

# Accelerated Cyclic Reduction: A distributed-memory fast solver for structured linear systems

Gustavo Chávez<sup>a,\*</sup>, George Turkiyyah<sup>b</sup>, Stefano Zampini<sup>a</sup>, Hatem Ltaief<sup>a</sup>, David Keyes<sup>a</sup>

<sup>a</sup>Extreme Computing Research Center (ECRC), King Abdullah University of Science and Technology (KAUST), Thuwal 23955, Saudi Arabia

<sup>b</sup>Department of Computer Science, American University of Beirut (AUB), Beirut, Lebanon

## ARTICLE INFO

### Article history:

Received 30 December 2016

Revised 2 October 2017

Accepted 9 December 2017

Available online 15 December 2017

### Keywords:

Cyclic reduction  
Hierarchical matrices  
Fast direct solvers  
Elliptic equations

## ABSTRACT

We present Accelerated Cyclic Reduction (ACR), a distributed-memory fast solver for rank-compressible block tridiagonal linear systems arising from the discretization of elliptic operators, developed here for three dimensions. Algorithmic synergies between Cyclic Reduction and hierarchical matrix arithmetic operations result in a solver that has  $O(k N \log N (\log N + k^2))$  arithmetic complexity and  $O(k N \log N)$  memory footprint, where  $N$  is the number of degrees of freedom and  $k$  is the rank of a block in the hierarchical approximation, and which exhibits substantial concurrency. We provide a baseline for performance and applicability by comparing with the multifrontal method with and without hierarchical semi-separable matrices, with algebraic multigrid and with the classic cyclic reduction method. Over a set of large-scale elliptic systems with features of nonsymmetry and indefiniteness, the robustness of the direct solvers extends beyond that of the multigrid solver, and relative to the multifrontal approach ACR has lower or comparable execution time and size of the factors, with substantially lower numerical ranks. ACR exhibits good strong and weak scaling in a distributed context and, as with any direct solver, is advantageous for problems that require the solution of multiple right-hand sides. Numerical experiments show that the rank  $k$  patterns are of  $O(1)$  for the Poisson equation and of  $O(n)$  for the indefinite Helmholtz equation. The solver is ideal in situations where low-accuracy solutions are sufficient, or otherwise as a preconditioner within an iterative method.

© 2018 The Authors. Published by Elsevier B.V.  
This is an open access article under the CC BY license.  
(<http://creativecommons.org/licenses/by/4.0/>)

## 1. Introduction

Cyclic reduction, introduced in [1], is a direct solver for tridiagonal linear systems. It is effective for the solution of (block) Toeplitz and (block) tridiagonal matrices that arise from the discretization of elliptic PDEs [2,3]. For the constant-coefficient Poisson equation, since each of the blocks of the discretized system is Fourier diagonalizable, cyclic reduction can be used in combination with the fast Fourier transform (FFT) to deliver optimal complexity, as proposed in the FACR method [4]. However, in the presence of variable coefficients, the FFT-enabled version of cyclic reduction can not be used. The purpose

\* Corresponding author.

E-mail addresses: [gustavo.chavezchavez@kaust.edu.sa](mailto:gustavo.chavezchavez@kaust.edu.sa) (G. Chávez), [gt02@aub.edu.lb](mailto:gt02@aub.edu.lb) (G. Turkiyyah), [stefano.zampini@kaust.edu.sa](mailto:stefano.zampini@kaust.edu.sa) (S. Zampini), [hatem.ltaief@kaust.edu.sa](mailto:hatem.ltaief@kaust.edu.sa) (H. Ltaief), [david.keyes@kaust.edu.sa](mailto:david.keyes@kaust.edu.sa) (D. Keyes).

<https://doi.org/10.1016/j.parco.2017.12.001>

0167-8191/© 2018 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license.  
(<http://creativecommons.org/licenses/by/4.0/>)

of this work is to address the time and memory complexity growth in the presence of heterogeneous blocks with a variant called Accelerated Cyclic Reduction (ACR). The main observation is that elliptic operators have a hierarchical structure of off-diagonal blocks that can be approximated with low-rank matrices. Thus we approximate appropriate blocks of the initially sparse matrix with hierarchical matrices and operate on these blocks with hierarchical matrix arithmetics, instead of the usual dense operations, to obtain a direct solver of log-linear arithmetic and memory complexities. This philosophy follows recent work discussed below, but to our knowledge, this is the first demonstration of the utility of complexity-reducing hierarchical substitution in the context of cyclic reduction.

Cyclic reduction can be thought of as a direct Gaussian elimination on a permuted system that recursively computes the Schur complement of half of the unknowns until a single block remains or the system is small enough to be inverted directly. Schur complement computations have a complexity that is dominated by the cost of the inverse; by applying a red/black re-ordering of the unknowns, the linear system separates into two halves with block diagonal structure. This decoupling addresses the most expensive step of the Schur complement computation regarding operation complexity and does so in a way that launches independent subproblems. This concurrency feature, in the form of recursive bisection, can be naturally implemented in a distributed-memory parallel environment. The stability of the block cyclic reduction has been studied in [5], where the author presents error bounds for strictly and nonstrictly diagonally dominant matrices.

In order to simplify the description of the algorithm, in this work we consider structured linear systems arising from the discretizations or scalar PDEs on three-dimensional Cartesian grids. For three-dimensional problems of size  $N = n^3$ , where  $n$  is the number of discretization points in the linear dimension of the target domain, the synergy of cyclic reduction and hierarchical matrices leads to a parallel fast direct solver of  $O(k N \log N (\log N + k^2))$  arithmetic complexity, and  $O(k N \log N)$  memory footprint, where  $k \ll N$  represents the numerical rank of compressed blocks. This is in contrast to  $O(N^2)$  and  $O(N^{1.5})$  respectively, if hierarchically low-rank matrices are not used.

In this manuscript, we present ACR and its distributed-memory implementation, and we demonstrate its performance on a set of problems with various symmetry and spectral properties in three dimensions. These problems include the Poisson equation, the convection-diffusion equation, and the indefinite Helmholtz equation. We show that ACR is competitive in memory consumption and time-to-solution when compared to methods that rely on a global factorization and do not exploit the cyclic reduction structure.

### 1.1. Related work

Recent years have seen increasing interest in the use of hierarchical low-rank approximations to accelerate the direct solution of linear systems. In this section, we briefly describe some of this literature focusing primarily on efforts that target distributed-memory environments.

Arguably the most common approach for using hierarchical matrix representations in matrix factorizations is to use low-rank approximations to compress the dense frontal blocks that arise in the multifrontal variant of Gaussian elimination. The enabling property is that under proper ordering, many of the off-diagonal blocks of the Schur complement of discretized elliptic PDEs have an effective low-rank approximation [6] that improves the memory and arithmetic estimates of conventional multifrontal solvers [7]. Furthermore, there are efficient low-rank approximation methods to perform the necessary arithmetic operations and preserve the low-rank representation during the factorization and solution stages of the solver. Within this general approach, various methods that differ in the particular data-sparse format used and in the algorithms for the computation of low-rank approximations have been developed.

In Wang et al. [8] the authors investigate the use of the HSS format [9] to accelerate the parallel geometric multifrontal method, which results in a method known as the HSS-structured multifrontal solver (HSSMF). The general approach uses intra-node parallel HSS operations within a distributed-memory implementation of the multifrontal sparse factorization. This approach lowers the complexity of both arithmetic operations and memory consumption of the resulting HSS-structured multifrontal solver by leveraging the underlying numerically low-rank structure of the intermediate dense matrices appearing within the factorization process, driven by an optimal nested dissection ordering.

In a similar line of work, Ghysels et al. [10] also investigate a combination of the multifrontal method and the HSS-structured hierarchical format, extending the range of applicability of the solver to general non-symmetric matrices. Using the task-based parallelism paradigm, they introduce randomized sampling compression [11] and fast ULV HSS factorization [12]. Under the assumption of the existence of an underlying low-rank structure of the frontal matrices, randomized methods deliver almost linear complexity; this reduces the asymptotic complexity of the solver, which is mainly attributed to the frontal matrices near the root of the elimination tree. The effectiveness of these task-based algorithms in combination with a distributed-memory implementation of the multifrontal method is available in an early stage software release of the package STRUMPACK [10], which we will consider in the numerical experiments section of this article. The HSS format assumes a weak admissibility condition (see Section 2.1.1), which in practice requires the use of large numerical ranks even for approximations with modest relative accuracy. Consequently, this stresses the memory requirements and increases overall execution time.

The hierarchical interpolative factorization [13,14] is another method for finding low-rank approximations that has proved to be a fast solver for symmetric elliptic PDEs and integral equations. This decomposition relies on a “skeletonization” procedure to eliminate a redundant set of points from a symmetric matrix to further compress the dense fronts. The key step in skeletonization uses the interpolative decomposition of low-rank matrices to achieve a quasi-linear overall complexity in

factorization. The performance of hierarchical interpolative decomposition in a distributed-memory environment is reported in [15].

A fast direct method for high-order discretizations of elliptic PDEs has been proposed by Martinsson et al. [16–18]. The method is based on a multidomain spectral collocation discretization scheme and a hierarchy of nested grids, similar to nested dissection. It exploits analytical properties of elliptic PDEs to build Dirichlet-to-Neumann operators, by hierarchically merging these operators originating from smaller grids. When computations are done using the HSS data-sparse format, an asymptotic complexity of  $O(N^{4/3})$  can be reached. The high-order discretizations used in this method makes it quite powerful in practice as they allow it reach the same accuracy with fewer degrees of freedom compared to second order discretizations. A distributed-memory implementation of this algorithm is in progress.

Even though this approach has larger asymptotic estimates than the log-linear performance of the methods above, because of the high-order discretization of the PDE, this method is quite powerful in practice as they can reach the same accuracy with fewer degrees of freedom as compared to second order discretizations. A distributed-memory implementation of this algorithm is in progress.

The BLR format [19] has also been used to compress blocks into low-rank approximations to accelerate the factorization process of the multifrontal method. This format is compatible with numerical pivoting and is well-suited for the reuse of existing high-performance implementations of dense linear algebra kernels. Even though this format is not hierarchical, it has proven to be useful for a wide range of problems [20] within the distributed-memory implementation of the multifrontal method provided by the MUMPS library [21,22].

Rather than compressing and identifying individual blocks of the decomposition, another hierarchy-exploiting approach considers the system as a whole and seeks to construct a holistic decomposition of the full linear system. An example of such decomposition is the recursive computation of the inverse of a hierarchical matrix [23,24], or the computation of its Cholesky or LU factorization [25,26]. These methods have generally much higher prefactors than methods that compress individual matrix blocks of the factorizations and are not usually competitive for large-scale problems; as an example, we refer the reader to [23] for a discussion of the challenges of scaling the construction of the inverse of a hierarchical matrix.

