

Performance of \mathcal{H} -Matrix-Vector Multiplication with Floating Point Compression

Ronald Kriemann
MPI for Mathematics i.t.S.
Leipzig, Germany
rok@mis.mpg.de

May 7, 2024

Abstract Matrix-vector multiplication forms the basis of many iterative solution algorithms and as such is an important algorithm also for hierarchical matrices. However, due to its low computational intensity, its performance is typically limited by the available memory bandwidth. By optimizing the storage representation of the data within such matrices, this limitation can be lifted and the performance increased. This applies not only to hierarchical matrices but also for other low-rank approximation schemes, e.g. block low-rank matrices.

AMS Subject Classification: 65Y05, 65Y20, 68W10, 68W25, 68P30

Keywords: hierarchical matrices, low-rank arithmetic, data compression, matrix-vector multiplication

1 Introduction

Introduced in [15] was a version of (hierarchical) low-rank arithmetic where the matrix data, i.e., dense and low-rank blocks, was compressed using floating point compression methods. As the standard \mathcal{H} -arithmetic is typically based on dense arithmetic functions defined by the BLAS and LAPACK function set [5], the modified \mathcal{H} -arithmetic in [15] was based on the idea of decompressing all input data of arithmetic kernel functions, executing the arithmetic kernel in standard double precision and then compressing the output data. This way, the actual arithmetic functions remain unchanged. Already this approach showed superior performance for the \mathcal{H} -matrix-vector multiplication (\mathcal{H} -MVM), which is often memory bandwidth limited and as such, any reduction of the memory size will increase performance.

Another reason for this *semi-on-the-fly* approach was the general compression approach in [15], i.e., for floating point

data *any* compressor could be used. This prevents direct arithmetic within the (unknown) compression format.

However, some of the compression schemes in [15] allow random access of entries in the compressed storage and hence, special arithmetic functions can be implemented. The aim of this work is to investigate the benefit of such an approach for \mathcal{H} -matrix vector multiplication.

An analog strategy was used in [6] with the idea of a *memory accessor*, i.e., transparent conversion between a storage and a computation format within a sparse matrix computation. This work is therefore the application of this concept for \mathcal{H} -matrix arithmetic.

A different strategy is used by mixed precision schemes ([20, 1]) where combinations of hardware provided floating point formats are used to reduce the memory footprint and to increase performance, partly due to faster execution of such smaller data formats. However, these approaches are limited in the reduction of the \mathcal{H} -matrix memory which is crucial on computer systems with memory bandwidth limitations. Also, not all floating point formats are (yet) hardware supported on all platforms, e.g., half precision formats like BF16 or FP16.

The rest of this work is structured as follows: in Section 2 basic definitions and algorithms for \mathcal{H} -matrices are introduced together with the introduction of compression schemes for dense and low-rank data. Section 3 will discuss different strategies for \mathcal{H} -matrix-vector multiplication with and without compression. Numerical experiments will be presented in Section 4, followed by a conclusion in Section 5.

2 \mathcal{H} -Matrices

For an indexset I we define the *cluster tree* (or \mathcal{H} -tree) as the hierarchical partitioning of I into disjoint sub-sets of I :

Definition 2.1 (Cluster Tree) Let $T_I = (V, E)$ be a tree with $V \subset \mathcal{P}(I)$. T_I is called a cluster tree over I if

1. $I = \text{root}(T_I)$ and
2. for all $v \in V$ with $\text{sons}(v) \neq \emptyset$: $v = \dot{\cup}_{v' \in \text{sons}(v)} v'$.

A node in T_I is also called a cluster and we write $\tau \in T_I$ if $\tau \in V$. The set of leaves of T_I is denoted by $\mathcal{L}(T_I)$.

Similar to a cluster tree we can extend the hierarchical partitioning to the product $I \times J$ of two index sets I, J , while restricting the possible set of nodes by given cluster trees T_I and T_J over I and J , respectively. Furthermore, the set of leaves will be defined by an *admissibility condition*. In the literature, various examples of admissibility can be found, e.g. standard [11], weak [12] or off-diagonal admissibility [9, 2].

