

Fast multipole preconditioners for sparse matrices arising from elliptic equations

Huda Ibeid¹ · Rio Yokota² · Jennifer Pestana³ · David Keyes⁴

Received: 12 January 2016 / Accepted: 22 August 2017 / Published online: 9 November 2017
© The Author(s) 2017. This article is an open access publication

Abstract Among optimal hierarchical algorithms for the computational solution of elliptic problems, the fast multipole method (FMM) stands out for its adaptability to emerging architectures, having high arithmetic intensity, tunable accuracy, and relaxable global synchronization requirements. We demonstrate that, beyond its traditional use as a solver in problems for which explicit free-space kernel representations are available, the FMM has applicability as a preconditioner in finite domain elliptic boundary value problems, by equipping it with boundary integral capability for satisfying conditions at finite boundaries and by wrapping it in a Krylov method for extensibility to more general operators. Here, we do not discuss the well developed applications of FMM to implement matrix-vector multiplications within Krylov solvers of boundary element methods. Instead, we propose using FMM for the volume-to-volume contribution of inhomogeneous Poisson-like problems, where the boundary integral is a small part of the overall computation. Our

method may be used to precondition sparse matrices arising from finite difference/element discretizations, and can handle a broader range of scientific applications. It is capable of algebraic convergence rates down to the truncation error of the discretized PDE comparable to those of multigrid methods, and it offers potentially superior multicore and distributed memory scalability properties on commodity architecture supercomputers. Compared with other methods exploiting the low-rank character of off-diagonal blocks of the dense resolvent operator, FMM-preconditioned Krylov iteration may reduce the amount of communication because it is matrix-free and exploits the tree structure of FMM. We describe our tests in reproducible detail with freely available codes and outline directions for further extensibility.

Keywords Fast multipole method · Preconditioner · Krylov subspace method · Poisson equation · Stokes equation

Communicated by Gabriel Wittum.

✉ Huda Ibeid
huda.ibeid@kaust.edu.sa
Rio Yokota
rioyokota@gsic.titech.ac.jp
Jennifer Pestana
jennifer.pestana@strath.ac.uk
David Keyes
david.keyes@kaust.edu.sa

- ¹ Extreme Computing Research Center, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia
- ² Tokyo Institute of Technology, Tokyo, Japan
- ³ University of Strathclyde, Glasgow, UK
- ⁴ King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

1 Introduction

Elliptic PDEs arise in a vast number of applications in scientific computing. A significant class of these involve the Laplace operator, which appears not only in potential calculations but also in, for example, Stokes and Navier-Stokes problems [25, Chapters 5 and 7], electron density computations [53, Part II] and reaction-convection-diffusion equations [43, Part IV]. Consequently, the rapid solution of PDEs involving the Laplace operator is of wide interest.

Although many successful numerical methods for such PDEs exist, changing computer architectures necessitate new paradigms for computing and the development of new algorithms. Computer architectures of the future will favor algorithms with high concurrency, high data locality, high arithmetic intensity (Flops/Byte), and low synchronicity.

This trend is manifested on GPUs and co-processors, where some algorithms are accelerated much less than others on the class of architectures that can be extended to extreme scale. There is always a balance between algorithmic efficiency in a convergence sense, and how well an algorithm scales on parallel architectures. This balance is shifting towards increased parallelism, even at the cost of increasing computation. Since the processor frequency has plateaued for the last decade, Moore's law holds continued promise only for those who are willing to make algorithmic changes.

Among the scientific applications ripe for reconsideration, those governed by elliptic PDEs will be among the most challenging. A common solution strategy for such systems is to discretize the partial differential equations by fairly low-order finite element, finite volume or finite difference methods and then solve the resulting large, sparse linear system. However, elliptic systems are global in nature, and this does not align well with the sweet spots of future architectures. The linear solver must enable the transfer of information from one end of the domain to the other, either through successive local communications (as in many iterative methods), or a direct global communication (as in direct solvers with global recurrences and Krylov methods with global reductions). In either case, avoiding synchronization and reducing communication are the main challenges. There has been considerable effort in this direction in the dense linear algebra community [22]. The directed-acyclic-graph-based technology developed in such efforts could be combined with iterative algorithms of optimal complexity for solving elliptic PDEs at extreme scale.