Pouransari et al. approximate fill-in via low-rank approximations with the  $\mathcal{H}^2$  format; see [27]. This format guarantees linear complexity provided that blocks correspond to well-separated clusters and have a data-sparse property. The algorithm starts by recursively bisecting the computational domain, implicitly forming a binary tree. The leaf nodes correspond to independent subdomains, and the internal nodes correspond to Schur complements to computed with low-rank arithmetic operations. The bottom-up elimination process is performed with a procedure referred to as “extended sparsification” in which the original matrix dimension grows by introducing auxiliary variables but nonetheless remains sparse. Alternatively, elimination can be performed with an in-place algorithm that keeps the matrix size constant. A related method with similar strategies as in this work is the so-called “compress and eliminate” solver [28]. A recent extension of this line of work into a distributed memory environment documented in [29], demonstrates that concurrent processors can work on independent subdomains defined by their corresponding subgraphs, where interior vertices are eliminated concurrently. Communication is needed at the boundary vertices, but additional concurrency at the boundary is exploited through graph coloring.

## 1.2. Contributions

The contribution of this work is the development of a parallel, robust and efficient method for the solution of block tridiagonal linear systems, with emphasis on systems that arise from the discretization of elliptic PDEs. ACR is a fast solver in the sense that it has a log-linear arithmetic complexity in operations count and memory consumption. The algorithm arrives at the solution in a finite number of steps, rather than iteratively converging to a solution, which makes it a direct solver with a tunable accuracy. The fact that ACR is entirely algebraic extends its range of applicability to problems with an arbitrary coefficient structure including nonsymmetry within the block tridiagonal sparsity structure, subject to their amenability to rank compression. This entirely algebraic property gives the method robustness on problems that are challenging for iterative methods, while still maintaining asymptotic efficiency.

Two key features of the algorithm from a computational perspective are the simplicity of its parallelization and the regularity of its communication patterns in a distributed memory environment. The communication pattern is well-established beforehand and it is based on recursive bisection, as opposed to nested dissection with different block sizes at different levels of the factorizations. The amount of inter-node concurrency is proportional to the size of the blocks and it fits readily into a distributed-memory parallel environment. The algorithm also exhibits substantial intra-node concurrency, both in processing multiple blocks and within its hierarchical operations on individual blocks, which fits the multi-core architecture of modern supercomputers.

We demonstrate that our implementation is well suited for modern parallel multi-core systems and scalable in a distributed-memory environment. We also compare our implementation against other state-of-the-art direct solvers over a relevant class of problems and show competitive time to solution and memory requirements.

## 2. Preliminaries

In this section, we review the building blocks of the proposed solver, namely hierarchical low-rank approximations and the cyclic reduction algorithm.

## 2.1. Hierarchical matrices

A hierarchical matrix is a data-sparse representation that enables fast linear algebraic operations by using a hierarchy of off-diagonal blocks, each represented by a low-rank approximation, that can be tuned to guarantee an arbitrary precision. The approximation, sometimes referred to as compression, is performed via singular value decomposition, or with a related method that delivers a low-rank approximation with fewer arithmetic operations than the traditional SVD method. For the representation to be effective in terms of arithmetic operations and memory requirements, numerical ranks significantly smaller than the sizes of the various matrix blocks are required.

There are several hierarchical and non-hierarchical low-rank approximation formats available in the literature. In this work, we consider the  $\mathcal{H}$ -matrix format introduced by Hackbusch et al. in [30]. Being modular by design, ACR is not limited to the  $\mathcal{H}$ -format. In fact, the use of the  $\mathcal{H}^2$ -format would immediately translate to an additional reduction of one logarithmic factor in terms of arithmetic and memory complexity estimates, from  $O(k N \log N (\log N + k^2))$  to  $O(k N \log N)$  in terms of operations, and  $O(k^2 N \log N)$  to  $O(k N)$  in terms of memory requirements, however, we require a complete set of hierarchical matrix operations and fast construction, which at the time of this publication is still ongoing work within our group. Our implementation uses the  $\mathcal{H}$ -format arithmetic and its arithmetics operations provided by the HLibPro library. We refer to the reader to [31,32] for a discussion of the shared-memory scalability of these hierarchical matrix operations, their relative costs, and their performance on modern manycore architectures. HLibPro does not feature a distributed memory solver. We use its shared-memory kernels in combination with MPI to orchestrate parallel workload across nodes in a distributed memory environment as we will discuss in Section 4.

### 2.1.1. $\mathcal{H}$ -matrix construction

The structure of a hierarchical matrix in the  $\mathcal{H}$  format can be described by four components: an index set, a cluster tree, a block cluster tree, and the choice of an admissibility condition. The index set  $\mathcal{I} = \{0, 1, \dots, N - 1\}$  represents the number of degrees of freedom  $N$ . The cluster tree represents row/column groupings, and it is constructed by recursively subdividing the index set. Once the cluster tree is formed, the block cluster tree defines matrix sub-blocks over the index  $\mathcal{I} \times \mathcal{I}$ . Its leaves are either low-rank blocks or small dense ones. Finally, the admissibility condition determines whether a given block should be represented as a low-rank approximation or a dense block.<sup>1</sup>

The first step for the construction of an  $\mathcal{H}$ -matrix is the definition of the cluster tree of unknowns. In this work, since each block row of the sparse matrix represents a plane from a three-dimensional regular discretization, we leverage the geometry information by selecting a binary space partitioning strategy to cluster the unknowns considering the two-dimensional domain representing the planes.

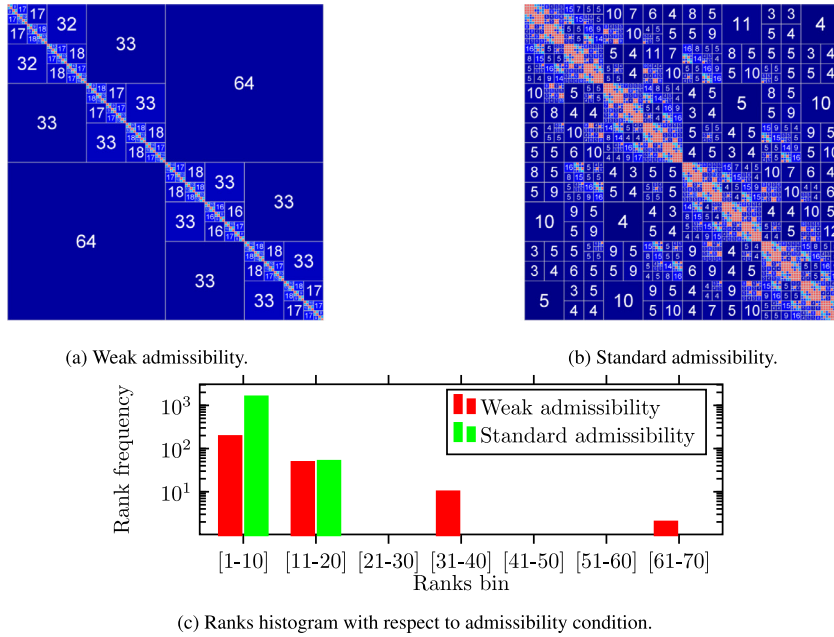
The next step is the definition of a block cluster tree for these two-dimensional domains, which together with the admissibility condition determines the structure of the hierarchical representation of the plane-block. We chose a standard admissibility condition, as opposed to the simpler weak admissibility condition, because it provides the flexibility of selecting a range of coarser or finer blocks, tuned by an admissibility parameter  $\eta$ . Weak admissibility refers to a matrix decomposition where the (1, 2) and (2, 1) blocks are single low-rank blocks and the (1, 1) and (2, 2) blocks are recursively decomposed in a similar way. On the other hand, standard admissibility allows a more refined blocking of the matrix; the  $\eta$  parameter appears in the inequality  $\min(\text{diameter}(\tau), \text{diameter}(\sigma)) \leq \eta \cdot \text{distance}(\tau, \sigma)$ , where  $\tau$  and  $\sigma$  denote two geometric regions defined as the convex hulls of two separate point sets  $t$  and  $s$  (nodes in cluster tree). A matrix block  $A_{ts}$  satisfying the previous inequality is represented in a low rank form.

The motivation for choosing a standard admissibility condition is that, by further refining the off-diagonal blocks, it is possible to achieve a similar accuracy with smaller numerical ranks, that are crucial to ensure economic memory consumption and overall high performance. The impact of the admissibility condition is illustrated in Fig. 1, which depicts the  $\mathcal{H}$ -inverse of the variable-coefficient two-dimensional Poisson operator discretized on a  $N = 64 \times 64$  grid using a finite difference scheme. In the right panel, the use of a few small dense blocks in the off-diagonal regions allows much smaller ranks to be used in the remaining low-rank blocks, without compromising accuracy.

Table 1 shows the storage gains by representing the inverse of a 2D Poisson problem with an  $\mathcal{H}$ -matrix with weak admissibility versus standard admissibility. Table 1 also shows the difference in terms of number of operations between these two structures. The cost of the  $\mathcal{H}$ -matrix inversion requires  $56C_{sp}^3kn(\log n + 1)^2 + 184C_{sp}k^3n(\log n + 1)$  operations, where  $k$  represents the average rank of the low-rank blocks, and  $C_{sp}$  represents the sparsity of the structure of the hierarchical matrix inverse, see [33]. Since the weak admissibility condition requires larger ranks than the standard admissibility condition, at scale, this tends to increase the memory requirements and the number of floating-point operations.

A low-rank approximation for a given off-diagonal block can be found in a variety of ways. Several strategies, ranging from randomized algorithms to heuristics for pivoting, are available in the literature. Every block of the  $\mathcal{H}$ -matrix stored as a low-rank approximation has the form of an outer product  $AB^T$ . The goal of efficient hierarchical matrix processing is to construct the best possible low-rank factorization as matrix operations are performed. This routine is often referred to as

<sup>1</sup> The word “block” is overloaded in this discussion. It is used to denote the partitions of the block tridiagonal coefficient matrix of the problem. It is also used to denote the partitioning of a matrix into low-rank and dense subdivisions. When necessary to avoid confusion, we will use the word “plane” or “plane-block” to refer to the first meaning.



