

Algorithmic patterns for \mathcal{H} -matrices on many-core processors

P. Zaspel

Departement Mathematik und Informatik
Fachbereich Mathematik
Universität Basel
CH-4051 Basel

Preprint No. 2017-12
August 2017

www.math.unibas.ch

Algorithmic patterns for \mathcal{H} -matrices on many-core processors

Peter Zaspel

Received: date / Accepted: date

Abstract In this work, we consider the reformulation of hierarchical (\mathcal{H}) matrix algorithms for many-core processors with a model implementation on graphics processing units (GPUs). \mathcal{H} matrices approximate specific dense matrices, e.g., from discretized integral equations or kernel ridge regression, leading to log-linear time complexity in dense matrix-vector products. The parallelization of \mathcal{H} matrix operations on many-core processors is difficult due to the complex nature of the underlying algorithms. While previous algorithmic advances for many-core hardware focused on accelerating existing \mathcal{H} matrix CPU implementations by many-core processors, we here aim at totally relying on that processor type. As main contribution, we introduce the necessary parallel algorithmic patterns allowing to map the full \mathcal{H} matrix construction and the fast matrix-vector product to many-core hardware. Here, crucial ingredients are space filling curves, parallel tree traversal and batching of linear algebra operations. The resulting model GPU implementation `hmglib` is the, to the best of the authors knowledge, first entirely GPU-based Open Source \mathcal{H} matrix library of this kind. We conclude this work by an in-depth performance analysis and a comparative performance study against a standard \mathcal{H} matrix library, highlighting profound speedups of our many-core parallel approach.

Keywords Hierarchical matrices · GPU · Batched Linear algebra · Many-core parallelization · Space filling curves · Kernel ridge regression

Mathematics Subject Classification (2000) 65Y05 · 68T05 · 65Y10 · 68W10 · 65Y20 · 65F30 · 65F10 · 15A06 · 015-04 · 65N38

P. Zaspel
Departement Mathematik und Informatik
Universität Basel
Spiegelgasse 1
4051 Basel, Switzerland
E-mail: peter.zaspel@unibas.ch

1 Introduction

In many fields of applications we are required to solve large dense linear systems of equations of the form

$$A_{\phi, \mathcal{Y} \times \mathcal{Y}} \mathbf{x} = \mathbf{b} \quad (1)$$

with

$$A_{\phi, \mathcal{Y} \times \mathcal{Y}} = \begin{pmatrix} \phi(\mathbf{y}_1, \mathbf{y}_1) & \cdots & \phi(\mathbf{y}_1, \mathbf{y}_N) \\ \vdots & \ddots & \vdots \\ \phi(\mathbf{y}_N, \mathbf{y}_1) & \cdots & \phi(\mathbf{y}_N, \mathbf{y}_N) \end{pmatrix}, \quad \mathbf{x}, \mathbf{b} \in \mathbb{R}^N. \quad (2)$$

where $\mathcal{Y} := \{\mathbf{y}_1, \dots, \mathbf{y}_N\} \subset \Omega$ is a set of N points in a space $\Omega \subset \mathbb{R}^d$ and $\phi : \Omega \times \Omega \rightarrow \mathbb{R}$ is a bivariate *kernel* function operating on that domain. In *kernel-based interpolation* [40], the linear system (1) arises in the computation of interpolation coefficients. In *Gaussian Process Regression* (GPR) [34] kernel ϕ is a covariance function and $A_{\phi, \mathcal{Y} \times \mathcal{Y}}$ is replaced by $(A_{\phi, \mathcal{Y} \times \mathcal{Y}} + \sigma^2 I)$ with σ^2 a (scalar) variance and I the unity matrix. The same modified system also shows up in *kernel ridge regression* [39]. Integral equations, discretized by e.g. collocation, lead to similar linear systems. Note that, even though we here stick to the model problem (1) with collocation matrices of type (2), all our developments can be equally applied e.g. in the context of boundary element method problems.

The problem size N might get very large. As an example, N could be the number of training samples in machine learning by kernel ridge regression. This can grow up to tens to hundreds of millions of samples or even more, depending on the application. At this point, obviously, linear solvers for (1) based on direct factorization get intractable due to their $O(N^3)$ complexity. This is overcome by iterative solvers with fast approximate dense matrix-vector product.

In this work, we address the topic of parallelization of the fast approximate dense matrix-vector product based on *hierarchical matrices* (\mathcal{H} matrices) [12, 4, 21, 22]. Using \mathcal{H} matrix techniques, a matrix-vector product for a fixed approximation accuracy is done in $O(N \log N)$ operations, given ϕ is *asymptotically smooth*, cf. Section 2. Similar to panel clustering [25] and multipole techniques [20], the core idea is to distinguish between subsets $\mathcal{Y}_i \times \mathcal{Y}_j \subset \mathcal{Y} \times \mathcal{Y}$, where \mathcal{Y}_i and \mathcal{Y}_j are “close” to each other or “far away”. In \mathcal{H} matrices, a *tree-based* spatial decomposition of $\mathcal{Y} \times \mathcal{Y}$ is done. Nodes in that tree correspond to subsets of $\mathcal{Y}_i \times \mathcal{Y}_j \subset \mathcal{Y} \times \mathcal{Y}$ and thus to sub-blocks of $A_{\phi, \mathcal{Y} \times \mathcal{Y}}$. Based on an *admissibility condition*, these sub-blocks are either identified as close and thus directly evaluated or as far and thus approximated. Approximation is done either using expansions of the kernel function ϕ or using low-rank approximations of the algebraically given matrix sub-block. In this work, low-rank approximations by adaptive cross approximation (ACA) [6] are considered, leading to a purely algebraic approach. A further refinement of \mathcal{H} matrix

techniques leads to \mathcal{H}^2 matrices [24,23,9] that even exhibit $O(N)$ time complexity. Nevertheless, due to a higher algorithmic complexity, we for now stick to the classical \mathcal{H} matrix techniques.

\mathcal{H} matrix techniques speed up the solution process of (1) significantly. Nevertheless, large to huge problem sizes still cannot be tackled using a single processor core or just one workstation with a limited amount of memory. Therefore parallelization of the \mathcal{H} matrix method is crucial. Parallelization of \mathcal{H} matrix methods on *standard processors* (CPUs) is an active research field. Research in this domain ranges from shared-memory to distributed-memory parallel \mathcal{H} matrix implementations on CPUs. The results of this research are a set of parallel \mathcal{H} matrix libraries, which include, but are not limited to $\mathcal{H}\text{-Lib}^{\text{pro}}$ [28,12,26,19], which is rather feature-complete with a shared-memory parallelization and limited distributed-memory support, AHMED (Another software library on hierarchical matrices for elliptic differential equations) [3] and DMHM (Distributed-Memory Hierarchical Matrices) [33] with a distributed-memory parallelization, H2Lib [10] with some support for shared-memory parallelism and work based on the related Hierarchically Semi-Separable (HSS) matrices [36] with the software STRUMPACK [35,18], where the latter one is parallelized for shared- and distributed-memory. Another related, strongly CPU-parallel software for problems of type (1), (2) is PetRBF [43]. In contrast to the above works, we here address parallelization on *many-core processors*.

Many-core processors such as graphics processing units (GPUs) or Intel Xeon Phi reflect recent developments in chip production and high performance computing (HPC): Future parallel computers might show a dramatic growth in the number of parallel processing units with a strong (negative) impact on scalability of current shared-memory and distributed-memory parallelizations. Many-core processors are often assumed to be an optimal testbed for reformulations of classical algorithms towards a massive amount of parallelism, preparing for future parallel computers.

In this work, we will discuss fundamental research on new formulations of standard \mathcal{H} matrix algorithms in order to expose as much parallelism as possible to many-core hardware. Our new algorithms are then implemented on a model hardware, namely GPUs (by NVIDIA). We claim that all of our algorithmic developments equivalently apply to GPU hardware of other vendors or to the Xeon Phi architecture. There is a small set of related work for \mathcal{H} matrices on many-core hardware. In [11], the GPU-acceleration of the quadrature in a \mathcal{H}^2 matrix method for boundary element method problems is considered. Moreover, in [27] many-core parallel LU-factorization for \mathcal{H} matrices is presented and evaluated on a Xeon Phi device. However, these works have in common that many-core hardware is only used as an accelerator or for another computing task, and not as main computing device for the fast matrix-vector product. In contrast, we want to rely completely on many-core parallel hardware for the full \mathcal{H} matrix construction and the \mathcal{H} matrix-vector product.

Other works in the field of many-core hardware concentrating on matrices of type (2) or using other methods are the ASKIT library [30] which uses GPU

acceleration and some very specific tree-based approximation technique and fast multipole methods [42, 2] with e.g. the multi-GPU parallel library **ExaFMM** [42]. While these approaches are very promising for these specific matrices, our main intention is to parallelize the entirely algebraic \mathcal{H} matrix technique, allowing to be used in much more applications.

Fully relying on many-core hardware specifically requires us to provide many-core parallel reformulations of the underlying spatial data structure, the tree construction and traversal, bounding box computations and the construction and evaluation of both the dense matrix parts as well as the low-rank matrix approximations. We propose several algorithmic patterns for many-core processors in context of \mathcal{H} matrices. Space filling curves, i.e. Z order curves, are discussed as parallelized spatial data structure. This goes back to work on the fast construction and evaluation of bounding volume hierarchies on GPUs [29]. We use a parallel formulation of tree traversal using an array-based tree description (cf. [31] for a background on GPU-based tree traversal). Batching or work aggregation, cf. e.g. [15, 1] allows to express parallelism even for code parts in which many similar non-equally sized subtasks are done, strongly optimizing bounding box calculations and low-rank approximations.

As a result of these developments, the author provides an Open Source reference implementation on GPU, which is called **hmglib** [45]. To the best of the authors knowledge, this is the first entirely GPU-based \mathcal{H} matrix library of this kind. For completeness, we should state that there is ongoing research on multi-GPU parallel hierarchical matrices in a library called **KSPARSE** [13], which is however not published and not available for download. Since very recently, there exists a preprint [14] of the authors of [13], discussing the parallel, batched GPU-based implementation of matrix factorizations in context of hierarchical matrices. However, it does not become clear, whether the full algorithm (beyond the batched linear algebra) is performed on GPU. Moreover, the underlying code is not published. Therefore, we still claim that the proposed work is the first available entirely GPU-based \mathcal{H} matrix method.

From a technical point of view, we will show that our many-core parallel model implementation on one GPU outperforms a classical *sequentially* running CPU-based \mathcal{H} matrix library by more than two orders of magnitude in the \mathcal{H} matrix construction and by roughly one order of magnitude for the \mathcal{H} matrix-vector product for a discussed model problem. Nevertheless, our main intention is to show the changes that are to be done to get an entirely many-core parallel implementation. This shall lead to a better understanding and preparation for future intrinsically extremely parallel computing hardware.

Section 2 introduces hierarchical matrices and adaptive cross approximation. Thereafter, Section 3 discusses a simplified programming model for many-core processors. This model allows to formulate many-core parallel programming patterns such as tree traversal or batching of similar sized sub-tasks. These patterns are introduced in Section 4 and applied in Section 5 to provide many-core parallel algorithms for \mathcal{H} matrices. Section 6 treats the reference GPU implementation covering an in-depth benchmark and empirical performance analysis. Finally, Section 7 concludes this work by a short summary.

2 \mathcal{H} matrix background

In the following, we will briefly summarize the necessary algorithmic and mathematical aspects of \mathcal{H} matrices. This overview is partially based on [12]. For further reading see e.g. [21].

Let us start by identifying the points in $\mathcal{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ by their index set $I := \{1, \dots, N\}$. A single entry $\phi(\mathbf{y}_i, \mathbf{y}_j)$ of the system matrix $A_{\phi, \mathcal{Y} \times \mathcal{Y}}$ corresponds to an index tuple (i, j) . Later, we will build *clusters* τ , i.e. specific subsets $\tau \subset I$. We can identify the product of two clusters, e.g. $\tau \times \sigma \subset I \times I$, with a sub-matrix $A_{\phi, \mathcal{Y} \times \mathcal{Y}}|_{\tau \times \sigma}$ of the system matrix $A_{\phi, \mathcal{Y} \times \mathcal{Y}}$. We will need this dual view between sets of index tuples and matrix entries to better understand the basic algorithmic idea of \mathcal{H} matrices.

A kernel function $\phi : \Omega \times \Omega \rightarrow \mathbb{R}$ is called *asymptotically smooth* if there are constants $C_{as1}, C_{as2} \in \mathbb{R}^{>0}$ such that

$$|\partial_x^\alpha \partial_y^\beta \phi(\mathbf{y}, \mathbf{y}')| \leq C_{as1} (C_{as2} \|\mathbf{y} - \mathbf{y}'\|)^{-|\alpha| - |\beta|} \boldsymbol{\alpha} + \boldsymbol{\beta} |\phi(\mathbf{y}, \mathbf{y}')|$$

for all $\mathbf{y}, \mathbf{y}' \in \Omega$ with $\mathbf{y} \neq \mathbf{y}'$ and all multi-indices $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{N}_0^d$. Fixing $\mathbf{y} \in \Omega$, the kernel evaluation $\phi(\mathbf{y}, \mathbf{y}_{far})$ of an approximately smooth kernel function can be approximated with a controlled, small error, in case the point \mathbf{y}_{far} is *far* away from \mathbf{y} . In the \mathcal{H} matrix approach, an *admissibility* condition identifies matrix blocks $A_{\phi, \mathcal{Y} \times \mathcal{Y}}|_{\tau \times \sigma}$ that represent interactions of points with indices τ that are *far* away from points with indices σ . Admissible matrix blocks are traditionally approximated via series expansions of kernel ϕ . We here consider the well-known alternative approach to approximate the matrix blocks by low-rank approximations as e.g. in [6].