Definition 2.2 (Block Tree) Let T_I, T_J be two cluster trees and let $\text{adm} : T_I \times T_J \rightarrow \mathbb{B}$. The block tree $T = T_{I \times J}$ is recursively defined starting with $\text{root}(T) = (I, J)$:

$$\begin{aligned} \text{sons}(\tau, \sigma) = \\ \begin{cases} \emptyset, & \text{if } \text{adm}(\tau, \sigma) = \text{true} \vee \text{sons}(\tau) = \emptyset \vee \text{sons}(\sigma) = \emptyset, \\ \{(\tau', \sigma') : \tau' \in \text{sons}(\tau), \sigma' \in \text{sons}(\sigma)\} & \text{else.} \end{cases} \end{aligned}$$

A node in T is also called a block. Again, the set of leaves of T is denoted by $\mathcal{L}(T) := \{b \in T : \text{sons}(b) = \emptyset\}$.

The admissibility condition is used to detect blocks in T which can be efficiently approximated by low-rank matrices with a predefined rank k , i.e., blocks b with $\text{adm}(b) = \text{true}$. The set of all such matrices forms the set of \mathcal{H} -matrices:

Definition 2.3 (\mathcal{H} -Matrix) For a block tree T over cluster trees T_I, T_J and $k \in \mathbb{N}$, the set of \mathcal{H} -matrices $\mathcal{H}(T, k)$ is defined as

$$\mathcal{H}(T, k) := \{M \in \mathbb{R}^{I \times J} : \forall (\tau, \sigma) \in \mathcal{L}(T) : \text{rank}(M_{\tau, \sigma}) \leq k \vee \tau \in \mathcal{L}(T_I) \vee \sigma \in \mathcal{L}(T_J)\}$$

Here, $M_{\tau, \sigma}$ refers to the sub-block $M|_{\tau \times \sigma}$.

In practice the constant rank k is typically replaced by a fixed low-rank approximation accuracy $\varepsilon > 0$ as the resulting \mathcal{H} -matrices are often more memory efficient. For this we assume for an admissible block $M_{\tau, \sigma}$:

$$\|M_{\tau, \sigma} - U_{\tau, \sigma} V_{\tau, \sigma}^H\| \leq \varepsilon \|M_{\tau, \sigma}\|. \quad (1)$$

In an analog way to $\mathcal{H}(T, k)$, the set $\mathcal{H}(T, \varepsilon)$ can be defined as the set of \mathcal{H} -matrices with *local* low-rank approximation error of ε . We will also use $\mathcal{H}(T)$ if either a fixed rank or a fixed accuracy is used.

Remark 2.4 The set $\mathcal{H}(T)$ also includes various other formats like block low-rank (BLR) [3] or hierarchical off-diagonal low-rank (HODLR) [2], as only the clustering or the admissibility has to be chosen appropriately.

Remark 2.5 For \mathcal{H} -matrices with a full hierarchy, the set $M_\tau := \{M|_{\tau, \sigma} : (\tau, \sigma) \in \mathcal{L}(T) \wedge \text{rank}(M_{\tau, \sigma}) \leq k\}$ of low-rank blocks for a cluster $\tau \in T_I$ is bounded by the constant c_{sp} [10] for a particular application.

2.1 Compressed \mathcal{H} -Matrices

Floating point data in \mathcal{H} -matrices appears in inadmissible blocks as dense matrices holding the coefficients and in low-rank blocks in the form of the low-rank factors. Often these are stored in FP64 (or FP32) format. However, due to low-rank approximation with accuracy ε , already an error is introduced which is typically much larger than the unit roundoff of FP64 (or even FP32).

In [16, 15] the FP64 storage was replaced by error adaptive floating point compression, i.e., an optimized storage format was chosen with a representation error depending on ε . Different compressors are available to implement such a direct compression of floating point data, e.g., ZFP [19] or BLOSC [8]. Furthermore, different storage schemes based on the IEEE-754 floating point standard were examined, where the number of mantissa bits m_ε is chosen based on the low-rank approximation error ε as $m := \lceil -\log_2 \varepsilon \rceil$. Different choices for the number of exponent bits e were also examined, e.g., with 8 bits as in the FP32 or BF16 formats (called BFL), 11 bits as in FP64 (called DFL) and an adaptive choice based on the *dynamic range* of the data, i.e., the base 10 logarithm of the ratio between the largest and smallest (absolute) value (called AFLP). In all cases m_ε was increased such that the number of bits per value $1 + e + m_\varepsilon$ ¹ is a multiple of 8 for fast byte aligned storage.