**Fig. 1.**  $\mathcal{H}$ -inverse of the 2D Poisson operator, discretized with  $N = 64 \times 64$  grid points, using two different admissibility conditions. The number in each block is the numerical rank necessary to achieve a compression accuracy of  $1E-4$ . The color map is determined by the ratio of the numerical rank and the size of the block, deep blue indicates an effective low-rank approximation, while red depict dense blocks. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

**Table 1**

$\mathcal{H}$ -inverse of the 2D Poisson operator for  $N = 256^2$  grid points, using two different admissibility conditions. We document the memory and floating-point operations to build the  $\mathcal{H}$ -matrix inverse with weak and standard admissibility. The weak admissibility condition tends to require large ranks, which lead to increased memory requirements and more arithmetic operations than the standard admissibility condition. (Equivalent dense storage and arithmetic operations of the inverse operation would have required 2,147 MB and  $1.0E13$  operations.).

Operation	Format	Storage	Operations
Inverse	$\mathcal{H}$ (weak admissibility)	723 MB	$8.0E11$
Inverse	$\mathcal{H}$ (standard admissibility)	434 MB	$5.0E11$
Factorization	HSS (weak admissibility)*	40 MB	$3.3E07$

\*As a matter of comparison we show the equivalent storage requirements of the same problem by using the nested-basis HSS format which for this particular problem has an optimal complexity.

the compression step. For a comprehensive discussion of the construction of  $\mathcal{H}$ -matrices and its arithmetics, we refer the reader to [33].

## 2.2. Cyclic reduction

This section reviews the cyclic reduction algorithm in preparation for the following section describing the accelerated cyclic reduction variant that improves its arithmetic and memory complexity growth.

### 2.2.1. Model problem

Consider the seven-point stencil finite difference discretization with Dirichlet boundary conditions of the three-dimensional variable-coefficient Poisson equation on the unit cube.

$$-\nabla \cdot \kappa(x) \nabla u = f(x) \quad (1)$$



**Table 2**

Comparing the complexity estimates of storing and computing the inverse of a  $N \times N$  matrix block in dense format, versus approximating the matrix block with a hierarchical matrix with numerical rank  $k$ .

	Inverse	Storage
Dense Matrix	$\mathcal{O}(N^3)$	$\mathcal{O}(N^2)$
$\mathcal{H}$ Matrix	$\mathcal{O}(k N \log N (\log N + k^2))$	$\mathcal{O}(k N \log N)$

### 2.2.3. Back-substitution

Once elimination is completed, the solve stage starts from the last block of unknowns, as shown in Eq. (8):

$$D_0^{(3)} \cdot_{\mathcal{H}} u_0^{(3)} = f_0^{(3)}. \quad (8)$$

Once the solution at the last step  $u_0^{(3)}$  is computed, it is propagated backward in the hierarchy of the elimination tree. The formula to compute the solution at step  $q$  is given by

$$u^{(i)} = (D^{(i)})^{-\mathcal{H}} \cdot_{\mathcal{H}} (f^{(i)} - E^{(i)} \cdot_{\mathcal{H}} u^{(i+1)} - F^{(i)} \cdot_{\mathcal{H}} u^{(i+1)}). \quad (9)$$

This procedure continues until the solution of the entire linear system is computed.

Back-substitution is much more lightweight than the elimination algorithm regarding computation and communication volume, because it communicates parts of the solution in the form of vectors, and the only matrix operation performed is a matrix-vector multiplication. For large scale problems, this makes the solve phase orders of magnitude faster than the elimination phase. As with other direct solvers, the ability to efficiently solve for a given right-hand side given a factorization motivates the use of ACR for multiple right-hand sides at a minimal cost per new forcing term.

## 3. Accelerated Cyclic Reduction

This section describes how cyclic reduction can be used in combination with hierarchical matrices to result in a variant that improves the computational complexity and memory requirements of the classical cyclic reduction method.

### 3.1. Block-wise $\mathcal{H}$ -matrix approximation

ACR approximates each  $D_i$ ,  $E_i$  and  $F_i$  block of the original block tridiagonal matrix  $A$  given in Eq. (2) with a hierarchical matrix, and then proceeds with the cyclic reduction algorithm, as described in the previous section, by using hierarchical matrix arithmetics instead of the conventional dense linear algebra arithmetic operations.

In generating the structure of the hierarchical matrix representations of the blocks, we exploit the fact the domain is subdivided into  $n$  planes each consisting of  $n^2$  grid points and block rows of the matrix are identified with the planes of the discretization grid. We consider this geometry and use a two-dimensional planar bisection clustering when constructing each  $\mathcal{H}$ -matrix.

Cyclic reduction requires hierarchical matrix addition, subtraction, matrix-matrix multiplication, matrix-vector multiplication and matrix inversion. The relative accuracy of the approximation is specified during the compression of each block and while performing hierarchical matrix arithmetic operations. Committing to a given tolerance ensures that the numerical ranks are adjusted to preserve the specified accuracy during the elimination and solve phases. It is at the block level that the improvements in the complexity estimates take place.

Table 2 summarizes the advantages of a block-wise approximation of matrix blocks with  $\mathcal{H}$ -matrices in the computation of the inverse of a block, and its storage, as compared to their equivalent dense counterparts.

### 3.2. General algorithm

To simplify the exposition, we assume the size of the linear system is a power of two; the number of steps required by ACR is thus  $q = \log N$ . The size of the blocks for 2D problems is  $n^2$ .

As mentioned in Section 2.2, two procedures define cyclic reduction: elimination and back-substitution. The high-level algorithm of elimination is shown in listing Algorithm 1, whereas the high-level algorithm for back-substitution is shown in listing Algorithm 2. Even though Algorithms 1 and 2 show permutations and matrix operations at the level of the global system, our implementation operates at a per-block granularity, which means that permutations are part of the implementation's logic and that linear algebraic operations are performed block by block as shown in Eq. (6). This is possible since cyclic reduction preserves the block tridiagonal structure during elimination.

### 3.3. Sequential complexity estimates

Every cyclic reduction step requires two matrix-matrix multiplications, one matrix inversion and one matrix addition per block being eliminated. These kernels have arithmetic complexity of  $\mathcal{O}(k n \log n (\log n + k^2))$  operations [33]. For a

**Algorithm 1** ACR elimination.

---

```

0: Block-wise low-rank approximation of  $A$ :  $A^{(0)} = A$ 
1: for  $i = 0$  to  $q-1$  do
1: // Generate  $A^{(i+1)}$  block tridiagonal of size  $n/2^{i+1}$  using block-level operations (Equation 6)
1: // Requires  $O(kn \log n (\log n + k^2)/2^{i+1})$  operations
2:  $A^{(i+1)} = A_{22}^{(i)} -_{\mathcal{H}} A_{21}^{(i)} \cdot_{\mathcal{H}} (A_{11}^{(i)})^{-\mathcal{H}} \cdot_{\mathcal{H}} A_{12}^{(i)}$ 
2: // Forward substitution, requires  $O(kn \log n/2^{i+1})$  operations
3:  $f^{(i+1)} = f_2^{(i)} - A_{21}^{(i)} \cdot_{\mathcal{H}} (A_{11}^{(i)})^{-\mathcal{H}} \cdot_{\mathcal{H}} f_1^{(i)}$ 
4: end for

```

---

**Algorithm 2** ACR back-substitution.

---

```

1: Solve  $A^{(q)}u^{(q)} = f^{(q)}$ 
2: for  $i = q-1$  to  $0$  do
2: // Back-substitution, requires  $O(kn \log n/2^{i+1})$  operations
2: // This is performed at block-level (Equation 9)
3:  $u^{(i)} = (A_{11}^{(i)})^{-\mathcal{H}} \cdot_{\mathcal{H}} (f^{(i)} - A_{12}^{(i)} \cdot_{\mathcal{H}} u^{(i+1)})$ 
4: end for

```

---

**Table 3**Summary of the complexity of the  $\mathcal{H}$  matrix arithmetic operations.

Operation	Complexity
$A +_{\mathcal{H}} B$	$O(k^2 n \log n)$
$A \cdot_{\mathcal{H}} B$	$O(k n \log n (\log n + k^2))$
$A^{-\mathcal{H}}$	$O(k n \log n (\log n + k^2))$
$A \cdot_{\mathcal{H}} b$	$O(k n \log n)$

**Table 4**Summary of the sequential complexity estimates of the classic cyclic reduction method and the proposed variant, accelerated cyclic reduction;  $k$  represents the numerical rank of the approximation.

Method	Operations	Memory
Cyclic Reduction (CR)	$\mathcal{O}(N^2)$	$\mathcal{O}(N^{1.5} \log N)$
Accelerated Cyclic Reduction (ACR)	$\mathcal{O}(k N \log N (\log N + k^2))$	$\mathcal{O}(k N \log N)$

**Table 5**

Execution parameters, obtained relative residual, and ranks of the ACR factorization for the Poisson experiments.

$N$	$\mathcal{H}_\epsilon$	$\eta$	Leaf size	Relative residual	Average rank	Largest rank
$32^3$	8E-03	2	32	1.39E-02	3	4
$64^3$	1E-03	2	32	3.20E-02	3	5
$128^3$	1E-03	2	32	2.22E-02	4	7
$256^3$	1E-03	2	32	8.75E-02	4	7
$512^3$	1E-04	2	32	3.26E-02	5	11

problem size of  $N = n^3$  with  $n = 2^q$ , ACR requires  $n/2 + n/4 + n/8 + \dots \approx n$  steps to perform elimination. The most expensive computation in each step is the computation of an inverse of a block of size  $n^2 \times n^2$ , which in  $\mathcal{H}$ -format has a complexity of  $O(k n^2 \log n (\log n + k^2))$ , therefore, ACR results in a  $O(k N \log N (\log N + k^2))$  overall algorithm, with  $\mathcal{O}(k N \log N)$  memory requirements. Table 3 summarizes the complexity estimates of each of the  $\mathcal{H}$  matrix operations involved in ACR. Table 4 summarizes the complexity estimates of the classical cyclic reduction algorithms without exploitation of equal blocks versus ACR.

Because ACR effectively uses hierarchical representations only for a set of regular two-dimensional problems, the resulting constants appearing in the asymptotic complexity estimates tend to be smaller, as a virtue of lower rank requirements, and make it feasible to perform large scale computations. For instance, our limited experiments show that for the 3D Poisson problem (Table 5) ACR requires substantially lower numerical ranks than the ranks reported in the HSSMF literature [8,10].

In terms of practical usage, ACR has different concurrency properties than  $\mathcal{H}$ -LU or multifrontal HSS, enabling different amounts of independent work to be performed. The regularity of the computational patterns of ACR is valuable in terms of the ability to efficiently use current and future hardware architectures.

#### 4. Parallel accelerated cyclic reduction

This section describes how to leverage the concurrency features of the accelerated cyclic reduction method in a distributed-memory parallel environment.

##### 4.1. Parallel implementation

The parallel ACR elimination and back-substitution algorithms are listed in [Algorithms 3](#) and [4](#), respectively.