2.1 Clustering

The *cluster tree* $\mathcal{T}_I = (V_I, \gamma, \mu)$ is a hierarchical spatial data structure on I (or \mathcal{Y}). V_I is the set of nodes in the tree, γ a mapping $\gamma : V_I \rightarrow \mathcal{P}(V_I)$ of the nodes to their children and $\mu : V_I \rightarrow \mathcal{P}(I)$ a mapping of the nodes to their value. Here, the *value* of each node is a cluster in I , i.e. a subset of I . A cluster tree has to fulfill

- (C1) $\mu(v) \in \mathcal{P}(I) \setminus \{\emptyset\}$, for all $v \in V_I$,
- (C2) $\mu(\text{root}(\mathcal{T})) = I$,
- (C3) if $v \in V_I$ is a leaf, i.e. $\gamma(v) = \emptyset$, then $|\mu(v)| \leq C_{leaf}$ and
- (C4) if $v \in V_I$ is no leaf, then it has exactly two sons $\gamma(v) = \{v_1, v_2\}$ and $\mu(v) = \mu(v_1) \cup \mu(v_2)$.

Thereby, the cluster tree divides the full set I (C2) into a hierarchy of clusters, where non-empty clusters of I (represented by nodes in \mathcal{T}_I , C1) are disjointly partitioned into two smaller clusters (C4). In case a cluster is no longer partitioned (thus represented by a leaf), its size is bound from above by C_{leaf} (C3).

In cardinality-based clustering (CBC) [12], an algorithm to create the cluster tree decomposes the sets $\tau = \gamma(v)$ such that the subsets in the child nodes of v have similar size. Moreover, the subsets shall build geometrically distinct clusters. A CBC based on space filling curves will be introduced in Section 4.4. The splitting in the cluster tree construction is continued as long as $|\tau| > C_{leaf}$.

2.2 Bounding box admissibility

In this work, we will restrict ourselves to an admissibility condition based on bounding boxes for clusters. Other choices are possible [21]. For a cluster $\tau \subset I$, the bounding box Q_τ is given as

$$Q_\tau := \prod_{i=1}^d [a_\tau^{(i)}, b_\tau^{(i)}]$$

with $a_\tau^{(i)} := \min_{j \in \tau} y_j^{(i)}$, $b_\tau^{(i)} := \max_{j \in \tau} y_j^{(i)}$ and $\mathbf{y}_j := (y_j^{(1)}, \dots, y_j^{(d)})^\top$. One admissibility condition for an index block $\tau \times \sigma \subset I \times I$ is

$$\min \{ \text{diam}(Q_\tau), \text{diam}(Q_\sigma) \} \leq \eta \text{dist}(Q_\tau, Q_\sigma) \quad (3)$$

with $\eta \in \mathbb{R}^{\geq 0}$ a parameter balancing convergence and algorithmic complexity. Diameter $\text{diam}(Q_\tau)$ and distance $\text{dist}(Q_\tau, Q_\sigma)$ of bounding boxes are defined by

$$\begin{aligned} \text{diam}(Q_\tau) &:= \left(\sum_{i=1}^d (b_\tau^{(i)} - a_\tau^{(i)})^2 \right)^{1/2}, \\ \text{dist}(Q_\tau, Q_\sigma) &:= \left(\sum_{i=1}^d \left(\max \{ 0, a_\tau^{(i)} - b_\sigma^{(i)} \}^2 + \max \{ 0, a_\sigma^{(i)} - b_\tau^{(i)} \}^2 \right) \right)^{1/2}. \end{aligned}$$

2.3 Block cluster tree

A hierarchy over *blocks* $\tau \times \sigma \subset I \times I$ is induced by the *block cluster tree* $\mathcal{T}_{I \times I} = (V_{I \times I}, \gamma, \mu)$, with γ the child node map and $\mu : V_{I \times I} \rightarrow \mathcal{P}(I \times I)$ the map of nodes to their values, i.e. blocks. Note that we re-use here the same notation (γ, μ) as for the cluster tree. Algorithm 1 implicitly defines the block cluster tree. For given cluster tree nodes v_1, v_2 (corresponding to clusters $\mu(v_1) = \tau, \mu(v_2) = \sigma$), a block cluster tree node w with $\mu(w) := \mu(v_1) \times \mu(v_2)$ (corresponding to $\mu(w) = \tau \times \sigma$) and parameter C_{leaf} , this algorithm recursively constructs a block cluster tree. Procedure `build_block_cluster_tree` is initially launched with v_1 and v_2 each being a root of the cluster tree \mathcal{T}_I and node w is initialized to represent the index block $I \times I$. By construction, the leaves of $\mathcal{T}_{I \times I}$, namely $\mathcal{L}_{I \times I} := \{w \in V_{I \times I} \mid \gamma(w) = \emptyset\}$, correspond to index blocks that form a partition of $I \times I$.

Algorithm 1 Algorithm to build a block cluster tree

```

procedure BUILD_BLOCK_CLUSTER_TREE( $v_1, v_2, w, C_{leaf}$ )
  ( $\tau, \sigma$ )  $\leftarrow$  ( $\mu(v_1), \mu(v_2)$ )
  if  $\tau \times \sigma$  is not admissible and  $|\tau| > C_{leaf}$  and  $|\sigma| > C_{leaf}$  then
     $\gamma(w) \leftarrow \emptyset$ 
    for  $v'_1 \in \gamma(v_1)$  do  $\triangleright$  Loop over all combinations of children in both cluster trees.
      for  $v'_2 \in \gamma(v_2)$  do
         $\mu(w') \leftarrow \mu(v'_1) \times \mu(v'_2)$   $\triangleright$  Set block cluster of new node  $w'$ .
         $\gamma(w) \leftarrow \gamma(w) \cup \{w'\}$   $\triangleright$  Add new node to children of  $w$ .
        BUILD_BLOCK_CLUSTER_TREE( $v'_1, v'_2, w', C_{leaf}$ )
      end for
    end for
  else
     $\gamma(w) \leftarrow \emptyset$   $\triangleright$  No child nodes are created, i.e.  $w$  becomes a leaf.
  end if
end procedure

```

2.4 $\mathbf{R}k$ -matrices and adaptive cross approximation

If a node w in a block cluster tree corresponds to an index block $\tau \times \sigma \subset I \times I$ that is admissible, the corresponding sub-matrix $A_{\phi, \mathcal{Y} \times \mathcal{Y}}|_{\tau \times \sigma} \in \mathbb{R}^{|\tau| \times |\sigma|}$ is replaced by an $\mathbf{R}k$ matrix $R_{\tau \times \sigma} \in \mathbb{R}^{|\tau| \times |\sigma|}$. An $\mathbf{R}k$ matrix $R_{\tau \times \sigma}$ is given as

$$R_{\tau \times \sigma} = U_{\tau \times \sigma} V_{\tau \times \sigma}^\top, \quad U_{\tau \times \sigma} \in \mathbb{R}^{|\tau| \times k}, V_{\tau \times \sigma} \in \mathbb{R}^{|\sigma| \times k},$$

that is, it has a maximum rank of k . Moreover, using $U_{\tau \times \sigma}$ and $V_{\tau \times \sigma}$, a matrix-vector product involving $R_{\tau \times \sigma}$ can be computed in $O(r \cdot (|\tau| + |\sigma|))$ operations.

While there are many (problem-dependent) ways to approximate $A_{\phi, \mathcal{Y} \times \mathcal{Y}}|_{\tau \times \sigma}$, we here aim at using a purely algebraic low-rank approximation method to derive $R_{\tau \times \sigma}$. Our method of choice is the *adaptive cross approximation* (ACA) [6, 5]. This method builds a low-rank approximation by an iterative rank-one update process that is terminated based on the error ϵ in the Frobenius norm $\|\cdot\|_F$.

One version of adaptive cross approximation is given in Algorithm 2. It follows the lines of [5]. The algorithm computes for a general matrix $A \in \mathbb{R}^{m \times n}$ and error threshold ϵ matrices $U \in \mathbb{R}^{m \times k_{max}}$, $V \in \mathbb{R}^{n \times k_{max}}$ such that $A \approx UV^\top$. In case the algorithm terminates due to the stopping criterion, k_{max} becomes the (adaptively computed) rank such that $\|A - UV^\top\|_F \leq \epsilon$. Otherwise, the maximum rank of k_{max} is hit. The choice of a column pivot index j_r is strongly problem-dependent. For simplicity, we choose j_r such that $\|\hat{\mathbf{u}}_r\|_2 > \epsilon_0$ for small ϵ in the range of machine precision. In our practical implementation, we will, however, avoid to evaluate the stopping criterion and will only impose the maximum rank k_{max} . As we will see in Section 6.4, k_{max} can be chosen rather small due to the exponential convergence of ACA for appropriate kernel functions ϕ . For more details on ACA, see [5, 6].

Algorithm 2 Adaptive cross approximation (ACA) for $A \in \mathbb{R}^{m \times n}$ [5][6]

```

function COMPUTE_ADAPTIVE_CROSS_APPROXIMATION( $A, \epsilon$ )
   $k_{max} \leftarrow k$ 
  for  $r = 1, 2, \dots, k$  do
     $\hat{\mathbf{u}}_r = A_{1:m, j_r} - \sum_{l=1}^{r-1} \mathbf{u}_l (\mathbf{v}_l)_{j_r}$ ,  $\triangleright$  Col. index  $j_r$  depending on implementation
     $\mathbf{u}_r = (\hat{\mathbf{u}}_r)_{i_r}^{-1} \hat{\mathbf{u}}_r$ , with  $|(\hat{\mathbf{u}}_r)_{i_r}| = \|\hat{\mathbf{u}}_r\|_\infty$   $\triangleright$  Row index  $i_r$  given as pivot position
     $\mathbf{v}_r = (A_{i_r, 1:n})^\top - \sum_{l=1}^{r-1} (\mathbf{u}_l)_{i_r} \mathbf{v}_l$ 
    if  $(\|\mathbf{u}_r\|_2 \|\mathbf{v}_r\|_2 \leq \frac{\epsilon(1.0-\eta)}{1.0+\epsilon} \|\sum_{l=1}^r \mathbf{u}_l \mathbf{v}_l\|_F)$  then  $\triangleright$  Stopping criterion
       $k_{max} \leftarrow r$   $\triangleright k_{max}$  is adaptively found rank
      stop loop
    end if
  end for
   $U \leftarrow (\mathbf{u}_1, \dots, \mathbf{u}_{k_{max}})$ 
   $V \leftarrow (\mathbf{v}_1, \dots, \mathbf{v}_{k_{max}})$ 
  return  $U, V$ 
end function

```

Algorithm 3 Matrix-vector product with an \mathcal{H} matrix $L \in \mathbb{R}^{I \times I}$

```

function MATRIX_VECTOR_PRODUCT( $L, w, x, z$ )
  if  $\gamma(w) \neq \emptyset$  then
    for  $w' \in \gamma(w)$  do
      MATRIX_VECTOR_PRODUCT( $L, w', x, z$ )
    end for
  else
     $\tau \times \sigma \leftarrow \mu(w)$ 
    if  $\tau \times \sigma$  is admissible then
       $\mathbf{t} \leftarrow V_{\tau \times \sigma}^\top \mathbf{x}|_\tau$ 
       $\mathbf{z}|_\tau \leftarrow \mathbf{z}|_\tau + U_{\tau \times \sigma} \mathbf{t}$ 
    else
       $\mathbf{z}|_\tau \leftarrow \mathbf{z}|_\tau + L|_{\tau \times \sigma} \mathbf{x}|_\sigma$ 
    end if
  end if
  return  $z$ 
end function

```

2.5 \mathcal{H} -matrices and their matrix-vector product

Formally, a general matrix $L \in \mathbb{R}^{|I| \times |I|}$ is — for fixed $k \in \mathbb{N}$ and block cluster tree $\mathcal{T}_{I \times I}$ — called \mathcal{H} matrix of blockwise rank k , if

$$\text{rank}(L|_{\tau \times \sigma}) \leq k$$

for all index blocks $\tau \times \sigma$ in admissible leaves. The operation to transform an existing dense matrix, e.g. $A_{\phi, \mathcal{Y} \times \mathcal{Y}}$, to \mathcal{H} matrix form is called *truncation*. It involves the introduction of a cluster tree \mathcal{T}_I , a block cluster tree $\mathcal{T}_{I \times I}$ and the computation of a low-rank approximation of matrix blocks corresponding to admissible leaves.

The (fast) matrix-vector product of an \mathcal{H} matrix $L \in \mathbb{R}^{|I| \times |I|}$ with a vector $\mathbf{x} \in \mathbb{R}^{|I|}$, that is, the efficient evaluation of

$$\mathbf{z} := \mathbf{z} + L\mathbf{x},$$

is summarized in Algorithm 3. The algorithm recursively traverses the block cluster tree for an initially given (root) node w and applies a low-rank matrix-vector product for admissible blocks and the full dense matrix for non-admissible blocks. If we launch `MATRIX_VECTOR_PRODUCT` with w corresponding to $I \times I$ and L being the truncated version of $A_{\phi, \mathcal{Y} \times \mathcal{Y}}$, it can be shown that the algorithm has a complexity of $O(k \cdot N \log N)$ [21].

3 Programming model for many-core parallel algorithms

In this Section, we introduce the terminology to describe efficient and scalable parallel many-core algorithms. Note that, to the best of the author’s knowledge, a common abstract programming model for many-core architectures is still missing. Therefore, algorithmic work on GPUs or Xeon Phi often addresses many details of these architectures. In contrast, we use a strongly simplified programming model, avoiding most of the technical details of classical many-core literature. Our model is based on two observations. First, a crucial part of a lot of many-core parallel algorithms requires almost no interaction between the involved parallel compute units, that is, they are close to *embarrassingly* parallel. Second, vendors (or enthusiasts) provide extremely efficient many-core parallel implementations of base algorithms (reductions, scan operations, etc.) for more complex parallel algorithmic patterns. Therefore, we claim that we can build all algorithms of interest by combinations of almost embarrassingly parallel kernels and standardized parallel algorithms. They are defined in more detail in the following paragraphs.