Independent on the particular choice of the compression scheme, this *direct* compression mode is then applied to the dense data of inadmissible blocks $M_{\tau, \sigma} \in \mathcal{L}_{\text{inadm}}$ and the low-rank factors $U_{\tau', \sigma'}, V_{\tau', \sigma'}$ of admissible blocks $M_{\tau', \sigma'} \in \mathcal{L}_{\text{adm}}$.

Furthermore, as described in [15], low-rank matrices permit an advanced compression scheme with an adaptive accuracy choice for each column in the low-rank factors. This *adaptive precision* compression (APLR) is based on the mixed precision approach described in [4]. For a block $M_{\tau, \sigma}$ we assume a rank- k approximation $U \cdot V^H$ with $\|M_{\tau, \sigma} - UV^H\| \leq \delta$. Using the singular value decomposition we can find orthogonal matrices W and X and a diagonal matrix $\Sigma = \text{diag}(\sigma_0, \dots, \sigma_{k-1})$ with the singular values $\sigma_0 > \sigma_1 > \dots > \sigma_{k-1}$ of UV^H .

If the i 'th column w_i of W and x_i of X is stored with precision δ/σ_i , then the total approximation error is (see [15, Section 4])

$$\|M_{\tau, \sigma} - \widetilde{W} \Sigma \widetilde{X}^H\| \leq \delta + \left(2\delta k + \delta^2 \sum_{i=1}^k \frac{1}{\sigma_i} \right).$$

With this, any direct floating point compression method \mathcal{Z} can be used to yield an improved storage method for low-rank matrices, denoted APLR- \mathcal{Z} , e.g., APLR-AFLP, APLR-BFL or APLR-DFL. The main advantage of this scheme compared to direct compression is, that in the latter case the chosen precision is applied to the full data whereas with APLR even for a high accuracy, a low precision may be used for some part of the data.

3 \mathcal{H} -Matrix-Vector Multiplication

We consider the update operation

$$y := \alpha Ax + y$$

¹With the additional sign bit.

with an \mathcal{H} -matrix $A \in \mathcal{H}(T, k)$ and vectors x and y . The product is computed by looping over the leaf blocks of A and performing local matrix-vector multiplications, either with a dense matrix for inadmissible blocks or in low-rank format, i.e., $t := V_{\tau,\sigma}^H x|_\sigma$ followed by $y|_\tau := y|_\tau + \alpha U_{\tau,\sigma} t$. The full procedure is shown in Algorithm 1.

Algorithm 1: \mathcal{H} -Matrix-Vector Multiplication

```

procedure hmv( $\alpha, A, x, y$ )
  for  $(\tau, \sigma) \in \mathcal{L}$  do
    if  $(\tau, \sigma)$  is admissible then
       $y|_\tau := y|_\tau + \alpha U_{\tau,\sigma} V_{\tau,\sigma}^H x|_\sigma$ ;
    else
       $y|_\tau := y|_\tau + \alpha D_{\tau,\sigma} x|_\sigma$ ;

```

Versions of the \mathcal{H} -MVM for parallel systems need to consider load balancing due to different, not a priori known ranks in different low-rank blocks of the \mathcal{H} -matrix if a fixed accuracy ε is used. Also the block structure is typically not equal throughout the matrix. This poses a serious scalability issue for the distributed memory case (see [7, 17] or systems with a NUMA architecture.

On shared memory systems a task-based approach can avoid these problems if the scheduling algorithm is able to assign ready tasks to idle processors. However, this may lead to other problems as the memory layout of the blocks handled by a single processor may not be optimal for efficient execution. This is of special importance because of the low computational intensity of matrix-vector multiplication, which normally leads to a memory bandwidth limited performance. Different optimization strategies are discussed in [13], where especially the memory layout of the \mathcal{H} -matrix data is adjusted such that memory loads are faster.

Another issue with shared memory programming is handling potential collisions when writing to the same memory positions, e.g., with matrix blocks $A_{\tau,\sigma}$ and $A_{\tau,\sigma'}$ handled by different processors writing simultaneously to $y|_\tau$. Solutions to this problem involve *atomic* updates [14] or reduction of thread local results [13]. A reduction approach of local results is also the default choice for the distributed memory case [7, 17].

An alternative approach is a collision free design in which the memory blocks are scheduled to the processors in a way to prevent simultaneous writing to the same memory positions. Such a method is used in the following.

