smaller  $\text{iindex\_t}$  to be used. Using an index type half the size nearly halves the bookkeeping traffic.

**Example 3.** The matrix *Janna/Bump\_2911* from the SuiteSparse Matrix Collection [4] has dimension 2911 419 and 127 729 899 non-zeros, distributed over a bandwidth of only  $31\,343 < 2^{15} = 32\,768$ , which is only about 1% of the matrix dimension. On average, this matrix has 44 non-zeros per row. See Figure 2 for its sparsity pattern. Therefore, standard CSR storage requires 32 bit indices for both  $\text{oindex\_t}$  and  $\text{iindex\_t}$ , which require 11 MiB and 487 MiB in total, respectively. DA-CSR allows for 16 bit  $\text{iindex\_t}$  to be used, which requires only 244 MiB, thus reducing the bookkeeping traffic by  $1 - \frac{11+244}{11+487} \approx 48.8\%$ , or from 4.09 to 2.09 Byte per nnz, irrespective of  $\text{scalar\_t}$ . For 64 bit and 32 bit  $\text{scalar\_t}$ , e.g., IEEE 754 *double* and *float*, which in total require 975 MiB and 487 MiB, using DA-CSR instead of CSR results in an overall matrix-related traffic reduction of  $1 - \frac{11+244+975}{11+487+975} \approx 16.5\%$  and  $1 - \frac{11+244+487}{11+487+487} \approx 24.7\%$ , respectively.

For matrices with more than a few non-zeros per row, it is therefore reasonable to ignore the effect of  $\text{oindex\_t}$ , i.e. to assume  $\text{sizeof}(\text{oindex\_t}) = 0$ . The final percentages of the previous example would then be estimated by  $\frac{1}{6}$  and  $\frac{1}{4}$ . Figure 3 shows this approximate reduction in matrix-related traffic by means of formula (5). Note how smaller  $\text{iindex\_t}$ ; i.e. lower bookkeeping traffic; yield better approximations of the factor  $\frac{1}{2}$  observed for dense storage.<sup>4</sup> Recall that by equation (1) a traffic reduction is tightly coupled with expected performance gains for memory-bound operations.

While the goal of multi-precision algorithms is to exchange  $\text{scalar\_t}$  for a smaller data type (as in, e.g., [1]), i.e. traversing the rows of Figure 3, DA storage focuses on  $\text{iindex\_t}$ , i.e. traversing the columns of Figure 3. However, as the main motivation of multi-precision algorithms is the memory bottleneck, DA storage is expected to enable even larger speedups in that context. CSR using 64 bit scalars and 32 bit indices merely allows for a  $\frac{3}{2} \times$  performance speedup when switching to 32 bit scalars. Meanwhile, if the matrix has a representation in DA-CSR using 16 bit indices, the expected speedup is  $\frac{5}{3} \times$ . This speedup is much closer to the  $2 \times$  possible for dense storage, when using a scalar type half the size.

<sup>4</sup>Note that dense storage may be seen as having  $\text{sizeof}(\text{oindex\_t}) = \text{sizeof}(\text{iindex\_t}) = 0$ , i.e. having zero bit of bookkeeping (per nnz).

### 3 Selection of Matrices

The SuiteSparse Matrix Collection [4] contains 1367 square matrices having a CSR representation using 32 bit indices and a full structural rank.<sup>5</sup> Only 993 of these matrices (72.6 %) have a dimension less than  $2^{15}$ , i.e. fit into CSR using 16 bit column indices. However, for 1302 matrices (95.2 %) there exists a permutation that reduces the matrix bandwidth to below  $2^{15}$ , such that these matrices fit into DA-CSR using 16 bit column indices. These are the matrices we select for further investigation.

Some of the investigated matrices are already stored in a bandwidth-reduced way. We applied an RCM reordering [2] to the ones that are not. Unfortunately, our implementation of the RCM permutation has not been able to sufficiently reduce the bandwidth of two matrices (Janna/Long\_Coup\_dt0 and Janna/Long\_Coup\_dt6), which reduces the number of matrices to 1300 (95.1 %).

### 4 Sparse Matrix Vector Product

Let  $A$  denote a matrix,  $x$  and  $y$  be vectors, and  $\alpha$  and  $\beta$  be scalars. The SpMV product denotes the operation  $y \leftarrow \alpha Ax + \beta y$ , which requires

$$W := 2\text{nnz} + 2\text{nrows} \quad (6)$$

floating-point operations of *work*  $W$  [FLOP]. The *performance*  $P$  [FLOP/s] is then defined as the ratio of work  $W$  and runtime  $t$ , where  $t$  denotes the runtime measured in elapsed time. The *relative performance* w.r.t. some baseline is computed via

$$P_{\text{candidate}}/P_{\text{baseline}} = t_{\text{baseline}}/t_{\text{candidate}}, \quad (7)$$

assuming  $W_{\text{candidate}} = W_{\text{baseline}}$ . The *traffic* [Byte] of computing the SpMV accounts for  $x$  and  $y$  on top of the three components of  $A$  in (DA-) CSR storage, cf. Listing 1 and formula (5). The *throughput* [Byte/s] is given by the ratio of traffic and  $t$ , and the *relative throughput* is then computed via

$$\frac{\text{traffic}_{\text{candidate}}}{\text{traffic}_{\text{baseline}}} \cdot \frac{t_{\text{baseline}}}{t_{\text{candidate}}}, \quad (8)$$

which is a scaled form of the relative performance. Refer to Figure 3 for typical and approximate expected traffic ratios. In the following, we aim to verify the predicted  $\frac{6}{5} \times$  speedup when replacing 32 bit column indices by 16 bit ones.