3.1 Almost embarrassingly parallel kernels

The kind of compute kernels we discuss here are strongly related to the *bulk synchronous parallel model*, cf. [38]: We introduce an (in principle) infinite number of virtual parallel threads. In each parallel thread, the same piece of sequential code is executed. Different memory accesses / execution paths are realized by a thread index which is associated to each thread.

All threads are aggregated in a *kernel*, which gets the number of threads to execute at launch time. The kernel terminates when all threads have stopped the execution of the sequential code. The sequential code (per thread) can either use *local memory*, which can only be read by that single thread, or *global memory*, which is available to all threads. At the end of the kernel execution, all local memory data is lost while global memory entries remain available. Whenever a single thread writes to a given global memory entry, read or write operations on that memory entry (by another thread) are invalid / prohibited. Reading (without writing) from a common global memory location by multiple threads in one kernel is possible.

One exception to the write rule is available in case of *atomic* operations (usually *atomic_add* or *atomic_compare_and_swap*) on global memory. Atomic

operations issued by different threads on one common global memory location are all correctly executed, even if this means that threads get serialized. However, the ordering of the execution is not assured. Therefore, atomic operations are only useful in very few cases (e.g. counters).

Note that the actual mapping of threads to hardware processing units is not part of the model. This especially allows to define parallel programs independent of the number of available hardware threads. Moreover, the beforehand given definition of computing kernels does not give any hints towards the performance of their actual mapping to a given hardware platform. Let us give examples for GPUs. Here, global memory accesses are fast if they are done consecutively for consecutive thread indices, that is, threads $0, 1, 2, 3, \dots$ access memory entries $e, e + 1, e + 2, e + 3, \dots$. In contrast, random access has rather low performance. Moreover, conditionals in the thread-sequential code of a kernel might have a severe impact on performance on GPUs if thread execution paths diverge. Other architectures might have similar limitations.

3.2 Standardized parallel algorithms

As second ingredient to our many-core parallel algorithms, we expect to have access to a parallel library of standardized (many-core parallel) algorithms similar to the C++ Standard Template Library (STL) algorithms library. We e.g. need `reduce`, `stable_sort`, `scan`, ... These algorithms are expected to be realized as one function call that is executed on data in global memory. The many-core parallel implementation of these algorithms is assumed to be extremely optimized and given e.g. by the vendor. On GPUs an implementation of STL-like algorithms is available via the `Thrust` library [7]. Alternatives include, but are not limited to `ArrayFire` [41] (supporting GPUs, CPUs and Xeon Phi) and `Boost.Compute` [37] (supporting multi-core CPUs and GPUs). In addition, we assume to have appropriate BLAS libraries for a given many-core device.

4 Many-core parallel programming patterns for \mathcal{H} matrices

As motivated before, we introduce in the following a set of parallel programming patterns that are necessary for algorithms based on \mathcal{H} matrices.

4.1 Parallel tree traversal

In the following, we introduce a fully parallel tree traversal algorithm, which is related to ideas in [29, 31]. It on-the-fly builds and traverses a tree. The tree traversal algorithm is given in Algorithm 4. An input tree $\mathcal{T} = (V, \gamma, \mu)$ is assumed to have height $h(\mathcal{T})$, levels $l \in \{0, \dots, h(\mathcal{T})\}$ and nodes $v \in V$ of arbitrary order. The algorithm is designed such that we only store nodes $V(l) := \{v \in V \mid \text{level}(v) = l\}$ and $V(l + 1)$ for two consecutive levels l and

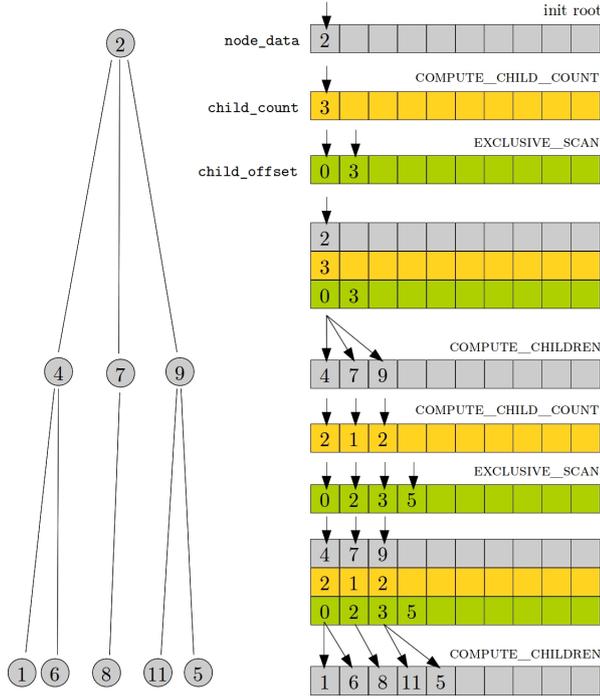


Fig. 1 The many-core parallel tree traversal parallelizes over the nodes on a given level of the tree. Our algorithm is here exemplified for nodes containing numbers.

$l + 1$. All other data is created level-wise and discarded after a new level has been successfully created. The nodes $v \in V(l)$ and $v' \in V(l + 1)$ are stored in global arrays `node_data_old` and `node_data`. In addition, we need for level $l + 1$ the number of children per node $|\gamma(v')|$ (stored in `child_count`) and the offset of the data of the child nodes (`child_offset`). Figure 1 illustrates these arrays.

The algorithm works as follows: Let us assume for now that the arrays per level can have arbitrary size and that we are on level $0 \leq l < h(\mathcal{T})$ and the only available data is the node data. We first invoke a kernel `compute_child_count` with the number of threads equal to the number $|V(l)|$ of nodes on that level, thus the number of valid entries in the `node_data` array. In each thread, we independently compute for each node $v \in V$ (based on the node data) the number of children $|\gamma(v)|$ that shall be created on the next level. This computation is problem-dependent. In case of the cluster tree, it e.g. holds $|\gamma(v)| \in \{0, 2\}$. The number of children is stored at the same offset in the array as the given node data. In a next step, we have to compute the offsets for the node data on the next level, i.e. `child_offset`. This can be done by an `exclusive_scan` operation initialized to 0. The entries of `child_offset` then become $[0, \text{child_count}[0], \text{child_count}[0] + \text{child_count}[1], \dots]$. The

Algorithm 4 Many-core parallel tree traversal

```

procedure TRAVERSE(root)           ▷ Traverse a tree with given root data
  allocate node_data, node_data_old, child_count, child_offset
  node_data[0] ← root_data
  l ← 0
   $|V(l)| \leftarrow 1$ 
  while  $|V(l)| > 0$  do           ▷ Handle tree levels as long a there are nodes
    COMPUTE_CHILD_COUNT $\langle |V(l)| \rangle$ (child_count, node_data)
    ▷ Problem-dependent kernel to compute the number of children per node
    EXCLUSIVE_SCAN(child_offset, child_count, 0,  $|V(l)|$ )
     $|V(l+1)| \leftarrow \text{child\_offset}[|V(l)|]$    ▷ Set total number of children of next level
    node_data_old ← node_data
    COMPUTE_CHILDREN $\langle |V(l)| \rangle$ (node_data, node_data_old, child_count, child_offset)
    ▷ Problem-dependent kernel to compute the content of the children

    l ← l + 1
  end while
end procedure

```

output of the scan operation contains as additional number (at the end of the set of valid entries) the total number of children $|V(l+1)|$. The last step on level l is the creation of the node data $V(l+1)$ on level $l+1$. This is again done using a kernel with the number of threads equal to $|V(l)|$. Each thread then independently computes the new entries taking the storage location in `node_data` for level $l+1$ from `child_offset`. This finishes the computation for one level. The whole process is iteratively proceeded over all levels $0 \leq l < h(\mathcal{T})$. To start the tree traversal on level zero, i.e. the root of the tree, the `node_data` array is initialized with a single entry. A full example of a tree traversal is given in Fig. 1 and the algorithm is stated in Algorithm 4.

We next have to discuss how to deal with the array allocation, knowing that the required size of the storage arrays differs between the tree levels. Here, we have two options. The first option would be to pre-allocate the arrays to a fixed size $\max_{l \in \{0, \dots, \text{height}(\mathcal{T})\}} |V(l)|$. This, of course, requires to know this number beforehand or to have a suitable upper bound for it. Very often, this is not the case. The second option is a *dynamic* allocation of the array size for the next level. This size can be predicted based on the information in the `child_count` array. In case a reallocation of memory is a very expensive operation on a given target architecture, one could also apply hybrid approaches such as adapting the size of the arrays only if a given array (of large size) would be too small for the next level. In our implementation on GPU, a global reallocation of the memory is a very efficient operation. This is why we have chosen to use the dynamic allocation approach.

Finally, we should have a look at the properties of the algorithm in terms of the use of the many-core processor. It becomes obvious that the number of utilized parallel threads on the first few levels is very low. That is, the proposed algorithm makes no full use of the many-core processor on the first levels. This might become an issue if many tree traversals on small trees are considered and if the tree traversal operation itself is the dominant operation in an application. However, both is not the case in our application: The trees are very large and,

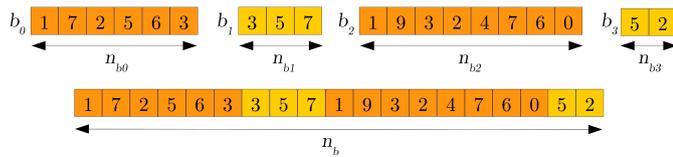


Fig. 2 By batching individual subproblems into one big array, it becomes possible to utilize a many-core processor much better.

as we will see in Section 6, the tree traversal operation makes only a very small fraction of the overall \mathcal{H} matrix setup / application process. Therefore, we consider our tree traversal method efficient enough for our needs. In case higher utilization of the many-core processor is needed, efficient solutions become very architecture-specific. In case of GPUs, there is work on tree traversal by work queues [17], which, however, makes explicit use of knowledge on the hardware and which somehow even breaks the programming model initially considered for GPUs.

4.2 Batching many similar non-equally sized compute tasks

We next want to discuss how to make optimal use of a many-core processor in case there is an identical computing task which shall be applied to m different, non-equally sized arrays b_0, b_1, \dots, b_{m-1} of sizes $n_{b_0}, n_{b_1}, \dots, n_{b_{m-1}}$. Figure 2 gives an example of such arrays. The easiest way to consider a parallelization on many-core hardware would be to loop over all arrays b_i and to perform the necessary many-core parallel operations individually to each array. This is efficient as long as the many-core processor is sufficiently utilized. However, we here consider arrays of changing and usually small size. In this case, a major part of the many-core processor is not used. Therefore we propose to use the technique of *batching* of the necessary computations, cf. [15, 1], in order to use the full processor while speeding up the calculation.

The first step in batching is to put all sub-arrays or *batches* b_i consecutively in a *batched array* of size $n_b := \sum_i n_{b_i}$, cf. Fig. 2. We next have to distinguish between *transformation* operations and *reduction* operations on that batched array. A transformation on each batch applies changes individually to each entry of each batch, i.e. there is no interaction between the data entries. Applying a transformation to each batch is therefore equivalent to applying the same transform to the full batched array. Therefore, in case of transformations, we apply one operation to the full batched array.

In contrast, reduction operations (such as sum, minimum, maximum, norm, etc.) require the interaction of all entries within a batch. Therefore, we need a different strategy. The STL-type algorithm `reduce_by_key` is applied to the full batched array and computes, in parallel, batch-wise reductions. The action of the method is shown in Fig. 3 for a maximum reduction operation. We introduce a `keys` array of integer values. A series of identical numbers in

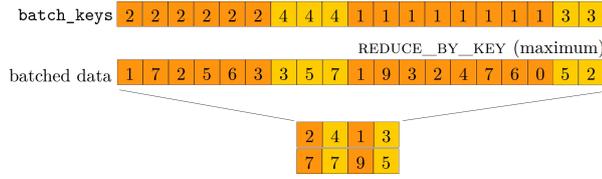


Fig. 3 Reduction operations (e.g. maximum computations) for several sub-problems can be handled in parallel by a `reduce_by_key` operation, where identical consecutive entries in `batch_keys` mark the individual sub-problems.

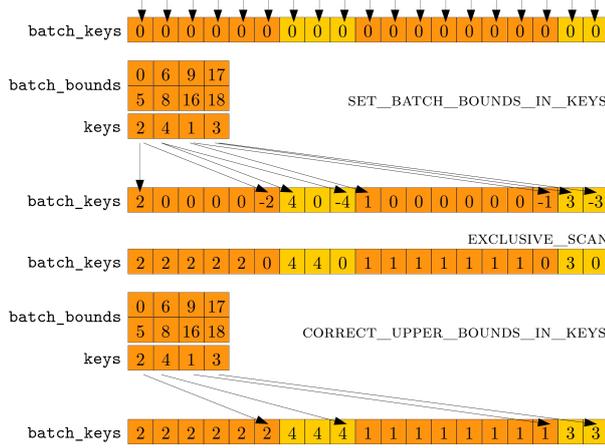


Fig. 4 The construction of an array of keys for batching involves marking boundaries of the batches and an `EXCLUSIVE_SCAN` operation.

Algorithm 5 Many-core parallel key generation for batching

```

procedure CREATE_KEYS(batch_bounds, batch_keys,  $n_b$ ,  $m$ )
  INIT $\langle n_b \rangle$ (keys, 0)
  SET_BATCH_BOUNDS_IN_KEYS $\langle m \rangle$ (keys, batch_bounds, batch_key)
  EXCLUSIVE_SCAN(keys, keys, 0,  $n_b$ )  $\triangleright$  Write exclusive scan on keys to keys
  CORRECT_UPPER_BOUNDS_IN_KEYS $\langle m \rangle$ (keys, batch_bounds, batch_key)
  return keys
end procedure

```

the `keys` array marks one batch. The method `reduce_by_key` then applies the reduction operation per subset and builds up a small array of size m containing the reduction results and the keys reduced to a single number.