Let $\mathcal{A}_\tau := \{A_{\tau,\sigma} : (\tau, \sigma) \in \mathcal{L}(T)\}$ be the set of all matrix blocks in A with identical row cluster τ and let $\mathcal{A} := \{\mathcal{A}_\tau : \tau \in T_I\}$ be the set of all such block lists. Since \mathcal{A} is defined based on T_I , it can be considered to be structurally identical to the cluster tree. Due to its definition, the number of matrix blocks in any \mathcal{A}_τ is bounded by c_{sp} (see 2.5) and therefore independent on the dimension of the matrix..

Now let τ_0, \dots, τ_ℓ be clusters of T_I with identical level,

i.e., $\text{depth}(\tau_i) = \text{depth}(\tau_j), 0 \leq i, j \leq \ell$. Then, for any $0 \leq i, j \leq \ell$ the matrix-vector products in the corresponding sets \mathcal{A}_{τ_i} and \mathcal{A}_{τ_j} can be computed in parallel since $\tau_i \cap \tau_j = \emptyset$.

For any $\tau, \sigma \in T_I$ with $\text{depth}(\tau) \neq \sigma$ the sets \mathcal{A}_τ and \mathcal{A}_{σ} can only be executed in parallel if $\tau \cap \sigma = \emptyset$. However, due to the definition of T_I if $\tau \cap \sigma \neq \emptyset$ then either $\tau \subseteq \sigma$ or $\sigma \subseteq \tau$ holds. Therefore, if T_I is traversed from root to bottom with execution of matrix blocks in a given \mathcal{A}_τ before proceeding to the sons in $\mathcal{S}(\tau)$, any race condition when accessing y is prevented. This procedure is implemented in Algorithm 2.

Algorithm 2: Parallel \mathcal{H} -Matrix-Vector Multiplication

```

procedure phmv( $\alpha, \tau, \mathcal{A}, x, y$ )
  for all  $A_{\tau,\sigma} \in \mathcal{A}_\tau$  do
    if  $(\tau, \sigma)$  is admissible then
       $y|_\tau := y|_\tau + \alpha U_{\tau,\sigma} V_{\tau,\sigma}^H x|_\sigma$ ;
    else
       $y|_\tau := y|_\tau + \alpha D_{\tau,\sigma} x|_\sigma$ ;
  parallel for  $(\tau' \in \mathcal{S}(\tau))$  do
    phmv( $\alpha, \tau', \mathcal{A}, x, y$ )

```




Figure 1: Roofline plot for \mathcal{H} -MVM.

When using this algorithm with the architecture used in Section 4, almost optimal performance is achieved as shown in the (empirical) roofline plot in Figure 1. Please note, that the plot shows values obtained for different problem sizes and therefore data sizes, demonstrating its performance consistency.

In principle, the computation of all products for matrix blocks in \mathcal{A}_τ in Algorithm 2 can be further parallelized using a reduction scheme. One could also combine memory layout optimizations from [13] easily with this approach. However, as already the performance limit is reached by the above procedure, only minor improvements by such modifications are possible.

In Algorithm 2 no parallelism of the per matrix block products was considered, which is a major drawback of this procedure if a block structure like HODLR [2] is used. There,

the most time consuming computations are performed on the upper levels of the block cluster tree, where only a few processors may be used in parallel. However, as the numerical results in Section 4 demonstrate, typical \mathcal{H} -matrix block structures do not show these problems, at least not for the number of processors cores considered in this work.

3.1 Compressed \mathcal{H} -Matrix-Vector Multiplication

The main interest of this work is in the performance of \mathcal{H} -MVM when using floating point compression. This was already the topic of [15]. There, for a dense or low-rank block, the compressed data was first fully converted into the computation format and only then the local matrix-vector multiplication was performed in double precision. This way, the standard arithmetic kernels, typically optimized by hardware vendors, could be reused.

In [6] the concept of a *memory-accessor* is described which implements on-the-fly conversion between the storage format and the computation format during the arithmetic. Since \mathcal{H} -MVM only uses decompression this approach is easier to apply in this case compared to the full \mathcal{H} -arithmetic. Also, since \mathcal{H} -MVM is often memory bandwidth limited it may be more forgiving for a (potentially) less heavily optimized implementation.

In any case, such an approach requires fast access to the compressed values, which holds for the above described compression schemes AFLP, BFL and DFL but much less so for ZFP or BLOSC. Furthermore, especially ZFP showed a low performance compared to other formats (see [15]). However, in principle by tightly coupling the compression scheme with the matrix-vector multiplication, any compression format could be used. As such, the restriction to AFLP, BFL and DFL in this work should be considered a proof-of-concept.