#### 4.1 Implementation Details and Methodology

A prototypical implementation of the SpMV product for a matrix in DA-CSR format is given in Listing 2. Instead of reversing the index translation in the innermost loop, i.e. computing `oindex_t col = row + colids[i]`, we instead compute a shifted view `xshift` into the factor `x` one level up. This replaces `nnz oindex_t`-additions by `nrows` pointer-additions. Recall that in C/C++ the memory access `x[row + col]` is equivalent to `*(x + (row + col))`. Applying associativity to the computation of the pointer address, we see that this access is also equivalent to `*((x + row) + col)` and `xshift[col]`.

**Remark 4** (Non-Square Matrices). For tall matrices, i.e. `nrows > ncols`, `xshift` points to memory outside `x`, and must therefore never be dereferenced directly. Within Listing 2 however, it will only be dereferenced at an offset `col` that yields a memory address within `x`.

The code has been compiled with GCC 10.3.0 using `-O3 -NDEBUG -mavx2 -mfma`. The benchmarks have been run on an Intel Xeon Skylake Silver 4110 running CentOS 7.9.2009,<sup>6</sup> with

<sup>5</sup>Eventually, we are interested in using DA storage when solving linear systems. We thus take full structural rank as a proxy for regularity, as the collection's metadata does not contain the numerical rank for all the matrices. Consequently, our selection of matrices contains irregular matrices as well.

<sup>6</sup><https://www.mpi-magdeburg.mpg.de/cluster/mechthild>

Listing 2: SpMV for DA-CSR

```

1 // Input:  oindex_t *rowptr, nrows;
2 //        iindex_t *colids;
3 //        scalar_t *values, *x, *y, alpha, beta;
4 // Output: scalar_t *y;
5 for (oindex_t row = 0; row < nrows; ++row) {
6     scalar_t accumulator = 0;
7     scalar_t *xshift = x + row;
8     for (oindex_t i = rowptr[row]; i < rowptr[row+1]; ++i) {
9         iindex_t col = colids[i];
10        scalar_t val = values[i];
11        accumulator += val * xshift[col];
12    }
13    y[row] = alpha * accumulator + beta * y[row];
14 }

```

threads pinned using `taskset`<sup>7</sup>. Runtime measurements have been taken using `nanobench` [8] with `minEpochTime` set to 100ms, `minEpochIterations` and `warmup` both set to 10, using the minimum over 11 epochs.

We measured the performance of a naive implementation (with `nnz` additions instead of `nrows`) as well as several implementations akin to Listing 2, optionally using OpenMP with 2, 4, 6, or 8 threads, both for CSR and DA-CSR matrices. Among all implementations executed on the given hardware, the best performing ones were the naive implementation, the one using 3 accumulators, and the one using a single AVX2 accumulator (4 scalars wide) accessing `values` using unaligned load instructions. In the following, for each matrix, and each storage format tested, we select the implementations and number of threads yielding the highest performance.

For the MKL [6] implementation of the CSR format, we measured single-threaded as well as multithreaded performance. Again, for each matrix we select the number of threads yielding the highest performance.

## 4.2 Numerical Results

For traffic within the size of the L1 cache, a single thread yields the best performance. Up to about 100 KiB, which is well within the size of a single L2 cache, the optimum number of threads increases gradually. For traffic larger than that, the maximum number of threads yields the best performance. This behavior is irrespective of the matrix format and the implementation vendor (ourselves or MKL [6]).

Our implementation of the SpMV product for the CSR format using 32 bit indices performs about the same as the MKL [6], see Table 1. Figure 4 shows the comparison of DA-CSR using 16 bit column indices to CSR. Using DA-CSR shows almost no change for traffic within the combined size of the L2 caches, i.e. up to 8 · 1 MiB. For traffic larger than that, up to the combined size of all caches, i.e. up to about 8 + 11 MiB, we observe a larger than 1.4× speedup. For traffic beyond that, we observe an average speedup of about +17.5%, which is reasonably close to the expected +20%. However, the throughput drops slightly, indicating some unused potential on the given hardware. See Table 2 for the comparison of DA-CSR to MKL [6].

**Remark 5** (CSR using 16 bit column indices). Recall that 933 matrices have a direct representation in CSR using smaller column indices. The DA-CSR format performs on par with CSR using the same index types for these matrices.