To compute the keys, we need an additional parallel algorithm, cf. Algorithm 5. It takes an array of boundaries (`batch_bounds`) of each batch b_i and an array (`batch_keys`) of keys k_{b_i} per batch as input. The procedure to create keys for batching is exemplified in Fig. 4. We initialize (by a kernel of n_b threads) the `keys` array to zeros. Then, the kernel `SET_BATCH_BOUNDS_IN_KEYS` of m threads is invoked, where each thread independently writes the key k_{b_i} and the

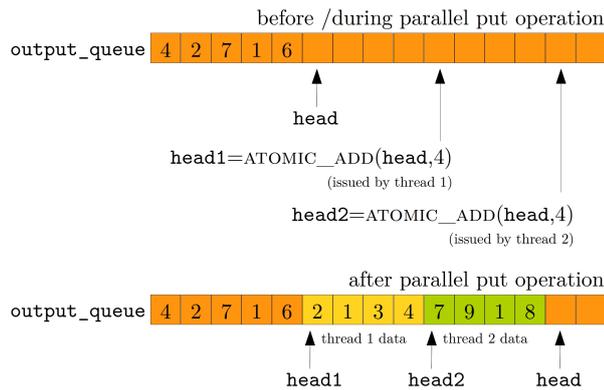


Fig. 5 By the use of atomic operations it is easily possible to create a write-only parallel output queue. In the above example, two threads concurrently add four entries to the head of the queue.

negative key $-k_{b_i}$ to the lower and upper bound of each batch in the batched array, cf. Fig. 4. Then an `exclusive_scan` operation (adding elements) is executed on the full batched array. This sets the correct keys almost everywhere, except at the upper boundary of each batch. Therefore a second kernel of m threads is invoked to correct the upper bounds of each batch b_i to k_{b_i} .

In some cases, the size n_b of the batched array is too large to be kept in the memory of the many-core processor. Such cases can be handled by appropriately partitioning the batches b_i into subsets of batches which are then handled as before.

A crucial property of the approach presented here is its independence of the size and the number of batches. The runtime for this approach is almost constant with the size n_b . This is a strong advantage over strategies that directly rely on the use of the different parallelization hierarchies (*thread blocks*, *grids* on GPUs and *vectorization*, *shared-memory parallelism*, etc. on Xeon Phi).

4.3 Parallel output queues

In some cases, we need to create what we define as *write-only parallel output queues*. Such queues can only be *filled* (in parallel). Removal of data or reading the head of the queue during the enqueueing process is not required. Instead, the stored queue data is handled as one array as post-processing step. As an example for such a queue, let us consider a parallel tree traversal in which (unordered) tasks are created in each leaf. Instead of executing the task during the tree traversal, we can, in parallel, put them in a queue. The actual execution of the tasks can be issued afterwards as new parallel operation.

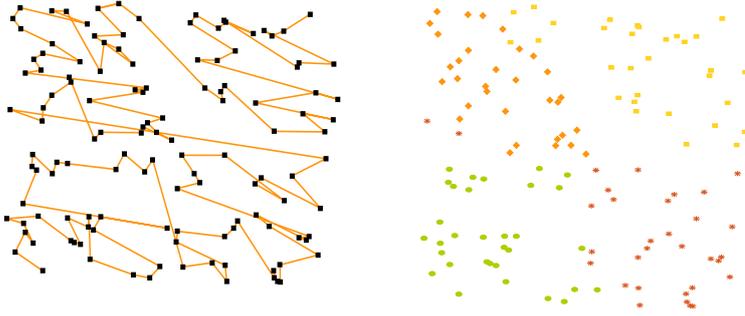


Fig. 6 Left: By sorting a set of arbitrary points following their Morton codes, a spatial data structure is imposed. Right: Dividing the set of ordered points in to equally sized subsets implicitly creates clusters.

The implementation of our parallel output queue, relies on an underlying global memory output array of appropriate size. If we cannot predict the size, we can also apply dynamic memory allocation approaches, as above. We store a pointer to the head and the tail of the queue in global memory. Whenever a put operation is issued in a thread of a kernel, the head pointer is moved accordingly by an atomic operation while storing the old head in the same operation. The old head is used as output address to write the data in the queue. Figure 5 summarizes and exemplifies the approach.

4.4 Spatial data structure by Z-order curves

We use a Z order space filling curve [32] to introduce a spatial data structure on top of the input point set \mathcal{Y} . This idea is based on [29]. The core idea is to assign each point in \mathcal{Y} a *Morton code*, which is an integer value. By ordering the elements of \mathcal{Y} following their Morton codes, two consecutively ordered points get spatially close to each other, cf. Fig. 6. The implicit spatial structure introduced by the Morton ordering strongly simplifies the construction of the cluster tree. Whenever we have to split up a given cluster into two spatially distinct clusters in cardinality-based clustering, we only have to divide a given ordered point array into two parts, i.e. the first half of the elements builds the first subset and the second half of the elements builds the second subsets. That is, spatial operations get reduced to array operations.

Our implementation follows the lines of [29]. We here assume that the reader has some knowledge about the construction of Morton codes. For details, see e.g. [8]. It is trivially parallel to compute Morton codes for a point set. Algorithm 6 summarizes the corresponding parallel kernel `COMPUTE_MORTON_CODES`. Per parallel thread / point coordinate, it iterates over the dimensions of the point coordinates, where it transforms the

Algorithm 6 Computation of Morton codes

```

procedure COMPUTE_MORTON_CODES( $\mathcal{Y}$ )(>(coords)
  for each thread  $t = 0, \dots, |\mathcal{Y}| - 1$  in parallel do
    current_code  $\leftarrow 0$ 
    for  $i = 1, \dots, d$  do
      code_current_dim  $\leftarrow$  COMPUTE_FIXED_POINT_REPRESENTATION(coords[i][t])
      code_current_dim  $\leftarrow$  STRETCH_BITS(code_current_dim, i, d)
      current_code  $\leftarrow$  INTERLEAVE(code_current_dim, current_code, i)
    end for
    morton_codes[t]  $\leftarrow$  current_code
  end for
  return morton_codes
end procedure

```

floating-point representation of the coordinate entry to a fixed-point representation. Next, the bits of the fixed-point representation are stretched. Finally, the stretched bits are interleaved dimension-wise such that the final Morton code is constructed. Sorting the points following their Morton codes is an operation of log-linear complexity for which we assume to have an STL-like operation, cf. Section 3.2.

5 Many-core algorithms for \mathcal{H} matrices

In the following, we use the beforehand discussed general parallel algorithmic patterns to construct algorithms for the many-core parallel construction of \mathcal{H} matrices and the \mathcal{H} matrix-vector product.

5.1 Data structures

We collect the points \mathcal{Y} in instances of a struct `point_set`. The struct contains a multi-dimensional array `coords` of coordinates, the dimension of the points and the number of points $|\mathcal{Y}|$. The ordering of the point coordinates in array `coords` follows the Morton order of \mathcal{Y} , cf. Section 4.4. Note that, since the data structure is constructed following the Morton order while the vector \mathbf{x} involved in the \mathcal{H} matrix-vector product is stored following the original point ordering, we have to permute the vector \mathbf{x} in the matrix-vector product or once at the beginning.

As described in Section 2, the \mathcal{H} matrix method strongly relies on sub-blocks $A_{\phi, \mathcal{Y} \times \mathcal{Y}}|_{\tau \times \sigma}$ of matrix $A_{\phi, \mathcal{Y} \times \mathcal{Y}}$, which are constructed over index blocks $\tau \times \sigma \subset I \times I$. As we will see, clusters $\tau \subset I$ will always correspond to points which are (by Morton ordering) consecutively stored in `coords`. Therefore, we can define τ by index ranges $\{i_{l,\cdot}, i_{l,\cdot} + 1, i_{l,\cdot} + 2, \dots, i_{u,\cdot}\}$ pointing to the storage location in `coords`. That is, each cluster τ is represented just by the lower and upper index bounds $i_{l,\cdot}$ and $i_{u,\cdot}$.

In our implementation, we collect the nodes $w \in V_{I \times I}$ of the block cluster tree $\mathcal{T}_{I \times I}$ in instances of structs `work_item`. In addition to the lower and upper

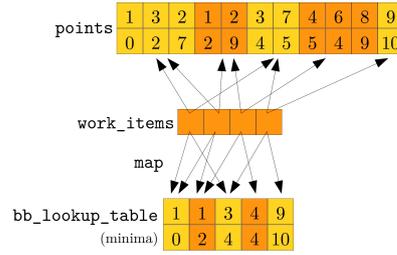


Fig. 7 The boundary box computation is sped up by computing boundary boxes of each subset once. They are stored in `bb_lookup_table` and accessed via a map between work items and the lookup table.

index bounds for sets τ and σ , this struct defines storage for bounding boxes for the points corresponding to clusters τ , σ and an admissibility flag.

5.2 Block cluster tree traversal

The construction and traversal of the block cluster tree is based on a modified version of the tree traversal procedure presented in Algorithm 4. Each node $w \in V_{I \times I}$ is an instance of a struct `work_item`, cf. Section 5.1. The root node is initialized to the set $I \times I$. Before computing the number of children via `COMPUTE_CHILD_COUNT`, we compute the bounding box lookup table and the map to the bounding box lookup table, cf. Section 5.3. A special instance of the `COMPUTE_CHILD_COUNT` method evaluates the admissibility condition (3) using the precomputed bounding boxes and writes the number of children according to that result. The generic `COMPUTE_CHILDREN` method is replaced by a method that – depending on the admissibility condition – either creates new children by splitting up the index sets corresponding to each cluster τ or puts the node as admissible or non-admissible leave node to a parallel work queue `work_queue` of `work_item` structs, cf. Section 4.3.

5.3 Batched bounding box computation

As part of the traversal of the block cluster tree, we have to evaluate the admissibility condition (3) for index blocks $\tau \times \sigma$ involving the bounding boxes of τ and σ in each node. In the following, we will discuss an algorithm to concurrently compute the bounding boxes for clusters τ , σ in all nodes on a given level l of the cluster tree. The algorithm is based on batching, cf. Section 4.2.

We collect the set of nodes on a level l of the block cluster tree, i.e. $V_{I \times I}(l)$, in the array `node_data` of length $|V_{I \times I}(l)|$ composed of structs `work_item` and have the input points \mathcal{Y} in an instance of struct `point_set`, cf. Section 4.4. As simplification, we only consider the concurrent computation of the bounding boxes for one cluster set, e.g. τ , in each node.

Algorithm 7 Compute bounding box lookup table

```

procedure COMPUTE_BOUNDING_BOX_LOOKUP_TABLE(node_data, coords, l,  $|V_{I \times I}(l)|$ )
  (lower_bounds, upper_bounds)  $\leftarrow$  GET_INDEX_BOUNDS(node_data)
  STABLE_SORT(lower_bounds)
  STABLE_SORT(upper_bounds)
  UNIQUE(lower_bounds, unique_lower_bounds)
  UNIQUE(upper_bounds, unique_upper_bounds)
  lookup_table_size  $\leftarrow$   $|lower\_bounds|$ 
  batch_bounds  $\leftarrow$  (unique_lower_bounds, unique_upper_bounds)
  SEQUENCE(unique_set_indices, lookup_table_size, 1)
  INIT $\langle |\mathcal{Y}| \rangle$ (batch_keys, 0)
  batch_keys  $\leftarrow$  CREATE_KEYS(batch_bounds, unique_set_indices,  $|\mathcal{Y}|$ , lookup_table_size)
  (coord_maxima, output_keys)  $\leftarrow$  REDUCE_BY_KEY(coords, batch_keys, maximum)
  (coord_minima, output_keys)  $\leftarrow$  REDUCE_BY_KEY(coords, batch_keys, minimum)
  REMOVE_BY_KEY(coord_maxima, output_keys, 0)  $\triangleright$  Remove invalid compute results
  REMOVE_BY_KEY(coord_minima, output_keys, 0)
  bb_lookup_table  $\leftarrow$  (coord_minima, coord_maxima)
  return bb_lookup_table
end procedure

```

By construction, many nodes $w \in V_{I \times I}(l)$, i.e. on the same level of the block cluster tree, contain identical clusters (not blocks), we e.g. have $\tau_1 \times \sigma_1 = \gamma(w_1)$, $\tau_2 \times \sigma_2 = \gamma(w_2)$ while $\tau_1 = \tau_2$. Therefore, we first identify the set of unique clusters. We then create a lookup table `bb_lookup_table` storing for each unique cluster the bounding box information. In addition, we need a map from a node in `node_data` to the entry in the lookup table. Figure 7 exemplifies this idea.

Algorithm 7 describes our approach to compute the entries of the lookup table `bb_lookup_table`. Function `COMPUTE_BOUNDING_BOX_LOOKUP_TABLE` gets as input the coordinate array `coords` of the input point set \mathcal{Y} , the nodes $V_{I \times I}(l)$ on level l in `node_data`, and further size information. First, the lower index bounds $i_{l,1}$ and upper index bounds $i_{u,1}$ are extracted from each node and stored in arrays `lower_index_bounds` and `upper_index_bounds`. By construction, the (block) cluster tree traversal based on Z-order curves only creates clusters that do not overlap and that, for a given lower index bound, will always have the same upper bound. Therefore, we can use parallel sorting and unification methods to identify the set of unique clusters. The unique clusters are collected (by their lower and upper index bounds) in `unique_lower_index_bounds` and `unique_upper_index_bounds`. The final step is to compute the coordinate minima and maxima in each subset. This step follows the ideas on batching, cf. Section 4.2. The *batched array* is the array of coordinates. The bounds for the batches are given by the unique lower and upper index bounds and the keys for the batches are the sequence of numbers $\{1, 2, \dots\}$. Results in the batched computation that are associated to points in \mathcal{Y} and not being part of any subset are finally removed by removing all batched compute results associated to the key 0.