Scalable algorithms for solving elliptic PDEs tend to have a hierarchical structure, as in multigrid methods [65], fast multipole methods (FMM) [35], and \mathcal{H} -matrices [38]. This structure is crucial, not only for achieving optimal arithmetic complexity, but also for minimizing data movement. For example, the standard 3-D FFT with three all-to-all communications requires $\mathcal{O}(\sqrt{P})$ communication for the transpose between pencil-shaped subdomains on P processes [20] and a recently published algorithm [46] with five all-to-all communication phases achieves $\mathcal{O}(P^{1/3})$ communication, whereas these hierarchical methods require $\mathcal{O}(\log P)$ communication [50]. This $\mathcal{O}(\log P)$ communication complexity is likely to be optimal for elliptic problems, since an appropriately coarsened representation of a local forcing must somehow arrive at all other parts of the domain for the elliptic equation to converge. In other words, an elliptic problem for which the solution is desired everywhere cannot have a communication complexity of $\mathcal{O}(1)$. However, the convergence of these hierarchical solvers can be fragile with respect to coefficient distribution in the second-order term, and, if present, with respect to the first-order and zeroth-order terms.

Krylov subspace methods provide another popular alternative to direct methods for general operators. We note that methods such as Chebyshev semi-iteration can require even less communication in the fortunate case when information about the spectrum of the coefficient matrix is known [28, Section 10.1.5], [29]. Among the best known Krylov methods are the conjugate gradient method [41], MINRES [56] and GMRES [60], although a multitude of Krylov solvers are available in popular scalable solver libraries. The great advantage of these solvers is their robustness—for any consistent linear system there exists a Krylov method that will converge, in exact arithmetic, for sufficiently many iterations. However, the convergence rate of unpreconditioned Krylov methods deteriorates as the discretization of an elliptic PDE is refined.

Mesh-independent convergence for Krylov methods applied to systems from elliptic PDEs can be obtained by sufficiently strong preconditioning. Among the best performing preconditioners are the optimal hierarchical methods or, for multiphysics problems such as Stokes and Navier–Stokes equations, block preconditioners with these methods as components. By combining these hierarchical methods and Krylov subspace solvers we get the benefits of both approaches and obtain a linear solver that is robust and fast. These hierarchical methods have multiple parameters for controlling the precision of the solution and are able to trade-off accuracy for speed, which is a useful feature for a preconditioner. Furthermore, in analogy to the pair of multigrid approaches denoted geometric and algebraic, \mathcal{H}^2 -matrices [39] can be thought of as an algebraic generalization of what FMMs do geometrically. There are advantages and disadvantages to using algebraic and geometric methods, and both have a niche as preconditioners.

There has been recent work on algebraic multigrid methods (AMG) in anticipation of the coming hardware constraints mentioned above. Gahvari et al. [26] developed a performance model for AMG and tested it on various HPC systems—Intrepid, Jaguar, Hera, Zeus, and Atlas. They showed that network distance and contention were both substantial performance bottlenecks for AMG. Adams presents a low-memory matrix-free full multigrid (FMG) with a full approximation storage (FAS) [2]. He revives an idea from the 1970s [12], which processes the multigrid algorithm vertically, and improves data locality and asynchronicity. Baker et al. [4] compared the scalability of different smoothers—hybrid Gauss–Seidel, l_1 Gauss–Seidel, and Chebyshev polynomial, and showed that l_1 Gauss–Seidel and Chebyshev smoothers scale much better. Vassilevski and Yang [66] present additive variants of AMG that are significantly improved with respect to classical additive methods and show their scalable performance on up to 4096 cores. Indeed, there is continuous progress to evolve multigrid to

future hardware constraints, and it is likely that multigrid will remain competitive.