---

#### Algorithm 3 Parallel ACR elimination.

---

```

1:  $j =$  Processor number
2: parallel for at all processors  $j, j \in 0 : 2^q - 1$ 
3:   Block-wise conversion to  $\mathcal{H}$ -matrix of  $A = \text{tridiagonal}(E_j^{(1)}, D_j^{(1)}, F_j^{(1)})$ 
4: end parallel for
5: for  $i = 1$  to  $q$  do
6:   parallel for at  $j$  even,  $j \in 0 : 2^{q-i} - 1$ 
7:     Compute  $(D_j^{(i)})^{-\mathcal{H}}$ 
8:     Communicate  $E_j^{(i)}, (D_j^{(i)})^{-1}, F_j^{(i)}, f_j^{(i)}$  to processors  $j - 1$ 
9:     Communicate  $E_j^{(i)}, (D_j^{(i)})^{-1}, F_j^{(i)}, f_j^{(i)}$  to processors  $j + 1$ 
10:   end parallel for
11:   parallel for at  $j$  odd,  $j \in 0 : 2^{q-i-1} - 1$ 
12:     Compute  $E_j^{(i+1)}, D_j^{(i+1)}, F_j^{(i+1)}, f_j^{(i+1)}$  from Equation 6
13:   end parallel for
14: end for

```

---

A number of concurrency features of the algorithms are evident. Each block row, identified by a plane in the discretiza-

---

#### Algorithm 4 Parallel ACR back-substitution.

---

```

1:  $n = 2^q$ 
2:  $j =$  Processor number
3: for  $i = q$  to 1 do
4:   parallel for at  $j, j \in 0 : 2^{q-i} - 1$ 
5:     Compute  $u_j^{(i)}$  from Equation 9
6:     Communicate  $u_j^{(i)}$  to processors  $j - 1$ 
7:     Communicate  $u_j^{(i)}$  to processors  $j + 1$ 
8:   end parallel for
9: end for

```

---

tion, is assigned to an MPI rank. This decomposition allows the initial conversion of each block into an  $\mathcal{H}$ -matrix in an embarrassingly parallel manner. The  $q = \log n$  levels of Schur complement computation exploit concurrent execution in two ways:

- The inverse of the block  $A_{11}$  of Eq. (4) can be computed concurrently in a block-wise fashion since  $A_{11}$  is block diagonal. This computation is embarrassingly parallel.
- Computing the Schur complement requires two matrix-matrix multiplications and one matrix addition. Since the linear system partition is formed out of matrix blocks, the computation of these block matrix-matrix multiplications and block matrix-addition can also be computed concurrently.

Fig. 2 depicts the concurrency through the various levels in ACR elimination. We note here that the ACR decomposition strategy bears a similarity to the slice decomposition [34], and also relate to the sweeping preconditioner strategy [35], with the key distinction being that rather than sweeping through the domain, ACR eliminates several planes at once, concurrently.

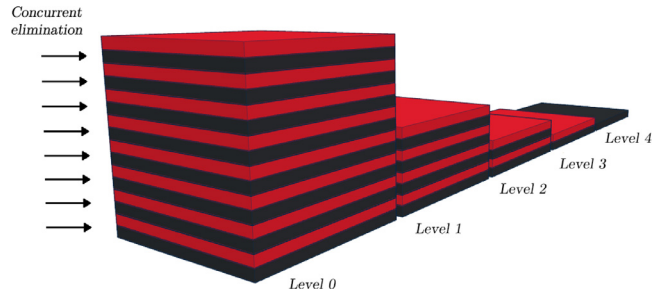


Fig. 2. Concurrency in ACR elimination for the 16-planes case. Level 0 can eliminate eight planes concurrently thus reducing the problem size to the next level by two; this process continues until one plane is left.

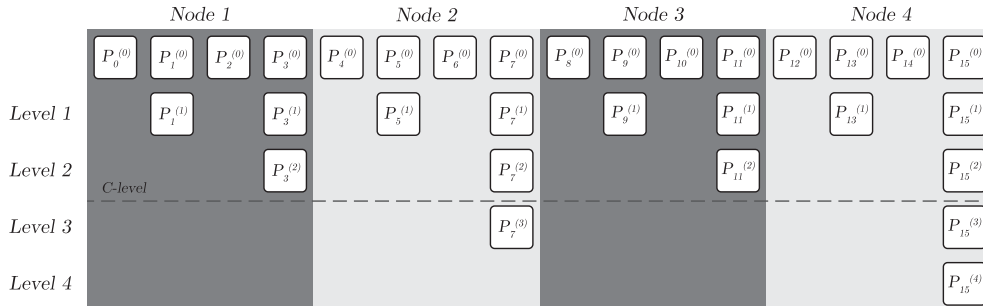


Fig. 3. Distribution of multiple planes per physical compute node for an example with  $n=16$  and  $p = 4$ .

4.2. Inter-node communication

In the current implementation, each plane is assigned to an MPI rank, and multiple planes are assigned to compute nodes. Let  $p$  be the number of physical compute nodes each storing  $n/p$  planes at the beginning of the factorization. After  $r$  steps of ACR, each compute node holds  $n/(2^r p)$  planes. At level  $r = \log(n/p)$ , a coarse level called the C-level, every node holds a single plane only. The remaining  $\log p$  steps of ACR beyond the C-level leave some compute nodes idle as illustrated in Fig. 3.

Distributed-memory communication occurs just at inter-node boundaries thanks to sorting at every step of the factorization, as the computation of the Schur complement for plane  $P_j$  just requires planes  $P_{j-1}$  and  $P_{j+1}$ , see Fig. 3. Thus up to the C-level there are  $O(p)$  communication messages per step, each transmitting planes of size  $O(kn^2 \log n)$ . Beyond the C-level, there are  $O(p/2 + \dots + 1) \approx O(p)$  communications messages, adding up to a total communication volume of  $O(k p n^2 \log n (\log \frac{n}{p} + 1))$  for ACR. The communication pattern with its bottom-up binary tree structure is depicted in Fig. 4.

4.3. Parallel time complexity

The regularity of the ACR algorithm makes it straightforward to estimate the parallel time of the factorization and assess its scalability characteristics. Consider the longest computing node which executes  $\log n$  ACR steps. In the  $\log(n/p)$  steps preceding the C-level, this node processes  $n/(2p) + n/(4p) + \dots + 1$  block rows in sequence. Beyond the C-level, it processes a single block row in every one of the sequential  $\log p$  steps. This results in an asymptotic parallel time complexity for ACR of  $O(kn^2 \log n (\log n + k^2)(n/p + \log p))$ . The sequential computational time gets reduced by the number of parallel compute nodes  $p$ , but at the expense of an additional  $\log p$  factor that inhibits perfect strong scaling. Fortunately, the amount of work above the C-level that introduces this  $\log p$  factor left is small and grows only as  $n^2 = N^{2/3}$ .

Finally, we note that beyond the parallelism across distributed computing nodes, there is additional concurrency available at the node level. This additional level of parallelism is possible, not only because elimination and back-substitution for multiple block rows can proceed concurrently, but also because parallel variants of the hierarchical matrix arithmetics can be used in performing operations on individual blocks. The two levels of intra-node parallelism are shown schematically in Fig. 5. In practice, programming models based on tasks and directed acyclic graphs have proven to be effective to parallelize hierarchical matrix arithmetics [10,32], but the optimal allocation of the multiple cores of a compute node to either block row processing or to individual operations on single blocks requires tuning. We do not describe this aspect of the parallel implementation further here.

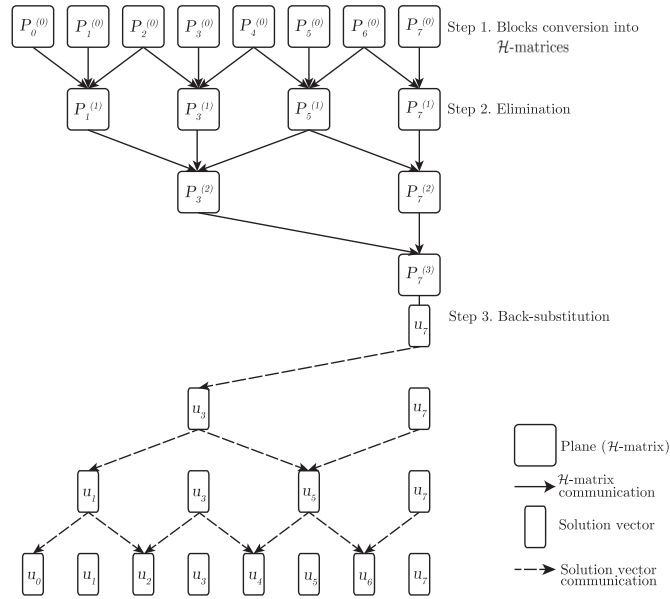


Fig. 4. Communication pattern for the 8-planes case.  $P_j$  depicts the plane being eliminated, and  $u_j$  its corresponding solution.

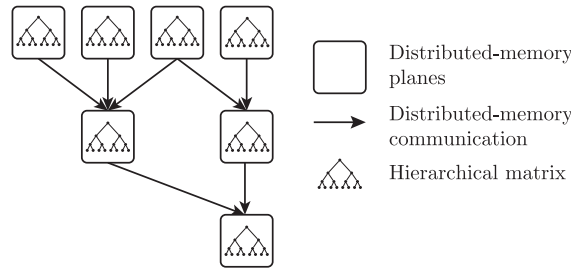


Fig. 5. Parallel ACR elimination tree depicting two-levels of concurrency using distributed-memory parallelism to distribute concurrent work across compute nodes, and shared-memory parallelism to perform  $\mathcal{H}$ -matrix operations within the nodes.

Table 6

Execution parameters, obtained relative residual, and ranks of the HSSMF factorization for the Poisson experiments.