Our approach to compute the map between the nodes in `node_data` and the lookup table is summarized in Algorithm 8. Again, we first get the lower and

Algorithm 8 Generator for map to bounding box table

```

procedure CREATE_MAP_FOR_BOUNDING_BOXES(node_data, l  $|V_{I \times I}(l)|$ )
  (lower_bounds, upper_bounds)  $\leftarrow$  GET_INDEX_BOUNDS(node_data)
  SEQUENCE  $\leftarrow$   $|V_{I \times I}(l)|$  (permutation,  $|V_{I \times I}(l)|$ )  $\triangleright$  permutation  $\leftarrow$   $\{0, 1, \dots, |V_{I \times I}(l)|\}$ 
  STABLE_SORT_BY_KEY(permutation, lower_bounds)
  INIT  $\leftarrow$   $|V_{I \times I}(l)|$  (map, 0)
  SET_BOUNDS_FOR_MAP  $\leftarrow$   $|V_{I \times I}(l)|$  (map, lower_bounds)
  INCLUSIVE_SCAN(map, map, 0,  $|V_{I \times I}(l)|$ )
  PERMUTE_MAP  $\leftarrow$   $|V_{I \times I}(l)|$  (map, permutation)
  return map
end procedure

```

	initial state											
	1	4	4	1	1	6	8	8	6	6	8	lower_bounds
	0	1	2	3	4	5	6	7	8	9	10	permutation
STABLE_SORT_BY_KEY	1	1	1	4	4	6	6	6	8	8	8	lower_bounds
	0	3	4	1	2	5	8	9	6	7	10	permutation
	0	0	0	0	0	0	0	0	0	0	0	map
SET_BOUNDS_FOR_MAP	0	0	0	1	0	1	0	0	1	0	0	map
	0	0	0	1	1	2	2	2	3	3	3	map
PERMUTE_MAP	0	3	4	1	2	5	8	9	6	7	10	permutation
	0	1	1	0	0	2	3	3	2	2	3	map

Fig. 8 Creating the map between a work item and an entry in the lookup table requires sorting a compute kernel and a scan operation.

upper index bounds. Then, without loss of generality, we sort the lower bounds of the subsets and keep the applied permutation in `permutation`. Next, we create a global array `map` of length $|V_{I \times I}(l)|$ and initialize it to “0”. The parallel kernel `set_bound_for_map` of $|V_{I \times I}(l)|$ threads then sets a “1” in `map` wherever there are two different subsequent entries in the sorted `lower_bounds`. By an inclusive scan on `map`, we create growing indices in `map` marking identical entries in `lower_bounds`. The result is exemplified in Fig. 8. We finally permute back `map` by kernel `permutation` with $|V_{I \times I}(l)|$ threads leading to the required `map`.

5.4 Numerical linear algebra

During the block cluster tree traversal, an array `work_queue` of `work_item` structs is constructed (via the parallel output queue), cf. Fig. 9. It contains the matrix sub-block information of blocks which are either approximated by ACA or directly constructed as dense matrices, i.e. admissible or non-admissible. Note that we did not evaluate a single matrix entry up to this point. So we only work on meta data. We initially decompose the `work_queue` into two according sub-arrays `aca_work_queue` and `dense_work_queue`, cf. Fig. 9. For

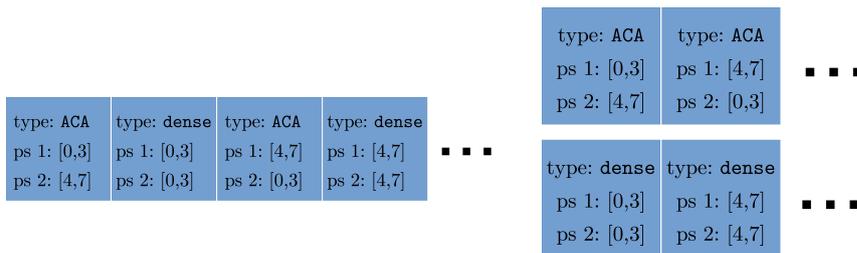


Fig. 9 *Left*: The work queue generated during the block cluster tree traversal only contains meta information. No matrix element has been evaluated, yet. *Right*: The work queue is split up into admissible and non-admissible work elements (ACA vs. dense) before it is used as ordering for the batched linear algebra operations.

the sub-matrices represented by the entries of these arrays, we either apply adaptive cross approximation or dense matrix-vector operations.

In classical (sequential) \mathcal{H} matrix implementations, both, the factors U and V of the adaptive cross approximation *and* the dense matrix blocks are precomputed during an initialization phase and then stored in memory. This is due to the fact, that often, e.g. in boundary element methods, the evaluation of a single matrix entry is already considered very expensive, a storage operation in memory is relatively cheap and large amounts of (CPU) memory are available. Using many-core processors, this balance is somewhat different. Here, evaluating matrix elements is often much faster. However storing data in global memory, i.e. not keeping it in the local memory of the kernel, is rather expensive. Moreover, the memory of many-core processors is often very limited. Therefore, we adapt the classical strategy to the abilities of many-core processors in the following way: We normally always re-compute all low-rank approximations and re-assemble dense matrices during each application of the fast matrix-vector product. Thereby we do not run into the very strong memory limitations of many-core processors. However, we also add the option to pre-compute the construction of the factors U and V in the adaptive cross approximation once, while using these factors during many matrix-vector products. Note however that this is very memory-consuming. A pre-computation of the dense sub-blocks is never done.

In the following, the details of the batched computation and application of the adaptive cross approximation and the dense matrix-vector products are presented.

5.4.1 Batched adaptive cross approximation

We apply batching to compute and apply adaptive cross approximations for all ACA elements in the `aca_work_queue`. The storage pattern is to consecutively store elements $\mathbf{u}_l^{(0)}, \mathbf{u}_l^{(1)}, \dots, \mathbf{u}_l^{(m-1)}$ in memory, where a single ACA sub-matrix $U^{(i)}$ is given as $U^{(i)} = (\mathbf{u}_1^{(i)} \mathbf{u}_2^{(i)} \dots \mathbf{u}_k^{(i)})$. The top index is the batch number and l is the index of the rank-one information. The blocks of batched rank-one

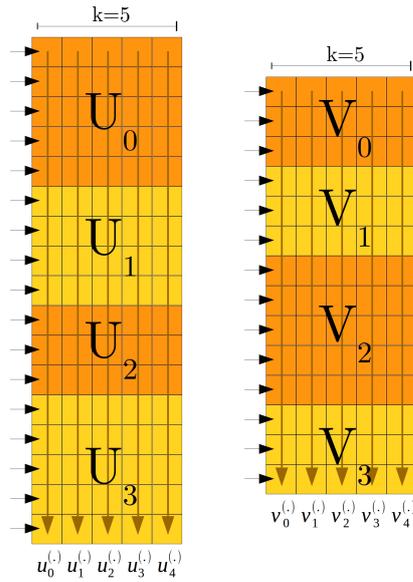


Fig. 10 In the batched version of the adaptive cross approximation, the columns of the low-rank factors U_l and V_l are stored consecutively in memory. The brown arrows indicate the order in memory while the black arrows indicate threads in the parallelization.

information is then stored consecutively for $l = 0, \dots, k - 1$, where k is the maximum number of ranks that is initially given as user argument. Figure 10 shows this storage principle.

In the batched ACA computation, we first set up several meta data arrays describing mainly mappings between the batched ACA data, indices of the input point set and the work items in the `aca_work_queue`. These mappings are used to have constant-time access in kernels being parallelized over the points, over the `aca_work_queue` entries or over the batched ACA data. We can compute these maps similar to the approaches presented e.g. in Algorithm 5. Then, we execute the classical ACA algorithm in a batched version. That is, simple transformations can be applied directly to the full batched array while batch-wise reductions are handled as described in Section 4.2.

Note that the ACA algorithm has an data-dependent iterative behavior, e.g. in case of pivoting. That is, the algorithm might need different numbers of iterations for different batches. We cope with this by introducing a voting mechanism which stops iterations on batched data, whenever all batches have finished their work. A drawback of this approach is that the runtime for the batched version is bound from above by the *slowest* batch. However, from our practical tests, this has never been a performance issue.

Depending on the choice of pre-computing or directly applying the low-rank factors U and V , we either keep these factors in global memory for later

use or we directly apply them using BLAS library calls for dense matrix-vector products.

If we choose to recompute the ACA during each matrix-vector product, we further have the opportunity to split up the whole batched ACA computation to several smaller batched ACA operations. This allows to approach much larger matrices, which would otherwise not fit into GPU memory. To make this possible, we have to choose the number m of matrix batches per batched matrix. We have designed a heuristics, which fills up a batched matrix with matrices of size $n_{b_i} \times k$ as long as $\sum_i n_{b_i}$ is smaller than a threshold bs_{ACA} , i.e. the batching size for ACA. As we will see in Section 6.6.1, the choice of this batching size parameter is important for the performance of the code.

5.4.2 Batched dense sub-matrix application

The application of the dense sub-matrix matrix-vector products is also done in a batched, parallel way. Analogously to the batched ACA computation, we first assemble, entirely in parallel, a larger number of dense sub-blocks using an appropriate compute kernel. The storage principle is similar to the one presented in the previous paragraph, i.e. we stack the dense matrices of size $n_{b_i} \times n'_{b_i}$ on top of each other. To get a simpler representation in memory, we pad all batched sub-blocks by zero columns such that they have all the same column count $\max_i n'_{b_i}$. Afterwards, we use a batched version of *BLAS* for the dense matrix-vector products.

As in the case of batched ACA computation, we have designed a heuristics to create batches of fixed maximum size. In case of the batched dense matrix-vector products, we choose to keep the total batch storage size smaller than a threshold bs_{dense} ,

$$\max_i n'_{b_i} \cdot \sum_i n_{b_i} \leq bs_{dense}.$$

6 Results

In this section, we evaluate the performance of the above described many-core parallel algorithms in the concrete GPU implementation `hmglib` [45] by the author. The library is available via GitHub and is licensed under LGPL License Version 3.0. This implementation only covers the \mathcal{H} matrix construction or setup and \mathcal{H} matrix vector product for a matrix $A_{\phi, \mathcal{Y}_1 \times \mathcal{Y}_2}$ for a given kernel function ϕ and sets \mathcal{Y}_1 and \mathcal{Y}_2 . It is not intended to be feature-complete, i.e. providing the full \mathcal{H} matrix algebra. Instead, it is a test bead for the above discussed many-core parallel algorithms. Nevertheless, it is possible to solve linear systems of type (1) by using the iterative dense linear solvers library `MPLA` [44] by the author (open source, available on GitHub), which has an interface to `hmglib`. However, the objective of this benchmark chapter is to stick to the discussion of the construction and the \mathcal{H} matrix-vector products,

avoiding to confuse the reader by solver details and with two different library implementations.

In the following, we start our discussion by giving brief details on the library `hmglib` with the targeted hardware and applied external many-core parallel libraries. Afterwards, we introduce a model problem and show empirically that the implemented approximate matrix-vector product converges exponentially in the number ranks used in the adaptive cross approximation for the given model problem. Since the main goal is, to get a code of optimal complexity, we check the runtime complexity of `hmglib` by numerical experiments. Thereafter, we give details about the performance improvements made by batching. In fact, these performance improvements are the most relevant ones for our final results. We finish this section, by comparing the runtimes of `hmglib` against a reference CPU implementation. Note here, that we will compare a *sequentially* used multi-purpose state-of-the-art, open source library for hierarchical matrices (`H2Lib` [10]) with a very specific, parallel many-core implementation. This comparison is non-optimal. Therefore, the results of this study are only treated as a rough hint towards the actual performance improvement by using `hmglib`.

6.1 GPU implementation `hmglib`

The library `hmglib` [45] is implemented for graphics processing units of *NVIDIA Corporation*. Our notion of a compute kernel from Section 3.1 can be easily mapped to the compute kernels in the C language extension *CUDA* for programming *NVIDIA* GPUs. Note however, that an implementation in *OpenCL* (for *NVIDIA* and *AMD* GPUs) or *OpenMP* with extensions for *Intel Xeon Phi* devices should be equally simple. Within our hand-implemented *CUDA* compute kernels, we always use a so-called *block size* of 512, i.e. 512 threads are bundled in a block with common shared memory (which we actually do not explicitly use). `hmglib` uses the *CUDA Toolkit 8.0*. It is compiled with optimization parameter `-O3`. As CPU code compiler, `gcc 4.8.5` is used.

Within our many-core parallel algorithms in Section 5, we launch, beside of compute kernels, library calls for general many-core parallel STL-type algorithms. In `hmglib`, the library `Thrust`, which is delivered as part of the *CUDA Toolkit*, provides these STL-type algorithms. `Thrust` contains all the necessary parallel algorithms and delivers decent performance for GPUs. Moreover, we use BLAS-type operations of the library `CUBLAS`, which is also delivered as part of the *CUDA Toolkit*. In case of the batched application of dense matrix-vector products, we apply the state-of-the-art GPU *Lapack* library `Magma 2.2.0`. There, we specifically use the batched multiplication `magmablas_dgemv_vbatched`.

`hmglib` allows to select, whether batching is applied in the matrix-vector product, or not. Moreover, it is possible to switch on the pre-computation of the low-rank factors in the adaptive cross approximation. This requires a lot of GPU memory. However, \mathcal{H} matrix-vector products can be applied faster if

the low-rank factors do not have to be recomputed for each multiplication. Remember that in CPU-based \mathcal{H} matrix implementations, the dense sub-blocks of the approximated matrix are often pre-computed, too. This is not done here, due to limited GPU memory and very fast matrix assembly on GPU. All calculations are done in double precision.