Complementing this evolution, hierarchical low-rank approximation (HLRA) of the off-diagonal blocks of a matrix leads to a whole new variety of $\mathcal{O}(N)$ solvers or preconditioners. HLRA based methods include FMM itself [35], \mathcal{H} -matrices [38], hierarchically semi-separable matrices [16], hierarchically off-diagonal low-rank technique [3], and recursive skeletonization [42], in an increasingly diverse pool. These techniques can be applied to a dense matrix or a Schur complement during a sparse direct solve, thus enabling an $\mathcal{O}(N)$ matrix-vector multiplication of a $N \times N$ dense matrix or an $\mathcal{O}(N)$ direct solve of a $N \times N$ sparse matrix to within a specified accuracy. These HLRA based methods exploit a cheaply approximable kernel in the far field, which yields a block low-rank structure. The distinguishing features of the variants come in the way the low-rank approximation is constructed—rank-revealing LU [57], rank-revealing QR [36], pivoted QR [48], truncated SVD [31], randomized SVD [51], adaptive cross approximation [58], hybrid cross approximation [10], and Chebychev interpolation [23] are all possibilities. Multipole/local expansions in the FMM constitute another way to construct the low-rank approximations. Many of the original developers of FMM are now working on these algebraic variants [34]. There are also several groups actively contributing to the field of the FMM algorithm, and its high-performance implementation to enable the algorithm migration to the exascale systems [8]. Furthermore, several performance models for the FMM have been developed to anticipate the challenges for FMM on future exascale architectures [45, 50, 64].

Literature on the HLRA-based methods mentioned above mainly focuses on the error convergence of the low-rank approximation and there is less investigation of their parallel scalability or of a direct comparison to multigrid. An exception is the work by Grasedyck et al. [32], where their \mathcal{H} -LU preconditioner is compared with BoomerAMG, Pardiso, MUMPS, UMFPACK, SuperLU, and Spooles. However, their executions are serial, and show that their \mathcal{H} -matrix code is not yet competitive with these other highly optimized libraries. Another is the work by Gholami et al. [27] where they compare FFT, FMM, and multigrid methods for the Poisson problem with constant coefficients on the unit cube with periodic boundary conditions. FMM has also been used as a continuum volume integral with adaptive refinement capabilities [52]. This approach defines the discretization adaptively inside the FMM, whereas in the present method a user defines the discretization and the preconditioner is provided.

In the present work, we consider the Laplace and Stokes boundary value problems and devise highly scalable preconditioners for these problems. Our Poisson preconditioner relies on a boundary element method in which matrix-vector multiplies are performed using FMM; the result is an $\mathcal{O}(N)$

preconditioner that is scalable, where N is the total degrees of freedom, not just those on the boundary. For the Stokes problem, we apply a block diagonal preconditioner, in which our Poisson preconditioner is combined with a simple diagonal matrix. FMM based preconditioners were first proposed by Sambavaram et al. [61]. Such methods lacked practical motivation when flops were expensive, since they turn a sparse matrix into a dense matrix of the same size before hierarchically grouping the off-diagonal blocks. But in a world of inexpensive flops relative to data motion, the notion of a “compute-bound preconditioner” is attractive. In the present work, we perform scalability benchmarks and compare the time-to-solution with state-of-the-art multigrid methods such as BoomerAMG in a high performance computing environment.

The rest of the manuscript is organized as follows. In Sect. 2 we present the model problems and in Sect. 3 we give an overview of Krylov subspace methods and preconditioning. The basis of our preconditioner is a boundary element method that is discussed in Sect. 4 and the FMM, the kernel essential for efficiency and scalability, is described in Sect. 5. Numerical results in Sect. 6 examine the convergence rates of FMM and multigrid for small Poisson and Stokes problems. In Sect. 7 we scale up the Poisson problem tests and perform strong scalability runs, where we compare the time-to-solution against BoomerAMG [40] on up to 1024 cores. Conclusions are drawn in Sect. 8.

2 Model problems

In this section we introduce the model Poisson and Stokes problems and describe properties of the linear systems that result from their discretization. We focus on low-order finite elements but note that discretization by low-order finite difference or finite volume methods give linear systems with similar properties.