<sup>7</sup><https://www.man7.org/linux/man-pages/man1/taskset.1.html> and <https://github.com/util-linux/util-linux/blob/master/schedutils/taskset.c>

Table 1: Average relative performance of our best SpMV implementation for CSR using 32 bit indices w.r.t. MKL [6] as the baseline. Values > 1 mean we are faster.

Traffic	Singlethreaded	Multithreaded
L1d	1.131	1.129
L2	1.073	1.044
L3	1.012	1.011
Large	0.982	0.994

Table 2: Average relative performance of our best SpMV implementation for DA-CSR using 16 bit column indices w.r.t. MKL [6] as the baseline. Values > 1 mean we are faster.

Traffic	Single-threaded	Multithreaded
L1d	1.103	1.098
L2	1.073	1.055
L3	1.052	1.032
Large	1.110	1.172




Figure 4: Relative performance and throughput of SpMV using the DA-CSR format with 16 bit column indices w.r.t. CSR using 32 bit column indices as the baseline (iso-scalar). The sizes of the L1d, L2, and L3 CPU caches are marked with vertical lines (left to right).

## 5 Conclusion and Outlook

Diagonally-Adressed (DA) storage allows to nearly halve the bookkeeping traffic in sparse matrix storage formats, when the matrix bandwidth allows for an index type half the size. On the hardware used, DA-CSR storage with 16 bit column indices improves the single-threaded SpMV performance over CSR storage with 32 bit column indices by more than 17 %, for both our implementation and MKL [6] if the traffic exceeds the size of the L3 cache of the CPU. Meanwhile, DA-CSR performs no worse than CSR when using the same data types.

## Code and Data Availability

The source code is available at:

DOI [10.5281/zenodo.8104335](https://doi.org/10.5281/zenodo.8104335)

The visualizations in this paper have been created using TikZ [11] and Makie.jl [3]. The SpMV performance measurements for the reported experiments are available at:

DOI [10.5281/zenodo.7551699](https://doi.org/10.5281/zenodo.7551699)

## References

- [1] A. Abdelfattah, H. Anzt, E. G. Boman, E. Carson, T. Cojean, J. Dongarra, A. Fox, M. Gates, N. J. Higham, X. S. Li, et al. A survey of numerical linear algebra methods utilizing mixed-precision arithmetic. *The International Journal of High Performance Computing Applications*, 35(4):344–369, 2021.
- [2] E. Cuthill and J. McKee. Reducing the bandwidth of sparse symmetric matrices. In *Proceedings of the 1969 24th national conference (ACM'69)*, New York, New York, USA, 1969. ACM Press. DOI [10.1145/800195.805928](https://doi.org/10.1145/800195.805928).
- [3] S. Danisch and J. Krumbiegel. Makie.jl: Flexible high-performance data visualization for Julia. *Journal of Open Source Software*, 6(65):3349, 2021. DOI [10.21105/joss.03349](https://doi.org/10.21105/joss.03349).
- [4] T. A. Davis and Y. Hu. The University of Florida sparse matrix collection. *ACM Transactions on Mathematical Software (TOMS)*, 38(1):1–25, 2011.
- [5] G. Goumas, K. Kourtis, N. Anastopoulos, V. Karakasis, and N. Koziris. Understanding the Performance of Sparse Matrix-Vector Multiplication. In *16th Euromicro Conference on Parallel, Distributed and Network-Based Processing (PDP 2008)*. IEEE, 2008. DOI [10.1109/pdp.2008.41](https://doi.org/10.1109/pdp.2008.41).
- [6] Intel. Math Kernel Library v2021.1, 2020. [https://en.wikipedia.org/wiki/Math\\_Kernel\\_Library](https://en.wikipedia.org/wiki/Math_Kernel_Library).
- [7] Z. Koza, M. Matyka, Ł. Mirosław, and J. Poła. *Sparse Matrix-Vector Product*, pages 103–121. Springer International Publishing, Cham, 2014. DOI [10.1007/978-3-319-06548-9\\_6](https://doi.org/10.1007/978-3-319-06548-9_6).
- [8] M. Leitner-Ankerl. ankerl::nanobench, 2022. <https://github.com/martinus/nanobench>.
- [9] X. Liu, M. Smelyanskiy, E. Chow, and P. Dubey. Efficient sparse matrix-vector multiplication on x86-based many-core processors. In *the 27th international ACM conference*, New York, New York, USA, 2013. ACM Press. DOI [10.1145/2464996.2465013](https://doi.org/10.1145/2464996.2465013).
- [10] M. Martone, S. Filippone, S. Tucci, M. Paprzycki, and M. Ganzha. Utilizing Recursive Storage in Sparse Matrix-Vector Multiplication - Preliminary Considerations. In T. Philips, editor, *Proceedings of the ISCA 25th International Conference on Computers and Their Applications, CATA 2010, March 24-26, 2010, Sheraton Waikiki Hotel, Honolulu, Hawaii, USA*, pages 300–305. ISCA, 2010.
- [11] T. Tantau. The TikZ and PGF Packages, 2020. Manual for version 3.1.5b. <https://github.com/pgf-tikz/pgf>.