6.2 Model problem

All benchmarks consider matrix-vector products of the form

$$A_{\phi, \mathcal{Y} \times \mathcal{Y}} \mathbf{x}$$

with

$$A_{\phi, \mathcal{Y} \times \mathcal{Y}} = \begin{pmatrix} \phi(\mathbf{y}_1, \mathbf{y}_1) & \cdots & \phi(\mathbf{y}_1, \mathbf{y}_N) \\ \vdots & \ddots & \vdots \\ \phi(\mathbf{y}_N, \mathbf{y}_1) & \cdots & \phi(\mathbf{y}_N, \mathbf{y}_N) \end{pmatrix}, \quad \mathbf{x} \in \mathbb{R}^N.$$

where $\mathcal{Y} := \{\mathbf{y}_1, \dots, \mathbf{y}_N\} \subset \Omega$ is a set of N points in a space $\Omega \subset \mathbb{R}^d$ and $\phi : \Omega \times \Omega \rightarrow \mathbb{R}$ is a bivariate kernel function operating on that domain. We specifically choose $\Omega = [0, 1]^d$ with $d = 2, 3$. Moreover, the point set is a Halton sequence, i.e. a quasi Monte-Carlo sequence, of length N in d dimensions. This choice corresponds to the typical setup in kernel-based approximation on the unit square / cube. We test the implementation with different (unscaled) kernel functions, namely the Gaussian kernel

$$\phi_G(\mathbf{y}, \mathbf{y}') = e^{-\|\mathbf{y} - \mathbf{y}'\|^2}$$

and a Matérn kernel [16, Section 4.4]

$$\phi_M(\mathbf{y}, \mathbf{y}') = \frac{K_{\beta - \frac{d}{2}}(\|\mathbf{y} - \mathbf{y}'\|) \|\mathbf{y} - \mathbf{y}'\|^{\beta - \frac{d}{2}}}{2^{\beta-1} \Gamma(\beta)},$$

where K_ν is the modified Bessel function of second kind of order ν and Γ is the gamma function. We choose $\beta - \frac{d}{2} = 1$. The resulting matrix $A_{\phi_M, \mathcal{Y} \times \mathcal{Y}}$ shows up in first-order convergent function interpolation schemes in kernel-based interpolation [16, Theorem 14.5, Example 15.4] for appropriately smooth functions. The norm $\|\cdot\|$ is the usual Euclidean norm of appropriated dimensionality.

This model represents the application fields of mesh-free kernel-based approximation, (non-regularized) kernel ridge regression and, in some cases, Gaussian process regression.

6.3 Hardware setup and time measurements

While a major part of the development work has been carried out on the cluster Titan at Oak Ridge National Lab, the benchmarking was done on the *PSG Cluster* of *NVIDIA Corporation*. On the latter one, IBM S822LC compute nodes with *IBM POWER8* architecture were used. They are each equipped with two 10-core *IBM POWER8* processors at 2.86 GHz, 512 GB RAM and four NVIDIA Tesla P100 SXM2. Only one out of these four GPUs was used. Our CPU performance comparison is done on the same platform. Additionally, we give timings for a standard *iMac* with *Intel Core i5* processor at 3.2 GHz and 16 GB RAM.

Whenever we use GPU-based calculations, we use *CUDA Events* to get very accurate time measurements. The time required by potentially necessary data transfers between GPU and CPU is always included. However, we assume the initial data, i.e. the point set \mathcal{Y} to reside in GPU memory. In case of CPU-based \mathcal{H} matrix benchmarks, we use the *gettimeofday* command to do the measurements. All measurements (GPU and CPU) are averaged results over five trials of a \mathcal{H} matrix construction or a \mathcal{H} matrix-vector product with different random vectors \mathbf{x} .

6.4 Convergence of the matrix-vector product approximation

We start our experiments by checking the convergence of our \mathcal{H} matrix implementation for growing ACA rank k for all discussed kernel functions in two and three dimensions and problem size $N = 32768$. Furthermore, we choose $C_{leaf} = 256$ and $\eta = 1.5$. All other parameters are not relevant for this convergence study. As for the performance measurements, we perform five runs and average over each result. The error in each run is the relative error

$$e_{rel} = \frac{\|\mathcal{H}(A_{\phi, \mathcal{Y} \times \mathcal{Y}})\mathbf{x}_{rand} - A_{\phi, \mathcal{Y} \times \mathcal{Y}}\mathbf{x}_{rand}\|_2}{\|A_{\phi, \mathcal{Y} \times \mathcal{Y}}\mathbf{x}_{rand}\|_2}$$

for a random input vector \mathbf{x}_{rand} . $\mathcal{H}(A_{\phi, \mathcal{Y} \times \mathcal{Y}})$ is the \mathcal{H} matrix approximation of the full system matrix $A_{\phi, \mathcal{Y} \times \mathcal{Y}}$. Note that we are strongly limited in the problem size N since we do all computations on GPU and therefore have to do the full matrix vector product $A_{\phi, \mathcal{Y} \times \mathcal{Y}}\mathbf{x}_{rand}$ in GPU memory.

Fig. 11 shows on the left-hand side the convergence results for $d = 2$ and the two different kernels from the model problem. Our implementation delivers exponential convergence in the number k of ranks used in the adaptive cross approximation. The same test is repeated for dimension $d = 3$ with similar results. Since the results for Gaussian and Matérn kernel are almost identical, we will, in the following paragraphs, restrict ourselves to performance studies for the Gaussian kernel.

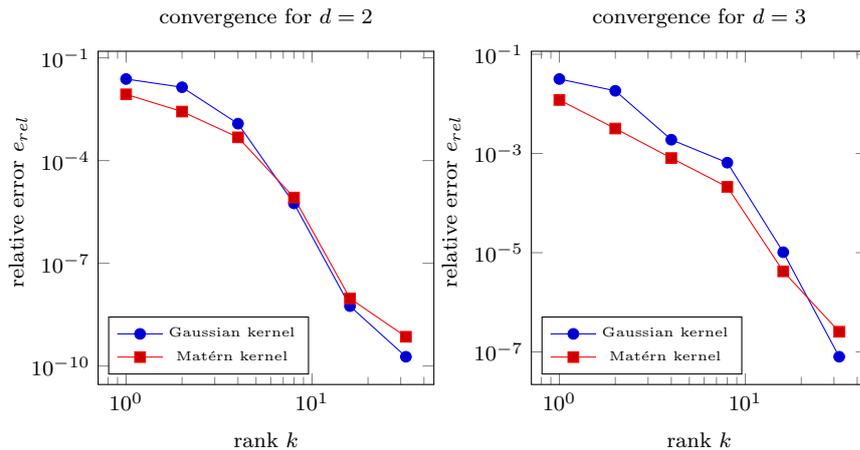


Fig. 11 For fixed problem size $N = 32768$ and growing number of ranks in the adaptive cross approximation, the \mathcal{H} matrix-vector product converges to the full matrix-vector product with exponential convergence for dimensions $d = 2$ (left) and $d = 3$ (right).

6.5 Runtime complexity and performance of the GPU-parallel code

The crucial objective of an implementation of the hierarchical matrix method is to achieve the optimal runtime complexity of $O(N \log N)$ for the matrix-vector product at fixed rank k . However, very often, high (pre-asymptotic) runtime performance on many-core hardware is only achieved by sticking to algorithmic simplifications of worse complexity but higher performance. The following empirical study shall show that the \mathcal{H} matrix implementation in `hmglib`, which is based on our many-core parallel \mathcal{H} matrix algorithms from Section 5, actually achieves the required $O(N \log N)$ runtime complexity. To study this, we choose $\eta = 1.5$, $C_{leaf} = 2048$, $k = 16$, $bs_{dense} = 2^{27}$ and $bs_{ACA} = 2^{25}$, use batching and carry out performance measurements for growing problem size N .

We first discuss the runtime complexity of the setup of the spatial data structure. While computing the Morton codes for all points \mathbf{y}_i is of complexity $O(N)$, sorting the points following the Z order curve is a $O(N \log N)$ operation. This is reflected by our empirical study shown on the left-hand side of Fig. 12. For $d = 2$ and $d = 3$ we observe a runtime complexity of $O(N \log N)$ after some pre-asymptotic range. The same behavior is observed for the construction and the traversal of the block cluster tree. Runtime results for this case are given on the right-hand side of Fig. 12. Note again that it is non-trivial to get the optimal complexity for such a complex many-core parallel code. Figure 12 further outlines that the spatial data structure setup and the tree traversal is actually very fast. Even for $N = 2^{26}$, i.e. an approximation of a dense matrix of roughly 67×67 million entries, we only need roughly 0.4 seconds for the spatial data structure and about 3 seconds for the tree traversal (for $d = 3$).

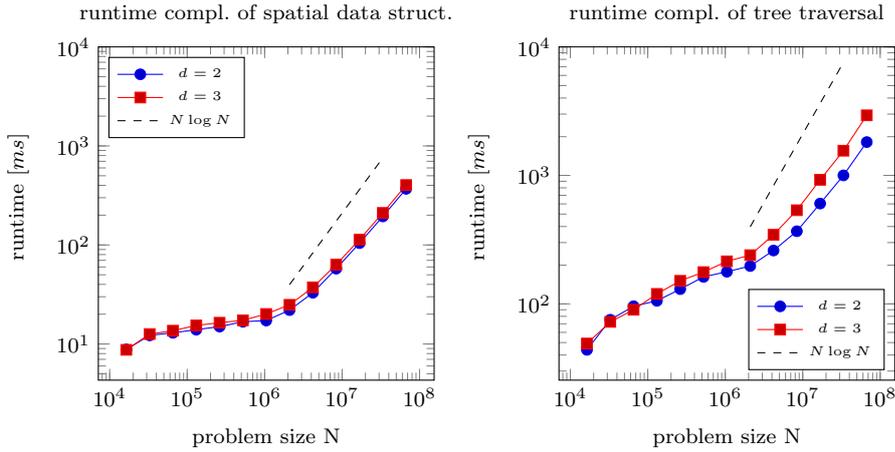


Fig. 12 Even for the largest problem size of $2^{26} \approx 67$ million unknowns, the time for the spatial data structure setup is below 0.5 seconds (*left*). The tree construction and traversal requires less than 3 seconds for 2^{26} unknowns, while matching the required runtime complexity of $N \log N$ (*right*).

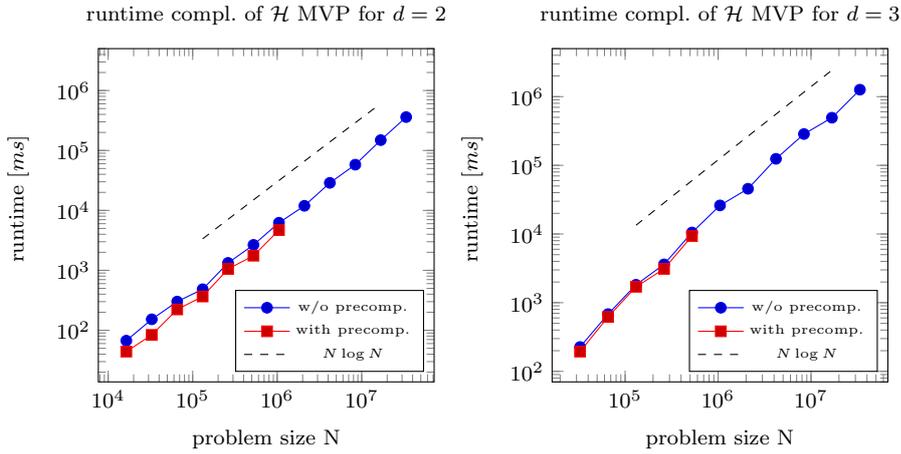


Fig. 13 The \mathcal{H} matrix-vector product shows the optimal algorithmic complexity of $O(N \log N)$. The operation is slightly more expensive if used on points in tree dimension (*right*) in contrast to points in two dimensions (*left*). Using pre-computation for the ACA factors leads to a performance improvement.

The second part of this runtime complexity study covers the application of the fast matrix-vector product. Figure 13 shows the measurements of the runtime for growing problem size N and different dimensionality d . Within each performance plot, we further distinguish between measurements that were done using a matrix-vector product with precomputed ACA factors and with

on-the-fly computation of the ACA factors. Pre-computing the ACA factors results in a performance improvement, which will be discussed in more detail in Section 6.7. In the plot, the impact is not clearly visible due to the logarithmic scaling of the axis. We cannot show runtime results with pre-computing for problem sizes beyond $N = 2^{19}$ or $N = 2^{20}$ due to the limited GPU memory.

In all cases, we observe a runtime complexity of $O(N \log N)$. Moreover, even for a problem size of $N = 2^{25}$, i.e. an approximated matrix-vector product for a dense matrix of 33×33 million entries, we see a runtime of only 6 minutes for a matrix-vector product on points in two dimensions.

6.6 Performance analysis of batching

Beforehand, we discussed the performance results of our implementation using batching in all linear algebra operations, as discussed in Section 5.4. However, it is important to know that batching is one of the crucial ingredients of this code allowing for high performance of the overall method. To show the actual impact of batching, we will analyse the performance with and without batching in the linear algebra operations. However, before we come to this point, we want to address the topic of parameter choice of the batching sizes bs_{dense} and bs_{ACA} .

6.6.1 Batching size influence

In Section 5.4, we introduced the parameters bs_{dense} and bs_{ACA} as batching sizes for the batching of the dense matrix-vector products and the batching of the adaptive cross approximation. These parameters balance the memory consumption against the performance improvement. To understand this relationship further, we benchmark the runtime of the batched dense matrix-vector products and the batched ACA in the \mathcal{H} matrix-vector product for different batching sizes. It is done for $N = 2^{20}$, $k = 16$, $\eta = 1.5$ and $d = 2$. We consider results for $C_{leaf} = 1024$ and $C_{leaf} = 2048$.

Figure 14 collects the results for the parameter studies in the batching size for the batched dense matrix-vector products on the left-hand side and for the batched ACA computation on the right-hand side. The choice of the leaf size C_{leaf} has a considerable influence on the performance balance between dense matrix-vector products and ACA. That is, larger leaf sizes lead to larger runtimes in the dense matrix-vector products. However the ACA runtime is reduced. The opposite holds for smaller leaf sizes. Moreover, the smaller leaf size of $C_{leaf} = 1024$ in batched ACA leads to a higher memory consumption for the batching, which limits us to a maximum of $bs_{ACA} = 2^{25}$.