2.1 Poisson model problem

The model Poisson problems we wish to solve are of the form

$$-\nabla \cdot (a \nabla u) = f \text{ in } \Omega, \quad (1a)$$

$$u = g \text{ on } \Gamma, \quad (1b)$$

where $\Omega \in \mathbb{R}^d$, $d = 2, 3$ is a bounded connected domain with piecewise smooth boundary Γ , f is a forcing term, g defines the Dirichlet boundary condition, and $a \geq a_0 > 0$ is a sufficiently smooth function of space.

Discretization of (1) by finite elements or finite differences leads to a large, sparse linear system of the form

$$Ax = b, \tag{2}$$

where $A \in \mathbb{R}^{N \times N}$ is the stiffness matrix and $b \in \mathbb{R}^N$ contains the forcing and boundary data. The matrix A is symmetric positive definite and its eigenvalues depend on the mesh size, which we denote by h , as is typical of discretizations of elliptic PDEs. In particular, the condition number $\kappa = \lambda_{\max}(A)/\lambda_{\min}(A)$, the ratio of the largest and smallest eigenvalues of A , grows as $O(h^{-2})$ (see, for example, [25, Section 1.6]).

2.2 Stokes model problem

Incompressible Stokes problems are important in the modeling of viscous flows and for solving Navier–Stokes equations by operator splitting methods [9, Section 2.1]. The equations governing the velocity $u \in \mathbb{R}^d$, $d = 2, 3$, and pressure $p \in \mathbb{R}$ of a Stokes fluid in a bounded connected domain Ω with piecewise smooth boundary Γ are [9,25]:

$$-\nabla^2 u + \nabla p = 0 \text{ in } \Omega, \tag{3a}$$

$$\nabla \cdot u = 0 \text{ in } \Omega, \tag{3b}$$

$$u = w \text{ on } \Gamma. \tag{3c}$$

Discretizing (3) by a stabilized¹ finite element or finite difference approximation leads to the symmetric saddle point system

$$\underbrace{\begin{bmatrix} A & B^T \\ B & -C \end{bmatrix}}_A \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} f \\ g \end{bmatrix}, \tag{4}$$

where $A \in \mathbb{R}^{N \times N}$ is the vector-Laplacian, a block diagonal matrix with blocks equal to the stiffness matrix from (2), $B \in \mathbb{R}^{M \times N}$ is the discrete divergence matrix, $C \in \mathbb{R}^{M \times M}$ is the symmetric positive definite pressure mass matrix and $f \in \mathbb{R}^N$ and $g \in \mathbb{R}^M$ contain the Dirichlet boundary data.

The matrix \mathcal{A} is symmetric indefinite and the presence of the stiffness matrix means that the condition number of \mathcal{A} increases as the mesh is refined. However, as we will see in the next section, the key ingredient in a preconditioner for \mathcal{A} that mitigates this mesh dependence is a good preconditioner for the Poisson problem. This allows us to use the preconditioner for the Poisson problem in the Stokes problem as well.

¹ Although we treat only stabilized discretizations here, stable discretizations are no more difficult to precondition and are discussed in detail in Elman et al. [25, Chapter 6].

3 Iterative solvers and preconditioning

3.1 Krylov subspace methods

Large, sparse systems of the form (2) are often solved by Krylov subspace methods. We focus here on two Krylov methods: the conjugate gradient method (CG) [41] for systems with symmetric positive definite coefficient matrices and MINRES [56] for systems with symmetric indefinite matrices. For implementation and convergence details, we refer the reader to the books by Greenbaum [33] and Saad [59].

The convergence of these Krylov subspace methods depends on the spectrum of the coefficient matrix which for the Poisson and Stokes problems, as well as other elliptic PDEs, deteriorates as the mesh is refined. This dependence can be removed by preconditioning. In the case of the Poisson problem (2), we can conceptually think of solving the equivalent linear system $P^{-1}Ax = P^{-1}b$ (left preconditioning), or $AP^{-1}y = b$, with $P^{-1}y = x$ (right preconditioning), for some $P^{-1} \in \mathbb{R}^{N \times N}$, and analogously for the Stokes equations (4). However, when the coefficient matrix is symmetric, we would like to preserve this property when preconditioning; this can be achieved by using a symmetric positive definite preconditioner [25, Chapters 2 and 6]. We also note that in practice we never need P^{-1} explicitly but only the action of this matrix on a vector. This enables us to use matrix-free approaches such as multigrid or the fast multiple method.