Since Algorithm 2 only uses dense matrix-vector multiplication, one only needs to focus on this function. The actual implementation of the compressed version is straightforward, without particular optimizations, aside from standard code reorganization due to the used column-major storage scheme and shown in Algorithm 3 for the application of a non-transposed $n \times m$ matrix D . Only the access to the coefficients in D is replaced by the corresponding coefficient decompression.

Algorithm 3: Matrix-vector multiplication with compressed dense $n \times m$ matrix

```

procedure zvmv(in:  $D, x$ , inout:  $y$ )
  for  $0 \leq j < m$  do
    for  $0 \leq i < n$  do
       $y_i := y_i + \text{decompress}(D_{ij})x_j$ 

```

For the mixed precision approach in [4] in the predecessor in [20] (using FP64 and FP32) which are based on

hardware supported floating point formats, one has the advantage of performing the computations for the corresponding floating point hardware natively, i.e., without converting to FP64, thereby potentially increasing the performance. As this had no negative side effects on the error for the model problems used in Section 4 this was also used in the following.

4 Numerical Experiments

The model problem is based on a boundary element discretization for the Laplace single layer potential (Laplace SLP) while the domain is defined by the unit sphere:

$$\int_{\Omega} \frac{1}{\|x - y\|} u(y) dy = f(x), \quad x \in \Omega \quad (2)$$

with $\Omega = \{x \in \mathbb{R}^3 : \|x\|_2 = 1\}$. Piecewise constant ansatz functions are used for the discretization. Furthermore, standard admissibility

$$\min \{\text{diam}(t), \text{diam}(s)\} \leq \eta \text{dist}(\tau, \sigma)$$

is applied for setting up the block tree.

All experiments are performed on an AMD Epyc 9554 CPU with 64 cores in total and 12 32GB DDR5-4800 memory DIMMs. For parallelization Intel TBB v2021.11 was used while Intel oneMKL v2024.0 provided the BLAS and LAPACK functions for the uncompressed case. Please note, that the sequential version was chosen as all parallelization is performed within the \mathcal{H} -arithmetic itself. Furthermore, the AVX512 code path in MKL was activated. All code was compiled using GCC v12.3.

The algorithms described in this work are implemented in the open source software HLR². For the numerical experiments version *9cdb804* was used.

For runtime results the median of ten runs is presented, while only little variations between each run was observed.

Aside from the compressors AFLP, BFL, DFL the mixed precision formats *MP-3* using FP64, FP32 and BF16³ (as in [4]) and *MP-2* with FP64 and FP32 (as in [20]) are used.

First, the compression ratio for the different schemes is shown in Figure 2 for a comparison on the memory savings. As can be seen the compression ration is slightly improving with a growing problem size if a fixed accuracy is chosen. This corresponds to the increasing portion of low-rank memory compared to memory associated to inadmissible blocks in \mathcal{H} -matrices with larger problems and is more pronounced for the mixed precision formats as there inadmissible blocks are not compressed. Due to a larger number of exponent bits in BFL and DFL, these formats have a higher memory consumption compared to AFLP.

²<http://libhlr.org>, program: *mpvm*

³BF16 was preferred over FP16 due to much faster conversion from FP64/FP32.




Figure 2: Compression rates compared to uncompressed \mathcal{H} -matrices for fixed accuracy of $\varepsilon = 10^{-6}$ (left) and a fixed problem size of $n = 2,097,152$ (right).

For a fixed problem size and varying accuracy, AFLP compression yields a high compression for low and high accuracies. In contrast to this MP-3 and MP-2 show a much worse compression if low accuracy is used, again due to the fact that it is applied to low-rank blocks only. However, for higher accuracies especially MP-3 is able to come closer to the other formats.

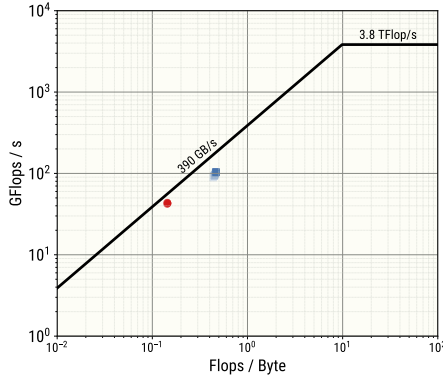


Figure 3: Roofline plot for \mathcal{H} -MVM without (red) and with compression using AFLP (blue) for a fixed accuracy of $\varepsilon = 10^{-6}$.