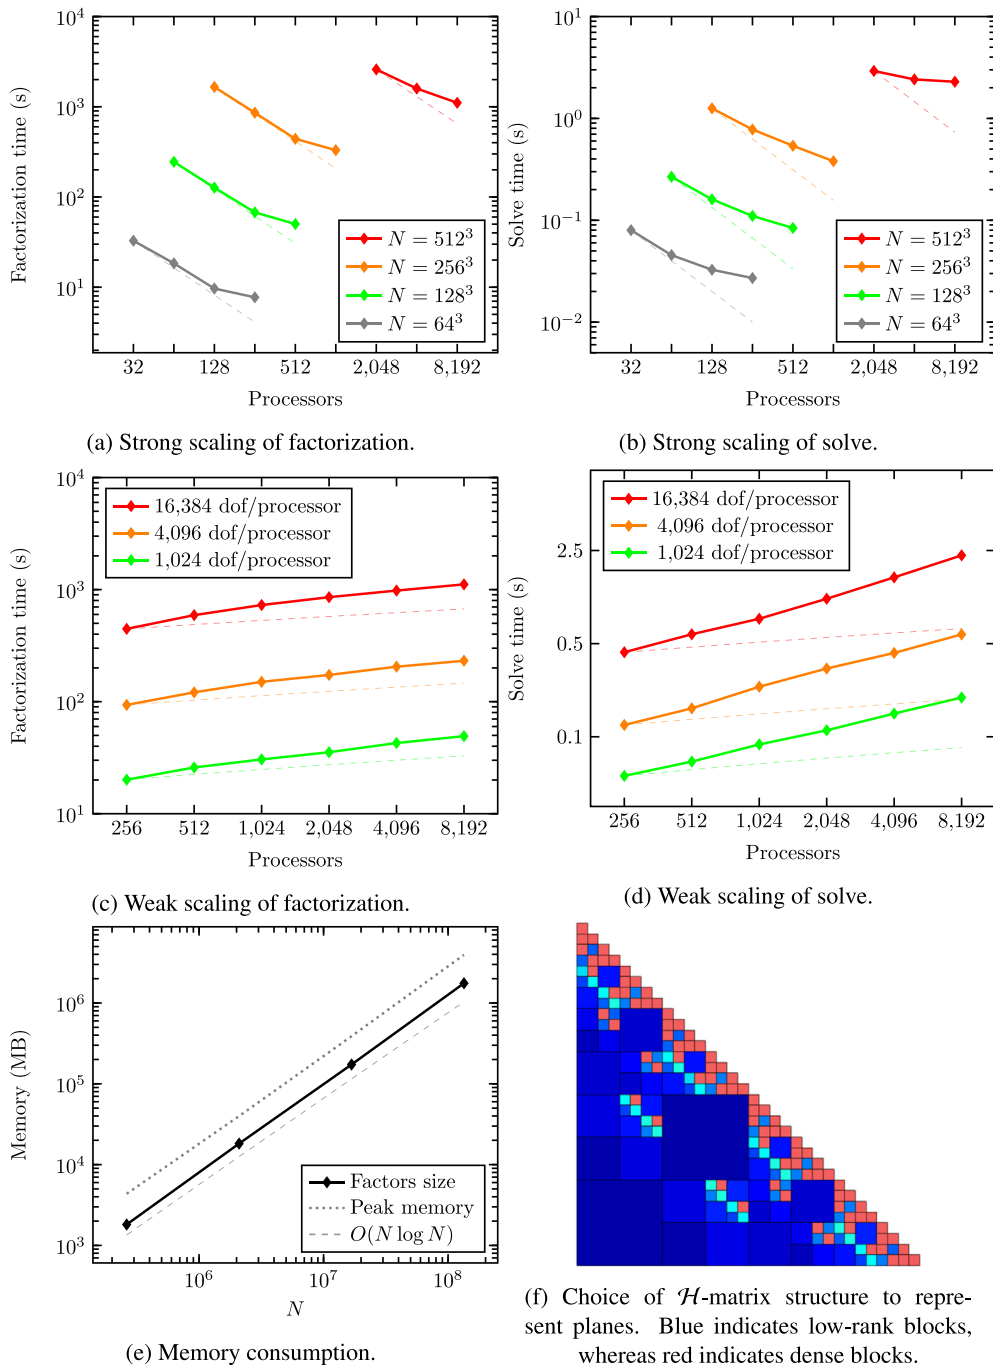
$N$	Compression tolerance	Relative residual	Leaf size	Minimum front size	Largest rank
$32^3$	1E-02	4.41E-02	128	256	82
$64^3$	1E-03	2.65E-02	128	1024	243
$128^3$	1E-03	8.40E-02	128	4096	532

### 5. Numerical results

This section documents the parallel performance and scalability of ACR in a distributed-memory environment. The source code is written in C and compiled with the Intel C compiler v15. External libraries utilized in the reference implementation include HLIBpro v2.2 with Intel TBB [31,36], and the sequential version of the Intel Math Kernel Library [37]. Experiments are conducted on the Cray XC40 Shaheen supercomputer at the King Abdullah University of Science & Technology. Each node has 128GB of RAM and two Intel Haswell processors, each with 16 cores clocked at 2.3Ghz.

To provide a baseline of performance we consider the solution of the same linear systems with STRUMPACK [10] v1.0.3, the open-source implementation of the HSS-structured multifrontal solver (HSSMF) developed at the Lawrence Berkeley National Laboratory. The HSSMF method can solve a broader class of linear systems compared to ACR, but the comparison is still of interest, as STRUMPACK is among the few available implementations of distributed-memory fast direct solvers that exploit hierarchically low-rank approximations.

The tuning parameters of ACR include the choice of the leaf node size  $n_{\min}$  for the  $\mathcal{H}$  matrices, the threshold parameter  $\eta$  used to decide which blocks will be approximated with a low-rank factorization, or as a dense, full-rank, block, and the accuracy of the approximation  $\mathcal{H}_\epsilon$  for the construction and algebraic operations of the  $\mathcal{H}$  matrices. The tuning parameters for



**Fig. 6.** Parallel scalability and memory consumption of ACR for the solution of the Poisson equation. Fig. 6f depicts the structure of the  $\mathcal{H}$ -matrices utilized for all the ACR blocks (in this case, extracted from the  $N = 32^3$  problem). The deep blue color indicates an effective compression, while red blocks indicate no compression. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

STRUMPACK include how many matrices from the nested-dissection elimination tree will be approximated as HSS, which is controlled by specifying the threshold at which frontal matrices will be represented as HSS matrices, the compression accuracy for the HSS matrices, and the minimum leaf size of the HSS frontal matrices. We recall here that the HSS matrix format uses the so-called weak admissibility condition, whereas ACR uses a standard admissibility condition, which does not limit the use of dense blocks exclusively at the matrix diagonal. Additionally, we also consider the algebraic multigrid (AMG) implementation of hypre [38,39]. Comparison experiments are set to deliver a solution with a relative error tolerance as

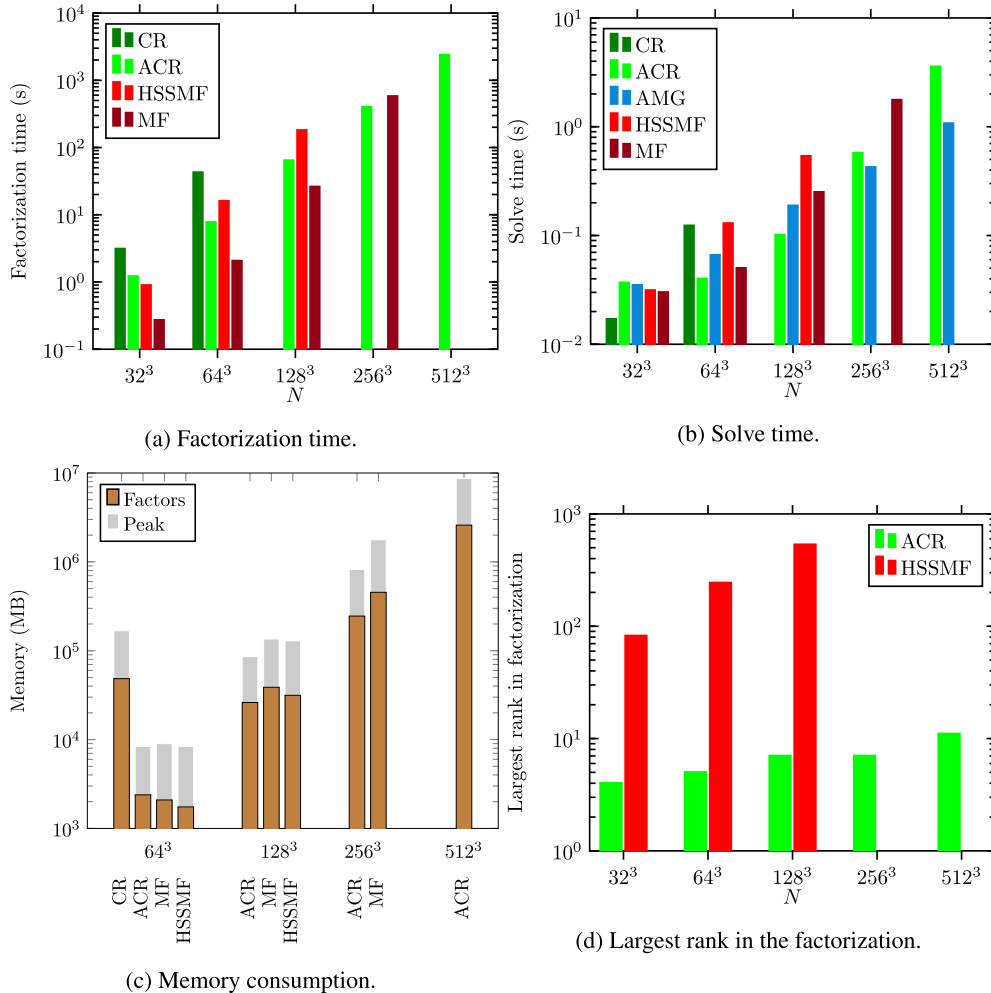


Fig. 7. Performance of the factorization and solve phases of ACR for the Poisson problem.

Table 7

Iterative solution of a  $N = 128^3$  Poisson problem with the conjugate gradients method and ACR preconditioner. Relative residual of the solution is  $1E-6$  in all cases.

$\mathcal{H}_\epsilon$	Factors (MB)	Largest rank	Factorization	Apply	Total time	Iterations
6E-01	17,280	31	18.55	0.050	20.72	43
3E-01	19,385	31	21.33	0.053	23.14	34
1E-01	22,328	31	26.56	0.058	28.01	25
1E-02	26,687	37	51.24	0.064	51.94	11
1E-03	32,212	53	89.32	0.104	89.73	4
1E-04	39,181	71	149.06	0.127	149.44	3

$\|Ax - b\|_2 / \|b\|_2 \approx 10^{-2}$ . For further comparisons, we also consider the multifrontal (MF) implementation of STRUMPACK, and our cyclic reduction (CR) implementation with dense matrix blocks.