Many preconditioners for the Poisson problem reduce the number of iterations, with geometric and algebraic multigrid among the most effective strategies [25,65]. However, to achieve a lower time-to-solution than can be obtained for the original system, it is also necessary to choose a preconditioner that can be cheaply applied at each iteration. Both geometric and algebraic multigrid methods are $O(N)$, and therefore exhibit good performance on machines and problems for which computation is expensive. However, stresses arise in parallel applications as discussed in the introduction.

For Stokes problems we consider the block diagonal preconditioner

$$\mathcal{P} = \begin{bmatrix} P_A & 0 \\ 0 & P_S \end{bmatrix}, \tag{5}$$

where $P_A \in \mathbb{R}^{N \times N}$ and $P_S \in \mathbb{R}^{M \times M}$ are symmetric positive definite matrices. The advantage of this preconditioner is that there is no coupling between the blocks, so \mathcal{P} is scalable provided the blocks P_A and P_S are.

Appropriate choices for P_A and P_S have been well studied and it is known that mesh-independent convergence of MINRES can be recovered when P_A is spectrally equivalent to A in (4) and P_S is spectrally equivalent to the pressure

mass matrix $Q \in R^{M \times M}$ [14], [25, Chapter 6]. These spectral equivalence requirements imply that the eigenvalues of $P_A^{-1}A$ and $P_S^{-1}Q$ are bounded in an interval on the positive real line independently of the mesh width h .

It typically suffices to use the diagonal of Q [25, Chapter 6], [71] or a few steps of Chebyshev semi-iteration [70] for P_S . Moreover, the diagonal matrix is extremely parallelizable. Thus, the key to obtaining a good preconditioner for A is to approximate the vector Laplacian effectively. This is typically the most computationally intensive part of the preconditioning process, since in most cases $M \ll N$.

3.2 The FMM–BEM preconditioner

In this paper we propose a preconditioner for Poisson and Stokes problems that heavily utilizes the fast multipole method (FMM). The FMM is $\mathcal{O}(N)$ with compute intensive inner kernels. It has a hierarchical data structure that allows asynchronous communication and execution. These features make the FMM a promising preconditioner for large scale problems on future computer architectures. We show that this preconditioner improves the convergence of Krylov subspace methods applied to these problems and is effectively parallelized on today’s highly distributed architectures.

The FMM in its original form relies on free-space Green’s functions and is able to solve problems with free-field boundary conditions. In Sect. 4 the FMM preconditioner is extended to Dirichlet, Neumann or Robin boundary conditions for arbitrary geometries by coupling it with a boundary element method (BEM). Our approach uses the FMM as a *preconditioner* inside a *sparse* matrix solver and the BEM solve is *inside* the preconditioner. Numerous previous studies use FMM for the matrix-vector multiplication inside the Krylov solver for the dense matrix arising from the boundary element discretization. In the present method we are calculating problems with non-zero sources in the volume, and the FMM is used to calculate the volume-to-volume contribution. This means we are performing the action of an $N \times N$ dense matrix-vector multiplication, where N is the number of points in the volume (not the boundary). Additionally, as discussed in Sect. 4.4, it is possible to extend the boundary element method to problems with variable diffusion coefficients, particularly since low accuracy solves are often sufficient in preconditioning.

Figure 1 shows the flow of calculation of our FMM-BEM preconditioner within the conjugate gradient method; its role in other Krylov solvers is similar. The FMM is used to approximate the matrix-vector multiplication of A^{-1} within the preconditioner. The BEM solver adapts the FMM to finitely applied boundary conditions. During each step of the iteration, the u vector from the previous iteration is used to determine $\partial u / \partial n$ at the boundary from (8), then (9) is used to compute the new u in the domain Ω .