The general tendency in the results in Fig. 14 is that increasing the batching size increases the performance to an optimum. Beyond this optimum, the performance of the batching gets slightly worse. This performance improvement up to an optimum is due to the improvement of the occupancy of the

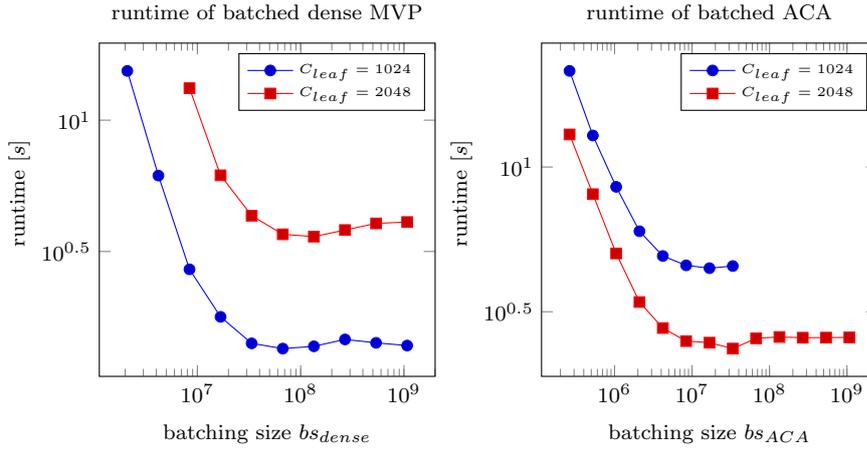


Fig. 14 The performance of batching strongly depends on the size of the batched array or matrices that are used. This is clearly visible for the batching of dense matrix-vector products (*left*) and adaptive cross approximation (*right*). The optimal batching size is only slightly influenced by the choice of the parameter C_{leaf} .

GPU. That is, the GPU gets more work to do. Thereby, it can hide latencies etc. *behind* parallel work. The slight performance degradation beyond the optimum for larger batching sizes is maybe due to a slight over-subscription of the GPU: The maximum throughput limit is hit, however, due to more batches per batched operation, the data structure overhead becomes visible. Note, however, that this latter reasoning is speculative.

Overall, choosing an appropriate batching size is rather simple. The rule of thumb is to take it as large as possible (in terms of memory consumption) and to accept the slight performance reduction for a too large batch size.

6.6.2 Performance improvement by batching

We next discuss the performance improvement for batched dense matrix-vector products and for batched ACA. We use parameters $N = 2^{20}$, $k = 16$, $\eta = 1.5$, $d = 2$, $C_{leaf} = 2048$, $bs_{dense} = 2^{27}$ and $bs_{ACA} = 2^{25}$. Figure 15 summarizes the results of this study with results for the batching of dense matrix-vector products on the left-hand side and results for the batched adaptive cross approximation on the right-hand side. For a problem size of $N = 2^{20}$, the batched application of the dense matrix-vector products is by more than a factor of 3 faster. We do not gain more, since, for $C_{leaf} = 2048$, we have a lot of large dense matrix sub-blocks which very soon fully occupy the GPU.

In contrast, the performance improvement for the adaptive cross approximation is about a factor of 32 for $N = 2^{20}$. This strong impact is due to the small amount of work that is done for each individual ACA computation and is a significant contribution of this work.

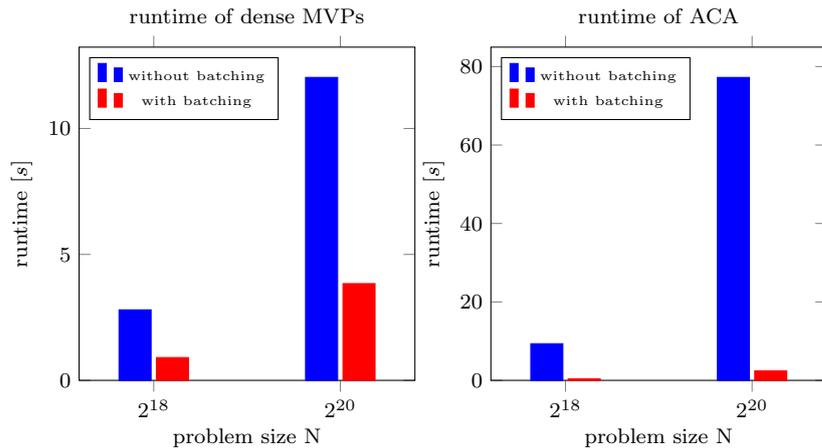


Fig. 15 We observe a significant performance improvement by roughly a factor of 32, when using batching in the ACA computation (*right*). Batching dense matrix-vector products still improves performance by roughly a factor of three (*left*).

To summarize, an efficient \mathcal{H} matrix-vector product would not be possible without ACA batching. However, it also pays off to do batching for the dense matrix-vector products.

6.7 Performance comparison against H2Lib

In the following, we aim at relating the performance of `hmglib` to the CPU \mathcal{H} and \mathcal{H}^2 matrix library H2Lib [10] in the, at time of writing this paper, latest available version. We have chosen H2Lib, since it is under active development and an Open Source library. The H2Lib library implements an algebra for \mathcal{H} matrices and \mathcal{H}^2 matrices. That is, the library allows to construct, add, multiply, factorize, etc. \mathcal{H} and \mathcal{H}^2 matrices. Moreover, it contains modules for the solution of problems discretized by the boundary element method. Recently, support for a GPU-accelerated \mathcal{H}^2 matrix setup for boundary element method problems was added [11], as discussed in Section 1. H2Lib also contains some support for shared-memory parallelism. However, it seemed to have no impact on the performance of the \mathcal{H} matrix construction and the \mathcal{H} matrix-vector product. Therefore, we used the sequential version, only.

As argued before, the comparison of our GPU implementation, which only implements the \mathcal{H} matrix-vector product, with this feature-complete sequential CPU implementation, which has been specifically optimized for \mathcal{H}^2 matrices and boundary element method problems, is non-optimal by construction. However, we add this comparison to somehow relate our performance results to currently available software in the field.

In our performance benchmarks, we try our best to fit the H2Lib implementation to our GPU implementation, even if this means that we have to

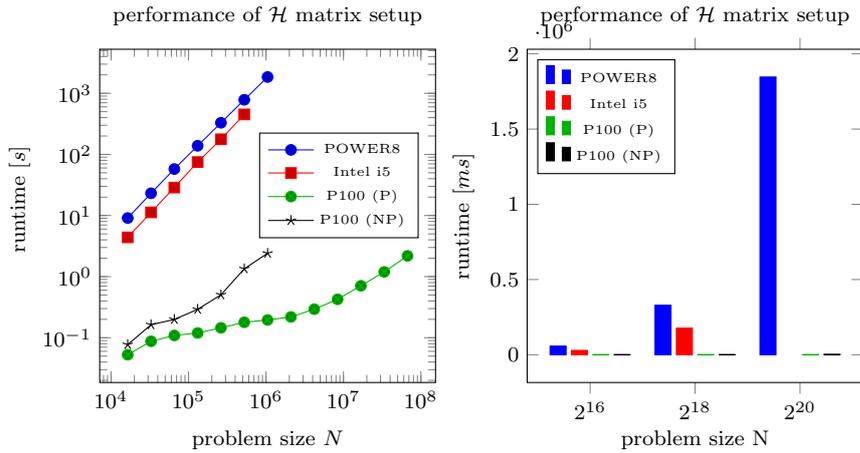


Fig. 16 We compare the runtime of the \mathcal{H} matrix setup in the *H2Lib* (including the computation of the ACA and all dense sub-blocks) with the setup in the *hmglib* library with (*P*) and without (*NP*) pre-computing the ACA. The GPU-based implementation outperforms the *sequential* CPU-based implementation by more than two orders of magnitude.

extend the *H2Lib* for this. To give an example, we added the ability to do ACA for a fixed rank k , which was not available in the library, before. On the IBM POWER8 platform, *H2Lib* is compiled with `gcc 4.8.5` and the usual optimizations and linked against the, at time of writing this article, latest available version of `OpenBLAS`. On the Intel architecture, it is compiled with the same compiler, however linked against the default *LAPACK* implementation of *macOS Sierra 10.12*.

We start the comparison with a benchmark of the \mathcal{H} matrix construction or setup phase. In case of the *H2Lib* this construction phase contains the spatial data structure setup, the block cluster tree traversal, the pre-computation of all low-rank factors and the assembly of all dense sub-blocks of the \mathcal{H} matrix. We choose $\eta = 1.5$ and $C_{leaf} = 128$. On the other hand, we choose $\eta = 1.5$, $C_{leaf} = 2048$, $bs_{dense} = 2^{27}$ and $bs_{ACA} = 2^{25}$ in the GPU implementation and analyse the construction phase including pre-computation (*P*) of the ACA factors or without (*NP*) such a pre-computation. Note that the leaf size C_{leaf} has a significant impact on the performance on the method. Therefore, we adapt it for the different architectures for best possible performance. All results in this paragraph are computed for a fixed rank of $k = 16$ and dimension $d = 2$. Since we observed very strong fluctuations of the runtime on the Intel workstation, we always take the smallest runtime out of five trials on that architecture to be as fair as possible.

Figure 16 gives the result for the first comparison. On the left-hand side, runtimes of the setup phase are given for growing problem size. The diagram on the right-hand side directly compares the results on the different architectures for fixed problem sizes. Due to limited memory, the benchmark is

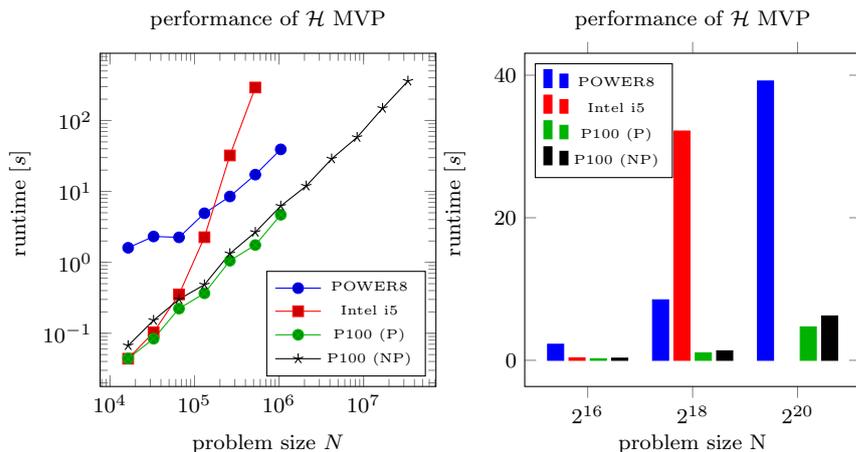


Fig. 17 The GPU-based \mathcal{H} matrix-vector product outperforms the *single*-threaded CPU-based \mathcal{H} matrix vector product by one order of magnitude. By precomputing (P) the ACA factors we get an increase in performance by about 60 %.

stopped for $N = 2^{19}$ on the Intel machine and for $N = 2^{20}$ on the GPU with pre-computing. The benchmark on the POWER8 CPU system is stopped for $N = 2^{20}$ due to large runtime. In case of the largest common problem size, i.e. $N = 2^{19} \approx 0.5$ million points, the CPU implementation requires 782 seconds and 451 seconds on the POWER8 system and the Intel system, respectively. On the other hand, the GPU implementation only needs 1.3 seconds with precomputing and 0.8 seconds without pre-computing. That is, it is more than two orders of magnitude faster. However, note again that the setup phase on CPU also pre-computes the dense matrix sub-blocks. Moreover, we compare a sequential implementation with a strongly parallelized GPU implementation. Even more, the single-threaded performance of the POWER8 system seems to be limited, which is why we also included the Intel workstation in the benchmark. A more fair comparison would e.g. compare a parallel CPU code on 16 CPU cores with the GPU code. However, even in this case (assuming perfect scalability on CPU), the GPU would outperform the CPU-based version by a factor of twenty.

Our second comparative study targets the \mathcal{H} matrix-vector product. It is done with the same parameters as before. The results are given in Fig. 17. While both GPU results and the POWER8-based result on CPU show the usual complexity behavior, we observe a significant increase in runtime for growing problem size on the Intel architecture. Right-now, the reason for this behavior is not clear. This is why we exclude this result from our discussion, here. Comparing the three remaining results, we observe a strong performance improvement of the GPU-based runs against the CPU-based results. Comparing again the results for $N = 2^{19}$, we see a runtime of about 17 seconds on the CPU and a runtime of 2.7 seconds without ACA pre-computing and an

improvement by about 60 % to 1.7 seconds with pre-computing on GPU. This is still a remarkable performance improvement by a factor of 10 on GPU. However, comparing a (fictive) 16-core parallel CPU implementation with the GPU results might result in a comparable performance between both architectures (depending on the scalability assumptions for the CPU). At this point, we still have to keep in mind that the CPU-based code assembles and stores all dense matrix sub-blocks of the approximated matrix, beforehand, while `hmglib` recomputes these on-the-fly due to memory limitations. Moreover, we still see some room for performance improvements of the GPU implementation.

Overall, we conclude that a perfectly fair comparison is hardly possible. CPU-based implementations rely much more on pre-computation and therefore might have a slight advantage for the \mathcal{H} matrix-vector product in a 16-core CPU to GPU comparison, while being much slower in the setup phase. The new GPU-based implementation tries to balance the strong memory restrictions of GPUs with a general performance improvement. Based on the raw numbers of the single-threaded CPU to GPU comparison, the GPU code outperforms the CPU code by two orders of magnitude for the setup and by one order of magnitude for the \mathcal{H} matrix-vector product.