### 5.1. Poisson equation

We consider a sequence of Poisson problems of up to  $512^3 \approx 134M$  unknowns, which is considered very large for this type of “direct” (as opposed to iterative) methods. We feature the Poisson equation with homogeneous Dirichlet boundary conditions in the unit cube, i.e.

$$-\nabla^2 u = 1, \quad \mathbf{x} \in \Omega = [0, 1]^3, \quad u(\mathbf{x}) = 0, \quad \mathbf{x} \in \Gamma, \tag{10}$$

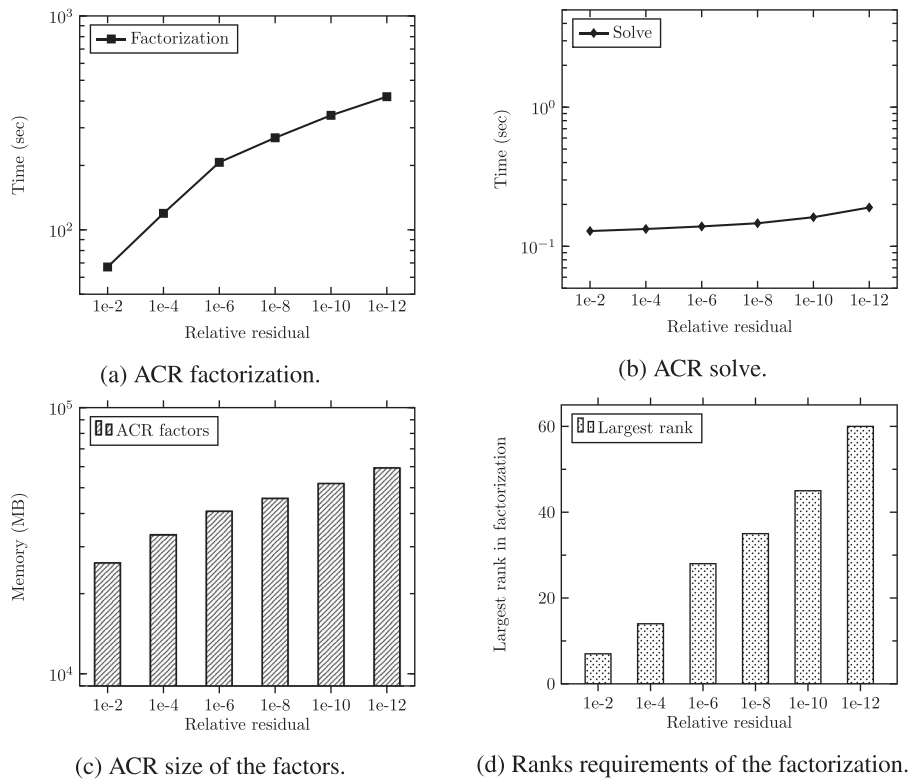


Fig. 8. Controllable accuracy solution of ACR for a  $N = 128^3$  Poisson problem.

discretized with the 7-point finite-difference star stencil, which leads to a symmetric positive definite linear system.

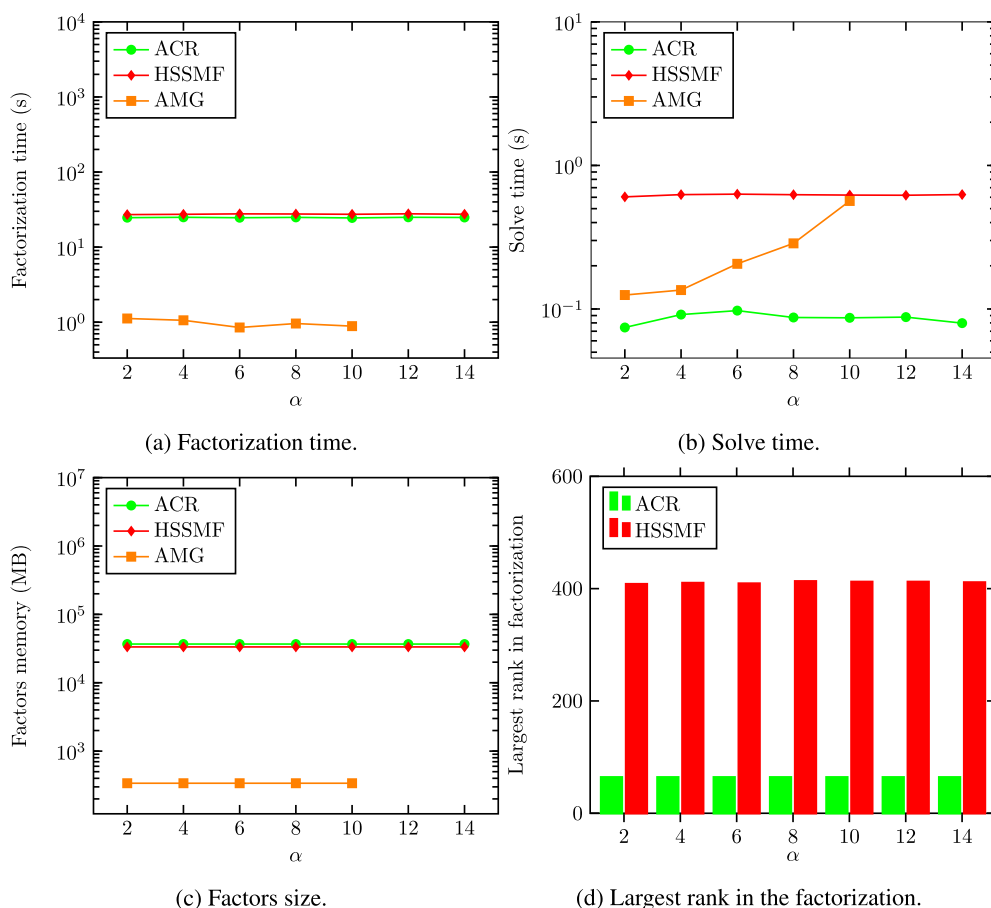
Although this problem can be solved with ACR, other methods such as multigrid or FFTs are ordinarily used instead; we consider it to report on a standard and well-known problem, to facilitate the exposition of ACR. Furthermore, the discretization of the Poisson equation has all positive eigenvalues with rapid decay in off diagonal, making it also an ideal case for hierarchically low-rank approximations analysis.

Fig. 6a and b show the total time in seconds for the factorization and solve phases of ACR in a strong scaling setting; dashed lines indicate ideal scaling. Ideal scaling of the factorization stage deteriorates at large processor counts as factors such as communication volume and hardware latency begin to play a significant role; the same factors tend to dominate even more during the solve phase, being the latter a sequence of fast  $\mathcal{H}$ -matrix-vector multiplications with limited availability of communication/computation overlap.

Fig. 6c and d depict the results of a weak scaling test for ACR with different numbers of degrees of freedom per processor, along with the ideal weak scaling reference lines depicted as dashed curves. The timings deviate from the ideal scaling due to the inherently load imbalance of the recursive bisection strategy of cyclic reduction as some processors become idle towards the end of the reduction. Communication latency further impacts the solve stage at large core counts due to the lower arithmetic intensity of this stage.

Fig. 6e depicts the memory requirements to store the ACR factorization, together with the expected asymptotic memory usage as  $O(N \log N)$ . We stress that the maximum rank of the factored matrices varies from 5 to 10 within all the combinations of problem sizes/number of processors considered in the strong and weak scaling tests (data not shown). Fig. 6f depicts the structure of the  $\mathcal{H}$ -matrices used to represent each plane, with the choice of standard admissibility condition. Dark blue blocks denote a low ratio between the numerical rank of the approximation and the full rank of the block, whereas red block indicates non-admissible blocks stored in dense format. For visualization purposes, the figure was taken from the  $N = 32^3$  problem, and represents the last diagonal block during the elimination phase of ACR. The prevalence of dark blue blocks indicate a good relative compression of each block, since the ratio of numerical rank of the approximation and the actual block size is very small. Most of the red blocks are clustered near the diagonal, where the smallest blocks reside.

Fig. 7 compares all solvers under consideration for a set of Poisson problems that range from  $N = 32^3$  to  $N = 512^3$  unknowns, with processor counts increased from 256 to 4096 respectively. We document the execution parameters, obtained relative residual, and ranks of the ACR and HSSMF factorization in Tables 5 and 6. We report factorization times in Figs. 7a



**Fig. 9.** Robustness of ACR for convection-diffusion problem. In convection dominated problems (large  $\alpha$ ), AMG fails to converge while direct solvers maintain a steady performance.

showing that ACR can competitively tackle these problems. Similarly, the solve timings in Fig. 7b show that ACR is able to solve for a given right-hand size in comparable times as the other methods under consideration. Fig. 7c documents the size of the factors required by the factorizations, and it shows that the cyclic reduction method (CR) cannot solve problems as small as  $N = 128^3$  due to memory limitations. Additionally, we report the peak memory utilization of each solver using the library PAPI v5.5 [40], which shows the largest memory usage that each solver required to produce the factorization. Also, the experiments confirm that the HSSMF method requires less memory to store its factors than the multifrontal method (MF). However, as Fig. 7d shows, the HSSMF method requires higher ranks than ACR, which translated into a larger size of the factors and prohibited the execution of HSSMF for problems of  $N = 256^3$  and above. The experiments show that ACR requires only  $O(1)$  ranks, as opposed to the  $O(n)$  rank requirements of the HSSMF factorization.

As expected for this particular problem, multigrid is the method of choice concerning performance and memory footprint for a single right-hand-side. However, for multiple right-hand-sides, the ability to reuse the factorization could give the advantage to solvers based on factorization. The factorization times for ACR and HSSMF are comparable, with the setup stage of HSSMF being faster for smaller problems; the smaller ranks required by ACR instead lead to a faster factorization step with large problem sizes and faster time to solution.

While ACR and HSSMF solvers can deliver a more accurate solution as direct solvers (i.e. without iterative procedures), this comes at the expense of more time and memory; it is common practice that this factorization is then used as a preconditioner or passed to an iterative refinement procedure. Numerical experiments confirm that ACR could be used as a direct solver if we tune its parameters with a higher accuracy for its  $\mathcal{H}$ -matrix representations and operations, as depicted in Fig. 8, at the expense of modest rank increases, albeit with higher memory requirements and time to solution. However, as Table 7 shows, a low-accuracy factorization in combination with an iterative procedure is best to minimize the total time-to-solution.