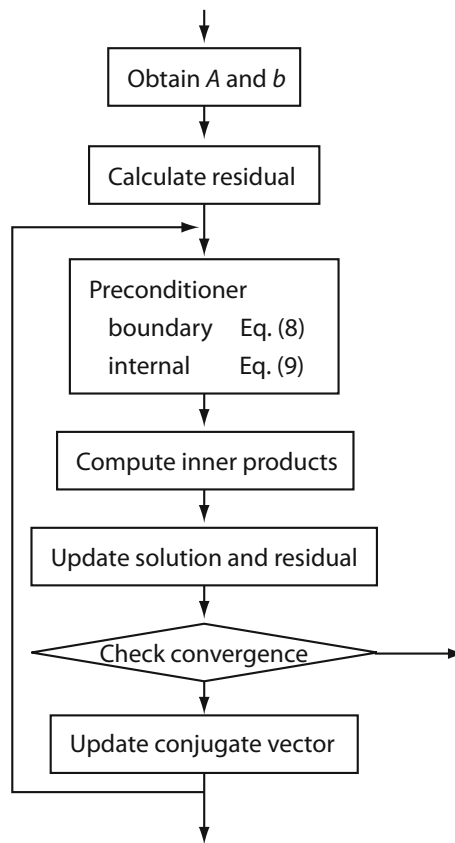


Fig. 1 Flow chart of the FMM-BEM preconditioner within the conjugate gradient method

4 Boundary element method

4.1 Formulation

We use a standard Galerkin boundary element method [62] with volume contributions to solve the Poisson equation. A brief description of the formulation is given here. Applying Green’s third identity to (1a) with $a \equiv 1$ gives

$$\int_{\Gamma} u \frac{\partial G}{\partial n} d\Gamma - \int_{\Gamma} \frac{\partial u}{\partial n} G d\Gamma - \int_{\Omega} u (\nabla^2 G) d\Omega = \int_{\Omega} f G d\Omega, \tag{6}$$

where G is the Green’s function of the Laplace operator, $\frac{\partial}{\partial n}$ is the derivative in the outward normal direction, and Γ is the boundary. Following the definition of the Green’s function $\nabla^2 G = -\delta$, the third term in (6) becomes

$$-\int_{\Omega} u (\nabla^2 G) d\Omega = \int_{\Omega} u \delta d\Omega = \begin{cases} \frac{1}{2}u & \text{on } \partial\Omega, \\ u & \text{in } \Omega. \end{cases} \tag{7}$$

Therefore, we may solve the constant coefficient inhomogeneous Poisson problem by solving the following set of equations

$$\int_{\Gamma} \frac{\partial u}{\partial n} G d\Gamma = \int_{\Gamma} u \left(\frac{1}{2} \delta + \frac{\partial G}{\partial n} \right) d\Gamma - \int_{\Omega} f G d\Omega \quad \text{on } \partial\Omega, \tag{8}$$

$$u = \int_{\Gamma} \frac{\partial u}{\partial n} G d\Gamma - \int_{\Gamma} u \frac{\partial G}{\partial n} d\Gamma + \int_{\Omega} f G d\Omega \quad \text{in } \Omega. \tag{9}$$

As an example, consider the case where Dirichlet boundary conditions are prescribed on $\partial\Omega$. The unknowns are $\partial u/\partial n$ on Γ and u in $\Omega \setminus \Gamma$, where (8) solves for the former and (9) can be used to determine the latter. For Neumann boundary conditions one can simply switch the two boundary integral terms in (8) and solve for u instead of $\partial u/\partial n$. In either case, we obtain both u and $\partial u/\partial n$ at each point on the boundary, then calculate (9) to obtain u at the internal points.

4.2 Singular integrals

The Laplace Green’s function in 2-D

$$G = -\frac{1}{2\pi} \log r \tag{10}$$

is singular. Therefore, the integrals involving G or $\partial G/\partial n$ in (8) and (9) are singular integrals. As described in the following subsection, these singular integral are discretized into piecewise integrals, which are evaluated using Gauss-Legendre quadratures with special treatment for the singular piecewise integral. For boundary integrals in (8) and (9), analytical expressions exist for the piecewise integral. However, for the volume integral an analytical expression does not exist [1]. For this reason, we used a smoothed Green’s function of the form

$$G = -\frac{1}{2\pi} \log \left(\sqrt{r^2 + \epsilon^2} \right) \tag{11}$$

where ϵ is a small number that changes with the grid resolution. An alternative approach is to instead approximate these singular integrals by hierarchical quadrature [11]. This approach expresses singular integrals in terms of regular ones by splitting the domain of integration into a hierarchy of suitable subdomains where standard quadrature can be performed.