7 Summary

This work considered the reformulation of algorithms in the construction and matrix-vector product of \mathcal{H} matrices for many-core parallelism. As core techniques, to get fast parallel performance of \mathcal{H} matrices on many-core hardware, we identified a parallel spatial data structure based on space filling curves, parallel tree traversal and batching of many small, non-equally sized compute tasks. On top of these basic building blocks, we designed algorithms for many-core parallel \mathcal{H} matrices. These algorithms were transferred to a reference implementation on a GPU, which results in the GPU \mathcal{H} matrix library `hmglib`. Our computation results section showed that the designed algorithms lead to a fast GPU implementation. Compared to the *sequential* version of the `H2Lib` library, we achieve more than two orders of magnitude performance improvement on one Tesla P100 SXM2 GPU for the \mathcal{H} matrix construction and roughly one order of magnitude performance in the \mathcal{H} matrix-vector product. Note however that comparing the libraries and the underlying hardware is somewhat difficult. Nevertheless, we tried our best to keep this comparison fair.

In the future, our new algorithms shall be extended to the use in a distributed-memory, thus e.g. multi-GPU, context. This, however, involves to build an appropriate load balancing for the work distribution of ACA computations and dense matrix-vector products on an entire cluster of compute nodes equipped with many-core hardware. Moreover, the heterogeneous nature, i.e. the existence of powerful CPUs and many-core devices, of current compute cluster should also be addressed, in order to get an even higher performance out of these systems.

Acknowledgements This work is funded by the Swiss National Science Foundation (SNF) under project number 407540_167186. Furthermore, code developments tasks in this research were done on resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. The IBM POWER8 system with the NVIDIA Tesla P100 SXM2 used in the benchmarks for this research was donated by the *NVIDIA PSG Cluster*. All funding and support is gratefully acknowledged.

References

1. Abdelfattah, A., Haidar, A., Tomov, S., Dongarra, J.: Novel HPC techniques to batch execution of many variable size BLAS computations on GPUs. In: Proceedings of the International Conference on Supercomputing, ICS '17, pp. 5:1–5:10. ACM, New York, NY, USA (2017)
2. Agullo, E., Bramas, B., Coulaud, O., Darve, E., Messner, M., Takahashi, T.: Task-based FMM for multicore architectures. *SIAM Journal on Scientific Computing* **36**(1), C66–C93 (2014)
3. Bebendorf, M.: *AHMED* Another software library on hierarchical matrices for elliptic differential equations URL <https://github.com/xantares/ahmed>
4. Bebendorf, M.: Hierarchical Matrices - A Means to Efficiently Solve Elliptic Boundary Value Problems, *Lecture Notes in Computational Science and Engineering*, vol. 63. Springer (2008)
5. Bebendorf, M., Kunis, S.: Recompression techniques for adaptive cross approximation. *J. Integral Equations Applications* **21**(3), 331–357 (2009)
6. Bebendorf, M., Rjasanow, S.: Adaptive low-rank approximation of collocation matrices. *Computing* **70**(1), 1–24 (2003)
7. Bell, N., Hoberock, J.: Thrust: A productivity-oriented library for CUDA. *GPU computing gems Jade edition* **2**, 359–371 (2011)
8. Bern, M., Eppstein, D., Teng, S.H.: Parallel construction of quadrees and quality triangulations. *International Journal of Computational Geometry & Applications* **09**(06), 517–532 (1999)
9. Börm, S.: \mathcal{H}^2 -matrices - Multilevel methods for the approximation of integral operators. *Computing and Visualization in Science* **7**(3), 173–181 (2004)
10. Börm, S.: *H2Lib*, a library for hierarchical matrices (2017). URL <http://www.h2lib.org>
11. Börm, S., Christophersen, S.: Approximation of BEM matrices using GPGPUs. *ArXiv e-prints* (2015)
12. Börm, S., Grasedyck, L., Hackbusch, W.: Introduction to hierarchical matrices with applications. *Engineering analysis with boundary elements* **27**(5), 405–422 (2003)
13. Boukaram, W., Ltaief, H., Litvinenko, A., Abdelfattah, A., Keyes, D.E.: Accelerating matrix-vector multiplication on hierarchical matrices using graphical processing units. URL <http://hdl.handle.net/10754/347275>
14. Boukaram, W.H., Turkiyyah, G., Ltaief, H., Keyes, D.E.: Batched QR and SVD Algorithms on GPUs with Applications in Hierarchical Matrix Compression. *ArXiv e-prints* (2017)
15. Charara, A., Keyes, D.E., Ltaief, H.: Batched triangular dense linear algebra kernels for very small matrix sizes on GPUs. *ACM Transactions on Mathematical Software*
16. Fasshauer, G.F.: *Meshfree Approximation Methods with MATLAB*. World Scientific Publishing Co., Inc., River Edge, NJ, USA (2007)
17. Garanzha, K., Pantaleoni, J., McAllister, D.: Simpler and faster HLBVH with work queues. In: Proceedings of the ACM SIGGRAPH Symposium on High Performance Graphics, HPG '11, pp. 59–64. ACM, New York, NY, USA (2011)
18. Ghysels, P., Li, X.S., Rouet, F., Williams, S., Napov, A.: An efficient multicore implementation of a novel HSS-structured multifrontal solver using randomized sampling. *SIAM J. Scientific Computing* **38**(5) (2016)
19. Grasedyck, L., Kriemann, R., Le Borne, S.: Parallel black box-LU preconditioning for elliptic boundary value problems. *Computing and Visualization in Science* **11**(4), 273–291 (2008)

20. Greengard, L., Rokhlin, V.: A new version of the fast multipole method for the Laplace equation in three dimensions. *Acta numerica* **6**, 229–269 (1997)
21. Hackbusch, W.: Hierarchical matrices : Algorithms and Analysis, *Springer series in computational mathematics*, vol. 49. Springer, Berlin (2015)
22. Hackbusch, W.: Survey on the technique of hierarchical matrices. *Vietnam Journal of Mathematics* **44**(1), 71–101 (2016)
23. Hackbusch, W., Börm, S.: \mathcal{H}^2 -matrix approximation of integral operators by interpolation. *Applied numerical mathematics* **43**(1-2), 129–143 (2002)
24. Hackbusch, W., Khoromskij, B., Sauter, S.A.: On \mathcal{H}^2 -matrices. In: *Lectures on Applied Mathematics: Proceedings of the Symposium Organized by the Sonderforschungsbereich 438 on the Occasion of Karl-Heinz Hoffmanns 60th Birthday*, Munich, June 30–July 1, 1999, p. 9. Springer Science & Business Media (2000)
25. Hackbusch, W., Nowak, Z.P.: On the fast matrix multiplication in the boundary element method by panel clustering. *Numerische Mathematik* **54**(4), 463–491 (1989)
26. Kriemann, R.: Parallel \mathcal{H} -matrix arithmetics on shared memory systems. *Computing* **74**(3), 273–297 (2005)
27. Kriemann, R.: \mathcal{H} -LU factorization on many-core systems. *Comput. Vis. Sci.* **16**(3), 105–117 (2013)
28. Kriemann, R.: \mathcal{H} -Lib^{PRO} (website) (2017). URL <http://www.hlibpro.com>
29. Lauterbach, C., Garland, M., Sengupta, S., Luebke, D., Manocha, D.: Fast BVH construction on GPUs. *Computer Graphics Forum* **28**(2), 375–384 (2009)
30. March, W.B., Xiao, B., Yu, C., Biros, G.: ASKIT: An efficient, parallel library for high-dimensional kernel summations. *SIAM Journal on Scientific Computing* **38**, S720–S749 (2016)
31. Merrill, D., Garland, M., Grimshaw, A.: Scalable GPU graph traversal. In: *Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP '12*, pp. 117–128. ACM, New York, NY, USA (2012)
32. Morton, G.: A computer oriented geodetic data base and a new technique in file sequencing. Tech. Rep. Ottawa, Ontario, Canada (1966)
33. Poulson, J.: DMHM - Distributed-Memory Hierarchical Matrices. URL <https://bitbucket.org/poulson/dmhm>
34. Rasmussen, C., Williams, C.: *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press (2005)
35. Rouet, F.H., Li, X.S., Ghysels, P., Napov, A.: A distributed-memory package for dense hierarchically semi-separable matrix computations using randomization. *ACM Trans. Math. Softw.* **42**(4), 27:1–27:35 (2016). DOI 10.1145/2930660
36. Sheng, Z., Dewilde, P., Chandrasekaran, S.: *Algorithms to Solve Hierarchically Semi-separable Systems*, pp. 255–294. Birkhäuser Basel, Basel (2007)
37. Szuppe, J.: Boost.Compute: A parallel computing library for C++ based on OpenCL. In: *Proceedings of the 4th International Workshop on OpenCL, IWOCCL '16*, pp. 15:1–15:39. ACM, New York, NY, USA (2016)
38. Valiant, L.G.: A bridging model for parallel computation. *Commun. ACM* **33**(8), 103–111 (1990)
39. Vovk, V.: *Kernel Ridge Regression*, pp. 105–116. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
40. Wendland, H.: *Scattered Data Approximation*. Cambridge University Press (2004)
41. Yalamanchili, P., Arshad, U., Mohammed, Z., Garigipati, P., Entschew, P., Kloppeborg, B., Malcolm, J., Melonakos, J.: ArrayFire - A high performance software library for parallel computing with an easy-to-use API (2015). URL <https://github.com/arrayfire/arrayfire>
42. Yokota, R., Barba, L.: FMM-based vortex method for simulation of isotropic turbulence on GPUs, compared with a spectral method. *Computers & Fluids* **80**, 17 – 27 (2013)
43. Yokota, R., Barba, L., Knepley, M.G.: PetRBF A parallel O(N) algorithm for radial basis function interpolation with Gaussians. *Computer Methods in Applied Mechanics and Engineering* **199**(25), 1793 – 1804 (2010)
44. Zaspel, P.: MPLA - Massively Parallel Linear Algebra (2017). URL <https://github.com/zaspel/MPLA>
45. Zaspel, P.: hmglib - Simple H matrix library on GPU (2017). URL <https://github.com/zaspel/hmglib>

LATEST PREPRINTS

- | No. | Author: Title |
|---------|--|
| 2016-22 | P. Habegger
<i>The norm of Gaussian periods</i> |
| 2016-23 | P. Habegger
<i>Diophantine approximations on definable sets</i> |
| 2016-24 | F. Amoroso, D. Masser
<i>Lower bounds for the height in Galois extensions</i> |
| 2016-25 | W. D. Brownawell, D. W. Masser
<i>Zero estimates with moving targets</i> |
| 2016-26 | H. Derksen, D. Masser
<i>Linear equations over multiplicative groups, recurrences, and mixing III</i> |
| 2016-27 | D. Bertrand, D. Masser, A. Pillay, U. Zannier
<i>Relative Manin-Mumford for semi-abelian surfaces</i> |
| 2016-28 | L. Capuano, D. Masser, J. Pila, U. Zannier
<i>Rational points on Grassmannians and unlikely intersections in tori</i> |
| 2016-29 | C. Nobili, F. Otto
<i>Limitations of the background field method applied to Rayleigh-Bénard convection</i> |
| 2016-30 | W. D. Brownawell, D. W. Masser
<i>Unlikely intersections for curves in additive groups over positive characteristic</i> |
| 2016-31 | M. Dambrine, H. Harbrecht, M. D. Peters, B. Puig
<i>On Bernoulli's free boundary problem with a random boundary</i> |
| 2016-32 | H. Harbrecht, J. Tausch
<i>A fast sparse grid based space-time boundary element method for the nonstationary heat equation</i> |
| 2016-33 | S. Iula
<i>A note on the Moser-Trudinger inequality in Sobolev-Slobodeckij spaces in dimension one</i> |
| 2016-34 | C. Bürli, H. Harbrecht, P. Odermatt, S. Sayasone, N. Chitnis
<i>Mathematical analysis of the transmission dynamics of the liver fluke, <i>Opisthorchis viverrini</i></i> |

LATEST PREPRINTS

- | No. | Author: Title |
|---------|---|
| 2017-01 | J. Dölz and T. Gerig, M. Lüthi, H. Harbrecht and T. Vetter
<i>Efficient computation of low-rank Gaussian process models for surface and image registration</i> |
| 2017-02 | M. J. Grote, M. Mehlin, S. A. Sauter
<i>Convergence analysis of energy conserving explicit local time-stepping methods for the wave equation</i> |
| 2017-03 | Y. Bilu, F. Luca, D. Masser
<i>Collinear CM-points</i> |
| 2017-04 | P. Zaspel
<i>Ensemble Kalman filters for reliability estimation in perfusion inference</i> |
| 2017-05 | J. Dölz and H. Harbrecht
<i>Hierarchical Matrix Approximation for the Uncertainty Quantification of Potentials on Random Domains</i> |
| 2017-06 | P. Zaspel
<i>Analysis and parallelization strategies for Ruge-Stüben AMG on many-core processors</i> |
| 2017-07 | H. Harbrecht and M. Schmidlin
<i>Multilevel Methods for Uncertainty Quantification of Elliptic PDEs with Random Anisotropic Diffusion</i> |
| 2017-08 | M. Griebel and H. Harbrecht
<i>Singular value decomposition versus sparse grids: Refined complexity Estimates</i> |
| 2017-09 | J. Garcke and I. Kalmykov
<i>Efficient Higher Order Time Discretization Schemes for Hamilton-Jacobi-Bellman Equations Based on Diagonally Implicit Symplectic Runge-Kutta Methods</i> |
| 2017-10 | M. J. Grote and U. Nahum
<i>Adaptive Eigenspace Regularization For Inverse Scattering Problems</i> |
| 2017-11 | J. Dölz, H. Harbrecht, S. Kurz, S. Schöps and F. Wolf
<i>A Fast Isogeometric BEM for the Three Dimensional Laplace- and Helmholtz Problems</i> |
| 2017-12 | P. Zaspel
<i>Algorithmic patterns for \mathcal{H}-matrices on many-core processors</i> |