Figure 3 shows the performance of uncompressed and compressed \mathcal{H} -MVM in an empirical roofline plot with a maximal memory bandwidth of 390 GB/s and maximal floating point performance of 3.8 TFlop/s, measured by the Likwid tool [18]. The arithmetic intensity of \mathcal{H} -MVM is well within the bandwidth limited regime. While the uncompressed multiplication is close to the limit, the compressed \mathcal{H} -MVM is slightly less optimal. However, the latter does not make use of the optimized matrix-vector product from the Intel MKL library and also has the additional overhead of the coefficient decompression.

Nevertheless, the performance is significantly increased by lowering the memory bandwidth requirements. The run-

time speedup compared to the uncompressed multiplication is shown in Figure 4 for all compression schemes. There a better compression ratio directly translates into a better performance with AFLP yielding best results.

To demonstrate, that this performance improvement is not restricted to \mathcal{H} -matrices, the same problem was computed using the Block Low-Rank approach from [3]. Here, the weak admissibility [12] was used. The results for BLR matrix-vector multiplication (BLR-MVM) are shown in Figure 5. The maximal problem sizes are smaller due to the higher memory demands of the BLR format.

While the general picture is similar to the \mathcal{H} -matrix case, the runtime improvements compared to the uncompressed case are even slightly bigger. Only the mixed precision formats do show a worse behavior for small problem sizes.

5 Conclusion

Matrix-vector multiplication for \mathcal{H} -matrices can benefit significantly from an optimized memory representation of the dense and low-rank data on platforms with slow memory access as is the case for many CPU based computations. With AFLP an efficient compressor is available which permits such optimizations if combined with floating point compression schemes with fast access to individual values. The such presented algorithm not only performs well for \mathcal{H} -matrices but may be used also for other forms of low-rank storage like BLR. It will be interesting to see if these concepts can be applied to more complex forms of \mathcal{H} -arithmetic, e.g. matrix-multiplication or LU factorization.

References

- [1] S. Abdulah et al. “Accelerating Geostatistical Modeling and Prediction With Mixed-Precision Computations: A High-Productivity Approach With PaRSEC”. In: *IEEE Transactions*




Figure 4: Speedup of compressed \mathcal{H} -MVM vs. uncompressed \mathcal{H} -MVM (base line) for fixed accuracy of $\varepsilon = 10^{-6}$ (left) and a fixed problem size of $n = 2.097.152$ (right).

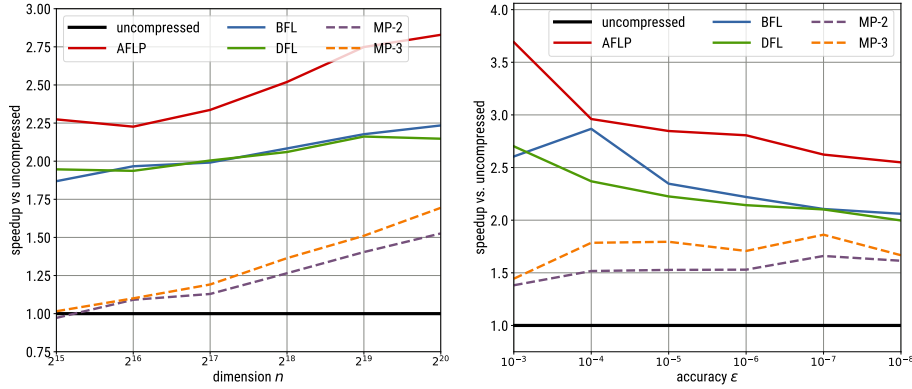


Figure 5: Speedup of compressed BLR-MVM vs. uncompressed BLR-MVM (base line) for fixed accuracy of $\varepsilon = 10^{-6}$ (left) and a fixed problem size of $n = 1.048.576$ (right).