4.3 Discretization

The integrals in Eqs. (8) and (9) are discretized in a similar fashion to finite element methods. In the following description of the discretization process, we will use the term on the left hand side in (8) as an example. The first step is to

break the global integral into a discrete sum of piecewise local integrals over each element

$$\int_{\Gamma} \frac{\partial u}{\partial n} G d\Gamma \approx \sum_{j=1}^{N_{\Gamma}} \int_{\Gamma_j} \frac{\partial u_j}{\partial n} G d\Gamma_j, \tag{12}$$

where N_{Γ} is the number of boundary nodes. These piecewise integrals are performed by using quadratures over the basis functions [62]. In the present case, we use constant elements so there are no nodal points at the corners of the square domain for the tests in Sects. 6 and 7. By applying this discretization technique to all terms in (8) we obtain

$$N_{\Gamma} \left\{ \begin{matrix} \overbrace{\left[\begin{matrix} \ddots & & \\ & G_{ij} & \\ & & \ddots \end{matrix} \right]}^{N_{\Gamma}} \left[\begin{matrix} \vdots \\ \frac{\partial u_j}{\partial n} \\ \vdots \end{matrix} \right] \\ \text{unknown} \end{matrix} \right\} = \left[\begin{matrix} \ddots & & \\ & \frac{1}{2} \delta_{ij} + \frac{\partial G_{ij}}{\partial n} & \\ & & \ddots \end{matrix} \right] \left[\begin{matrix} \vdots \\ u_j \\ \vdots \end{matrix} \right] - \left[\begin{matrix} \ddots & & \\ & G_{ij} & \\ & & \ddots \end{matrix} \right] \left[\begin{matrix} \vdots \\ f_j \\ \vdots \end{matrix} \right],$$

where N_{Ω} is the number of internal nodes. All values on the right hand side are known, and $\partial u/\partial n$ at the boundary is determined by solving the linear system. Similarly, we apply the discretization to (9) to have

$$N_{\Omega} \left\{ \begin{matrix} \left[\begin{matrix} \vdots \\ u_i \\ \vdots \end{matrix} \right] = \left[\begin{matrix} \ddots & & \\ & \frac{\partial G_{ij}}{\partial n} & \\ & & \ddots \end{matrix} \right] \left[\begin{matrix} \vdots \\ u_j \\ \vdots \end{matrix} \right] \\ \text{unknown} \end{matrix} \right\} - \left[\begin{matrix} \ddots & & \\ & G_{ij} & \\ & & \ddots \end{matrix} \right] \left[\begin{matrix} \vdots \\ \frac{\partial u_j}{\partial n} \\ \vdots \end{matrix} \right] + \left[\begin{matrix} \ddots & & \\ & G_{ij} & \\ & & \ddots \end{matrix} \right] \left[\begin{matrix} \vdots \\ f_j \\ \vdots \end{matrix} \right].$$

At this point, all values on the right hand side are known so one can perform three matrix-vector multiplications to obtain u at the internal nodes, and the solution to the original Poisson equation (1a). The third term on the right hand side involves an $N_{\Omega} \times N_{\Omega}$ matrix, and is the dominant part of the computational load. This matrix-vector multiplication can be approximated in $\mathcal{O}(N)$ time by using the FMM described in Sect. 5. We also use the FMM to accelerate all other matrix-vector multiplications.

4.4 Variable coefficient problems

A natural question that arises is how to extend the boundary element method, which is the basis of our preconditioner, to problems (1) with variable diffusion coefficients.

Several strategies for extending boundary element methods to problems with variable diffusion coefficients have been proposed (see, for example, the thesis of Brunton [13, Chapter 3]). Additionally, in this preconditioner setting we may not need to capture the variation in the diffusion coefficient to a high degree of accuracy; for a similar discussion in the context of additive Schwarz preconditioners see, for example, Graham et al. [30].