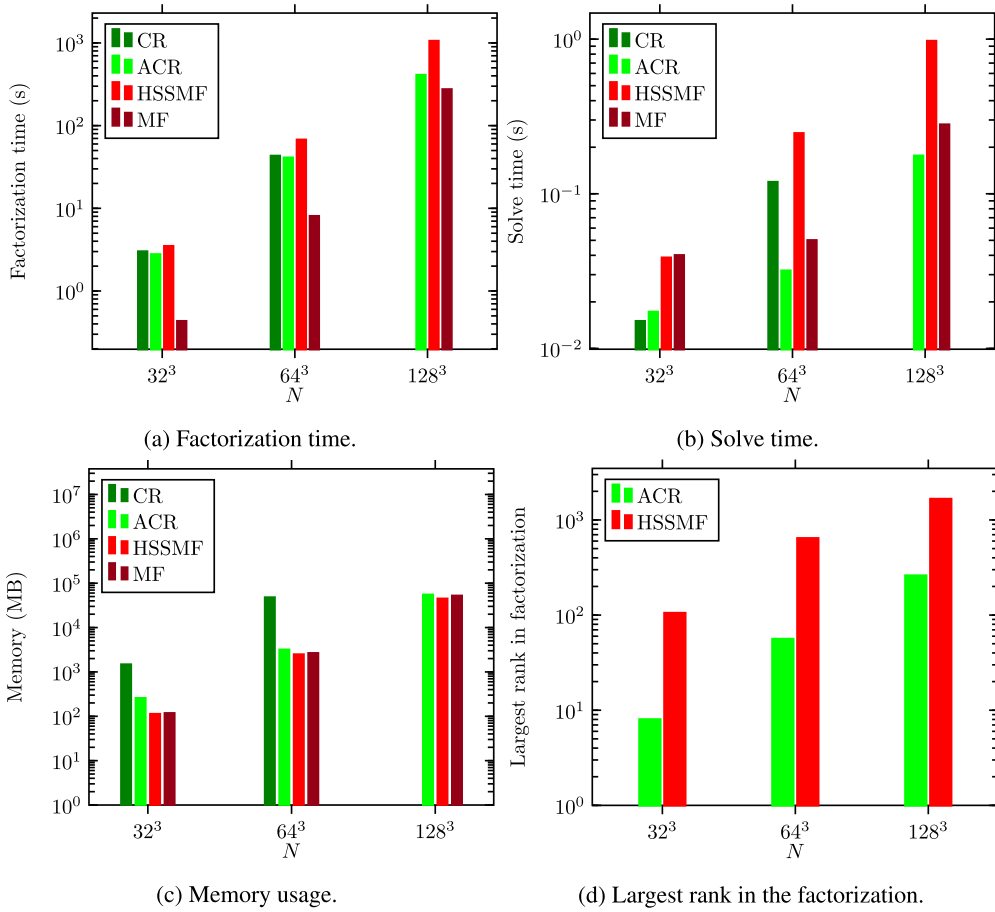


Fig. 10. Solution of increasingly larger indefinite Helmholtz problems consistently discretized with 12 points per wavelength.

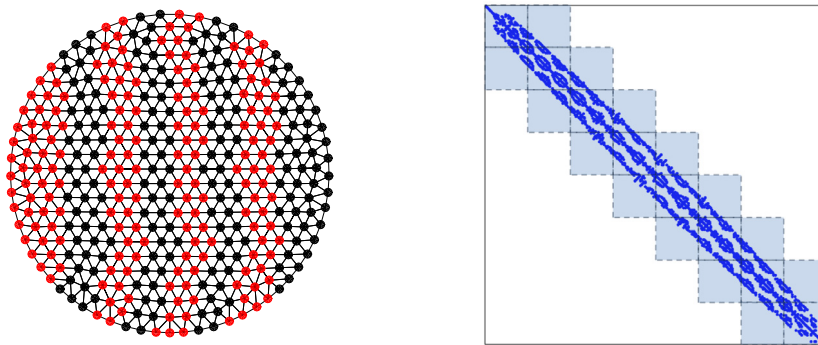


Fig. 11. Partitioning of an unstructured mesh that produces a block tridiagonal matrix structure, for the application of ACR.

5.2. Convection-diffusion equation

We next consider a standard convection-diffusion problem

$$-\nabla^2 u + \alpha b(\mathbf{x}) \cdot \nabla u = f(\mathbf{x}), \quad \mathbf{x} \in \Omega = [0, 1]^3,$$

$$b(\mathbf{x}) = \begin{bmatrix} \sin(a 2\pi x) \sin(a 2\pi (1/8 + y)) + \sin(a 2\pi (1/8 + z)) \sin(a 2\pi x) \\ \cos(a 2\pi x) \cos(a 2\pi (1/8 + y)) + \cos(a 2\pi (1/8 + y)) \cos(a 2\pi z) \\ \cos(a 2\pi x) \cos(a 2\pi (1/8 + z)) + \sin(a 2\pi (1/8 + y)) \sin(a 2\pi z) \end{bmatrix}, \tag{11}$$

discretized with a 7-point upwind finite difference scheme, that leads to a non-symmetric linear system which is challenging for classical iterative solvers, especially when the convection term dominates the equation. The  $b(\mathbf{x})$  term we consider is a

**Table 8**

Execution parameters, obtained relative residual, and ranks of the ACR factorization for the Helmholtz experiments.

$N$	$\mathcal{H}_\epsilon$	$\eta$	Leaf size	Relative residual	Average rank	Largest rank
$32^3$	5E-03	4	32	1.67E-02	5	8
$64^3$	5E-08	8	32	2.63E-02	30	56
$128^3$	5E-13	16	32	1.07E-02	113	260

**Table 9**

Execution parameters, obtained relative residual, and ranks of the HSSMF factorization for the Helmholtz experiments.

$N$	Compression tolerance	Relative residual	Leaf size	Minimum front size	Largest rank
$32^3$	5E-03	5.32E-02	128	256	105
$64^3$	1E-04	6.08E-02	128	1024	641
$128^3$	1E-06	1.13E-02	128	4096	1659

three-dimensional generalization of the two-dimensional vortex flow proposed by Wessel et. al.[41]. We adjust the forcing term and boundary conditions to meet the exact solution

$$u(\mathbf{x}) = \sin(\pi x) + \sin(\pi y) + \sin(\pi z) + \sin(3\pi x) + \sin(3\pi y) + \sin(3\pi z),$$

as proposed by Gupta and Zhang [42], as it is an archetypal challenging problem for multigrid methods.

To demonstrate the robustness of ACR and HSSMF for this problem, we fix the number of degrees of freedom at  $N = 128^3$  and we increase the dominance of the convective term; results are reported in Fig. 9. Consistently with the Poisson problem, multigrid methods remains the method of choice for diffusion dominated problems in terms of time to solution; however, when  $\alpha$  is increased, the performance of AMG deteriorates. On the other hand, both ACR and HSSMF prove to be able to solve convection-dominated problems, with ACR being consistently faster than HSSMF particularly in the back-substitution phase. The size of the factors generated by ACR and HSSMF are comparable, with ACR using significantly smaller ranks.

### 5.3. Helmholtz equation

We finally consider the indefinite Helmholtz equation with Dirichlet boundary conditions on the unit cube, i.e.

$$-(\nabla^2 u + \kappa^2 u) = 1, \quad \Omega = [0, 1]^3, \quad (12)$$

discretized with the 27-point trilinear finite element scheme on hexahedra. Results for ACR and HSSMF are reported in Fig. 10. The parameter  $\kappa$  is chosen to obtain a sampling rate of approximately 12 points per wavelength, specifically  $\kappa = \{16, 32, 64\}$  respectively, corresponding to approximately  $10 \times 10 \times 10$  for the  $N = 128^3$  problem. As opposed to the positive definite Helmholtz equation which models phenomena similar to diffusion, the indefinite variant, commonly denoted as the wave Helmholtz equation, has a solution that is oscillatory in nature. Multigrid methods are known to diverge without specific customizations for high-frequency Helmholtz problems, which we also confirmed via experimentation. For a detailed examination of the difficulties of solving the Helmholtz equation with classical iterative methods we refer the reader to [43].

We document the execution parameters, obtained relative residual, and ranks of the ACR and HSSMF factorization in Tables 8 and 9. Numerical experiments show that ACR features consistently lower factorization and solve times than HSSMF, as can be seen in Fig. 10a and b. The size of the factors of ACR and HSSMF are comparable, with a slightly higher memory requirements of ACR due to performance-oriented tuning, see Fig. 10c. Furthermore, as also shown in Section 5.1, HSSMF required less memory than MF, and CR quickly runs out of memory for problems larger than  $N = 64^3$ . Finally, the largest rank of ACR is consistently lower than that of HSSMF, even though both solvers require  $O(n)$  ranks, as shown in Fig. 10d. Nevertheless, lower ranks lead to faster time-to-solution in favor of ACR.

## 6. Conclusions and future work

We present a novel fast direct solver, Accelerated Cyclic Reduction, for block tridiagonal linear systems which commonly arise in the discretization of elliptic operators. The elimination strategy is based on a red/black ordering of the blocks that logically divides the grid into planes, approximates matrix blocks representing these planes with  $\mathcal{H}$ -matrices, and proceeds with elimination using hierarchical matrix operations. ACR achieves log-linear arithmetic complexity of  $O(k N \log N (\log N + k^2))$  and memory requirements of  $O(k N \log N)$  by approximating each block with a hierarchical matrix whose structure is refined using a spatial partitioning of the planar grid sections, employing a strong admissibility criterion that effectively limits the ranks of individual low rank blocks in the hierarchical matrix representations, and operating with hierarchical matrix arithmetics throughout. The average rank  $k$  of the blocks inside the hierarchical matrix representations controls the accuracy of the approximation and grows only modestly with problem size. A fair agreement with the rank estimate of [6] was found for the 3D Poisson equation of  $O(1)$  (Table 5), and for the 3D Helmholtz equation  $O(n)$  (Table 8).

The concurrency features of ACR are among its most important strengths. The regularity and structure of the decompositions allow efficient load balance. These features are demonstrated in a distributed-memory environment with numerical experiments that study the strong and weak scalability of our implementation. We provide a reference for performance and memory consumption using comparisons with state-of-the-art open-source implementations of the HSS-structured multi-frontal solver from the STRUMPACK library, and algebraic multigrid from hypre.

ACR, being essentially a direct solver with tunable accuracy, can tackle problems that lack definiteness, such as the indefinite high-frequency Helmholtz equation, or symmetry, such as the convection-diffusion equation. For these problems, stock versions of algebraic multigrid fail to produce convergent schemes. We demonstrated the robustness of ACR in dealing with such problems over a range of problem sizes and parameters.

While multigrid methods are generally superior for scalar problems possessing smoothness and definiteness, direct factorization methods such as ACR and HSSMF benefit where multiple right-hand sides are involved, as the time to solve per extra forcing term is orders of magnitude smaller than the factorization, which can be reused. The smaller ranks  $k$  of ACR result in solution times per new right-hand side that are smaller than those of HSSMF.

Although having the same asymptotic complexity as other solvers that use general hierarchical matrix representations in their factorizations, such as  $\mathcal{H}$ -LU, ACR has fundamentally different algorithmic roots which enable a novel alternative for a relevant class of problems with competitive performance, increasing concurrency as the problem grows and almost optimal memory requirements. Moreover, to the best of our knowledge, this is the first distributed-memory implementation of the synergies of cyclic reduction and hierarchical matrices, which scales up to 8192 cores for problems up to  $N = 512^3$  degrees of freedom.