- on Parallel and Distributed Systems 33.4 (2022), pp. 964–976. DOI: 10.1109/TPDS.2021.3084071.
- [2] S. Ambikasaran and E. Darve. “An $\mathcal{O}(N \log N)$ Fast Direct Solver for Partial Hierarchically Semi-Separable Matrices”. In: *Journal of Scientific Computing* 57.3 (2013), pp. 477–501.
 - [3] P. Amestoy et al. “Improving Multifrontal Methods by Means of Block Low-Rank Representations”. In: *SIAM Journal on Scientific Computing* 37.3 (2015), A1451–A1474. DOI: 10.1137/120903476.
 - [4] P. Amestoy et al. “Mixed precision low-rank approximations and their application to block low-rank LU factorization”. In: *IMA Journal of Numerical Analysis* (Aug. 2022). DOI: 10.1093/imanum/drac037.
 - [5] E. Anderson et al. *LAPACK Users’ Guide*. Third. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1999. ISBN: 0-89871-447-8 (paperback).
 - [6] Hartwig Anzt, Thomas Grützmacher, and Enrique S. Quintana-Orti. “Toward a modular precision ecosystem for high-performance computing”. In: *International Journal of High Performance Computing Applications* 33.6 (2019), pp. 1069–1078. DOI: 10.1177/1094342019846547.
 - [7] M. Bebendorf and R. Kriemann. “Fast parallel solution of boundary integral equations and related problems”. In: *Computing and Visualization in Science* 8.3 (2005), pp. 121–134. DOI: 10.1007/s00791-005-0001-x.
 - [8] Blosc Development Team. “A fast, compressed and persistent data store library”. <https://blosc.org>. 2009–2023.
 - [9] S. Chandrasekaran et al. “Some Fast Algorithms for Sequentially Semiseparable Representations”. In: *SIAM Journal on Matrix Analysis and Applications* 27 (2 2005), pp. 341–364.
 - [10] L. Grasedyck and W. Hackbusch. “Construction and arithmetics of \mathcal{H} -matrices”. In: *Computing* 70 (2003), pp. 295–334.
 - [11] W. Hackbusch and B.N. Khoromskij. “A sparse \mathcal{H} -matrix arithmetic. Part II. Application to multi-dimensional problems”. In: *Computing* 64.1 (2000), pp. 21–47.
 - [12] W. Hackbusch, B.N. Khoromskij, and R. Kriemann. “Hierarchical Matrices Based on a Weak Admissibility Criterion”. In: *Computing* 73.3 (2004), pp. 207–243. DOI: 10.1007/s00607-004-0080-4.
 - [13] T. Hoshino, A. Ida, and T. Hanawa. “Optimizations of \mathcal{H} -matrix-vector Multiplication for Modern Multi-core Processors”. In: *2022 IEEE International Conference on Cluster*

- Computing (CLUSTER)*. 2022, pp. 462–472. DOI: 10.1109/CLUSTER51413.2022.00056.
- [14] A. Ida, T. Iwashita, T. Mifune, and Y. Takahashi. “Parallel Hierarchical Matrices with Adaptive Cross Approximation on Symmetric Multiprocessing Clusters”. In: *Journal of Information Processing* 22.4 (2014), pp. 642–650. DOI: 10.2197/ipsjjip.22.642.
- [15] R. Kriemann. “Hierarchical Lowrank Arithmetic with Binary Compression”. 2023. arXiv: 2308.10960 [cs.MS].
- [16] R. Kriemann et al. “High-Performance Spatial Data Compression for Scientific Applications”. In: *Euro-Par 2022: Parallel Processing*. Ed. by José Cano and Phil Trinder. Cham: Springer International Publishing, 2022, pp. 403–418. ISBN: 978-3-031-12597-3.
- [17] Y. Li, J. Poulson, and L. Ying. “Distributed-Memory \mathcal{H} -Matrix Algebra I: Data Distribution and Matrix-Vector Multiplication”. In: *CSIAM Transactions on Applied Mathematics* 2.3 (2021), pp. 431–459. ISSN: 2708-0579. DOI: 10.4208/csiam-am.2020-0206.
- [18] LIKWID. 2024. URL: <https://github.com/RRZE-HPC/likwid/> (visited on 04/29/2024).
- [19] P. Lindstrom. “Fixed-Rate Compressed Floating-Point Arrays”. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2674–2683.
- [20] R. Ooi et al. “Effect of Mixed Precision Computing on \mathcal{H} -Matrix Vector Multiplication in BEM Analysis”. In: *Proceedings of the International Conference on High Performance Computing in Asia-Pacific Region. HPCAsia '20.*, Fukuoka, Japan, Association for Computing Machinery, 2020, pp. 92–101. ISBN: 9781450372367. DOI: 10.1145/3368474.3368479. URL: <https://doi.org/10.1145/3368474.3368479>.