ACR has been demonstrated for a regular grid discretization, but the generalization to arbitrary grids is possible and we intend to explore it in the future. Such a generalization would require an ordering of the mesh that produces a sequence of thin elongated regions (in 2D or 3D) where every region has only two neighbors so that the block tridiagonal structure is preserved. Such an ordering might be produced via a breadth-first search traversal of the mesh as shown in Fig. 11. In the unstructured case, the diagonal blocks do not necessarily have the same size, and the off-diagonal blocks might be of rectangular shape. The main algorithmic implication is that each block will now have its own hierarchical matrix structure generated from the geometry of the region it represents. Computationally however, the structure generation represents a small portion in the overall computation.

In addition, because of the tunable accuracy characteristics of ACR, there are complexity-accuracy trade-offs that would naturally lead to the development of a new scalable preconditioner which we present at [44].

## Acknowledgments

We thank the anonymous reviewers for their detailed comments and suggestions for this manuscript. The authors would also like to thank Ronald Kriemann from the Max-Planck-Institute for Mathematics in the Sciences for development and continuous support of HLibPro, Alexander Litvinenko from the King Abdullah University of Science and Technology (KAUST) for the enlightening discussions and advice, and Pieter Ghysels from the Lawrence Berkeley National Laboratory for his recommendations on the use of STRUMPACK. Support from the KAUST Supercomputing Laboratory and access to Shaheen is gratefully acknowledged. The work of all authors was supported by the Extreme Computing Research Center at KAUST.

## References

- [1] R.W. Hockney, A fast direct solution of Poisson's equation using Fourier analysis, *J. ACM* 12 (1) (1965) 95–113.
- [2] B.L. Buzbee, G.H. Golub, C.W. Nielson, On direct methods for solving Poisson equation, *SIAM J. Numer. Anal.* 7 (4) (1970) pp.627–656.
- [3] W. Gander, G.H. Golub, Cyclic Reduction History and Applications, *Scientific Computing (Hong Kong, 1997)*, 1997, pp. 73–85.
- [4] P.N. Swartztrauber, The methods of cyclic reduction, Fourier analysis and the FACR algorithm for the discrete solution of Poisson equation on a rectangle, *SIAM Rev.* 19 (3) (1977) 490–501.
- [5] P. Yalamov, V. Pavlov, Stability of the block cyclic reduction, *Linear Algebra Appl.* 249 (1–3) (1996) 341–358.
- [6] S. Chandrasekaran, P. Dewilde, M. Gu, N. Somasunderam, On the numerical rank of the off-diagonal blocks of Schur complements of discretized elliptic PDEs, *SIAM J. Matrix Anal. Appl.* 31 (5) (2010) 2261–2290.
- [7] I.S. Duff, J.K. Reid, The multifrontal solution of indefinite sparse symmetric linear equations, *ACM Trans. Math. Softw.* 9 (3) (1983) 302–325.
- [8] S. Wang, X.S. Li, F.-H. Rouet, J. Xia, M.V. De Hoop, A parallel geometric multifrontal solver using hierarchically semiseparable structure, *ACM Trans. Math. Softw.* 42 (3) (2016) 21:1–21:21.
- [9] R. Vandebril, M. Barel, G. Golub, N. Mastronardi, A bibliography on semiseparable matrices, *Calcolo* 42 (3–4) (2005) 249–270.
- [10] P. Ghysels, X.S. Li, F.-H. Rouet, S. Williams, A. Napov, An efficient multicore implementation of a novel HSS-structured multifrontal solver using randomized sampling, *SIAM J. Sci. Comput.* 38 (5) (2016) S358–S384.
- [11] P.G. Martinsson, A fast randomized algorithm for computing a hierarchically semiseparable representation of a matrix, *SIAM J. Matrix Anal. Appl.* 32 (4) (2011) 1251–1274.
- [12] J. Xia, Y. Xi, M. Gu, A superfast structured solver for Toeplitz linear systems via randomized sampling, *SIAM J. Matrix Anal. Appl.* 33 (3) (2012) 837–858.
- [13] K.L. Ho, L. Ying, Hierarchical interpolative factorization for elliptic operators: differential equations, *Commun. Pure Appl. Math.* (2015).
- [14] K.L. Ho, L. Ying, Hierarchical interpolative factorization for elliptic operators: integral equations, *Commun. Pure Appl. Math.* (2015).
- [15] Y. Li, L. Ying, Distributed-memory hierarchical interpolative factorization, *arXiv preprint arXiv:1607.00346* (2016).
- [16] P.G. Martinsson, A direct solver for variable coefficient elliptic PDEs discretized via a composite spectral collocation method, *J. Comput. Phys.* 242 (2013) 460–479.
- [17] A. Gillman, P.G. Martinsson, A direct solver with  $\mathcal{O}(N)$  complexity for variable coefficient elliptic PDEs discretized via a high-order composite spectral collocation method, *SIAM J. Sci. Comput.* 36 (4) (2014) A2023–A2046.
- [18] S. Hao, P.G. Martinsson, A direct solver for elliptic PDEs in three dimensions based on hierarchical merging of Poincaré-Steklov operators, *J. Comput. Appl. Math.* 308 (2016) 419–434.

- [19] C. Weisbecker, Improving multifrontal solvers by means of algebraic Block Low-Rank representations, Institut National Polytechnique de Toulouse-INPT, 2013 Ph.D. thesis.
- [20] P.R. Amestoy, C. Ashcraft, O. Boiteau, A. Buttari, J.-Y. L'Excellent, C. Weisbecker, Improving multifrontal methods by means of block low-rank representations, *SIAM J. Sci. Comput.* 37 (3) (2015) A1451–A1474.
- [21] P.R. Amestoy, I.S. Duff, J.-Y. L'Excellent, J. Koster, A fully asynchronous multifrontal solver using distributed dynamic scheduling, *SIAM J. Matrix Anal. Appl.* 23 (1) (2001) 15–41.
- [22] P.R. Amestoy, A. Guermouche, J.-Y. L'Excellent, S. Pralet, Hybrid scheduling for the parallel solution of linear systems, *Parallel Comput.* 32 (2) (2006) 136–156. *Parallel Matrix Algorithms and Applications (PMAA04)*
- [23] M. Izadi, Hierarchical matrix techniques on massively parallel computers, Max Planck Institute for Mathematics in the Sciences, 2012 Ph.D. thesis.
- [24] S. Ambikasaran, E. Darve, An  $\mathcal{O}(N \log N)$  fast direct solver for partial Hierarchically Semiseparable matrices, *J. Sci. Comput.* 57 (3) (2013) 477–501.
- [25] I. Ibragimov, S. Rjasanow, K. Straube, Hierarchical Cholesky decomposition of sparse matrices arising from curl-curl-equation, *J. Numer. Math.* 15 (1) (2007) 31–57.
- [26] L. Grasedyck, R. Kriemann, S. Le Borne, Parallel black box  $\mathcal{H}$ -LU preconditioning for elliptic Boundary Value Problems, *Comput. Vis. Sci.* 11 (4–6) (2008) 273–291.
- [27] H. Pouransari, P. Coulier, E. Darve, Fast hierarchical solvers for sparse matrices using extended sparsification and low-rank approximation, *SIAM J. Sci. Comput.* 39 (3) (2017) A797–A830, doi:10.1137/15M1046939.
- [28] D.A. Sushnikova, I.V. Oseledets, “Compress and eliminate” solver for symmetric positive definite sparse matrices, arXiv preprint arXiv:1603.09133 (2016).
- [29] C. Chen, H. Pouransari, S. Rajamanickam, E. Boman, E. Darve, A distributed memory hierarchical solver for sparse matrices, (Personal communication) (2017).
- [30] W. Hackbusch, A sparse matrix arithmetic based on  $\mathcal{H}$ -matrices. part i: introduction to  $\mathcal{H}$ -matrices, *Computing* 62 (2) (1999) 89–108.
- [31] R. Kriemann, Parallel  $\mathcal{H}$ -matrix arithmetics on shared memory systems, *Computing* 74 (3) (2005) 273–297.
- [32] R. Kriemann,  $\mathcal{H}$ -LU factorization on many-core systems, *Comput. Visualization Sci.* 16 (3) (2013) 105–117.
- [33] W. Hackbusch, *Hierarchical Matrices: Algorithms and Analysis*, 49, Springer, 2015.
- [34] R. Guivarch, L. Giraud, J. Stein, Parallel distributed fast 3D Poisson solver for meso-scale atmospheric simulations, *Int. J. High Perform. Comput. Appl.* 15 (1) (2001) 36–46.
- [35] B. Engquist, L. Ying, Sweeping preconditioner for the Helmholtz equation: hierarchical matrix representation, *Commun. Pure Appl. Math.* 64 (5) (2011) 697–735.
- [36] L. Grasedyck, W. Hackbusch, R. Kriemann, Performance of preconditioning for sparse matrices, *Comput. Methods Appl. Math.* 8 (4) (2008) 336–349.
- [37] A. Kalinkin, A. Anders, R. Anders, et al., Schur complement computations in Intel® Math Kernel Library PARDISO, *Appl. Math.* 6 (02) (2015) 304.
- [38] W. Briggs, V. Henson, S. McCormick, *A Multigrid Tutorial*, Second Edition, 2nd, Society for Industrial and Applied Mathematics, 2000.
- [39] R.D. Falgout, U.M. Yang, *hypre: A Library of High Performance Preconditioners*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 632–641.
- [40] S. Browne, J. Dongarra, N. Garner, G. Ho, P. Mucci, A portable programming interface for performance evaluation on modern processors, *Int. J. High Perform. Comput. Appl.* 14 (3) (2000) 189–204.
- [41] W.F. Ames, *Numerical Methods for Partial Differential Equations*, Academic press, 2014.
- [42] M.M. Gupta, J. Zhang, High accuracy multigrid solution of the 3d convection–diffusion equation, *Appl. Math. Comput.* 113 (2) (2000) 249–274.
- [43] O.G. Ernst, M.J. Gander, Why it is difficult to solve Helmholtz problems with classical iterative methods, in: *Numerical Analysis of Multiscale Problems*, Springer, 2012, pp. 325–363.
- [44] G. Chávez, G. Turkiyyah, S. Zampini, D. Keyes, Parallel accelerated cyclic reduction preconditioner for three-dimensional elliptic PDEs with variable coefficients, *J. Comput. Appl. Math.* (2017), doi:10.1016/j.cam.2017.11.035.