Although analytic fundamental solutions can sometimes be found for problems with variable diffusion (see, e.g., Cheng [17] and Clements [18]), in most cases numerical techniques are employed. One popular method is to introduce a number of subdomains, on each of which the diffusion coefficient is approximated by a constant function [49, 67].

A second option is to split the differential operator into a part for which a fundamental solution exists and another which becomes part of the source term. Specifically, starting from (1), a similar approach to that described in Banerjee [7] and Cheng [17] leads to

$$\int_{\Gamma} au \frac{\partial G}{\partial n} d\Gamma - \int_{\Gamma} a \frac{\partial u}{\partial n} G - \int_{\Omega} u \nabla a \cdot \nabla G d\Omega - \int_{\Omega} au \nabla^2 G d\Omega = \int_{\Omega} f G d\Omega,$$

where again G is the standard fundamental solution for the Laplace operator, i.e., not the fundamental solution for (1). We can then proceed as described above for (6). It is also possible (see Concus and Golub [19]) to change the dependent variable to soak up the variation in a prior to discretization, again resulting in a modified source FEM.

5 Fast multipole method

5.1 Introduction to FMM

The last term in Eq. (9) when discretized has the form

$$u_i = \sum_{j=1}^{N_{\Omega}} f_j G_{ij}, \tag{13}$$

where $i = 1, 2, \dots, N_{\Omega}$. If we calculate this equation directly, it will require $\mathcal{O}(N^2)$ operations. In Fig. 2, we show by schematic how the fast multipole method is able to calculate this in $\mathcal{O}(N)$ operations. Figure 2a, b show how the source particles (red) interact with the target particles (blue)

for the direct method and FMM, respectively. In the direct method, all source particles interact with all target particles directly. In the FMM, the source particles are first converted to multipole expansions using the P2M (particle to multipole) kernel. Figure 2c shows the corresponding geometric view of the hierarchical domain decomposition of the particle distribution. Then, multipole expansions are aggregated into larger groups using the M2M (multipole to multipole) kernel. Following this, the multipole expansions are translated to local expansions between well-separated cells using the M2L (multipole to local) kernel. Both Fig. 2b, c show that the larger cells interact if they are significantly far away, and smaller cells may interact with slightly closer cells. The direct neighbors between the smallest cells are calculated using the P2P (particle to particle) kernel, which is equivalent to the direct method between a selected group of particles. Then, the local expansions of the larger cells are translated to smaller cells using the L2L (local to local) kernel. Finally, the local expansions at the smallest cells are translated into the potential on each particle using the L2P (local to particle) kernel. The mathematical formulas for these kernels will be given in Sect. 5.3. Note, for simplification purposes, that each scale of the hierarchical summation can be translated asynchronously from source to target.

In order to perform the FMM calculation mentioned above, one must first decompose the domain in a hierarchical manner. It is common to use an octree in 3-D and quad-tree in 2-D, where the domain is split by a geometrical center plane. The splitting is performed recursively until the number of particles per cell reaches a prescribed threshold. The splitting is usually performed adaptively, so that the densely populated areas result in a deeper branching of the tree. A common requirement in FMMs is that these cells must be isotropic (cubes or squares and not rectangles), since they are used as units for measuring the well-separateness as shown in Fig. 2c during the M2L interaction. However, our FMM does not use the size of cells to measure the distance between them and allows the cells to be of any shape as long as they can be hierarchically grouped into a tree structure. Once the tree structure is constructed, it is trivial to find parent-child relationships between the cells/particles. This relation is all that is necessary for performing P2M, M2M, L2L, and L2P kernels. However, for the M2L and P2P kernels one must identify a group of well-separated and neighboring cells, respectively. We will describe an efficient method for finding well-separated cells in the following subsection.

5.2 Dual tree traversal

The simplest method for finding well-separated pairs of cells in the FMM is to “loop over all target cells and find their parent’s neighbor’s children that are non-neighbors,” as shown by Greengard and Rokhlin [35]. A scheme that permits the

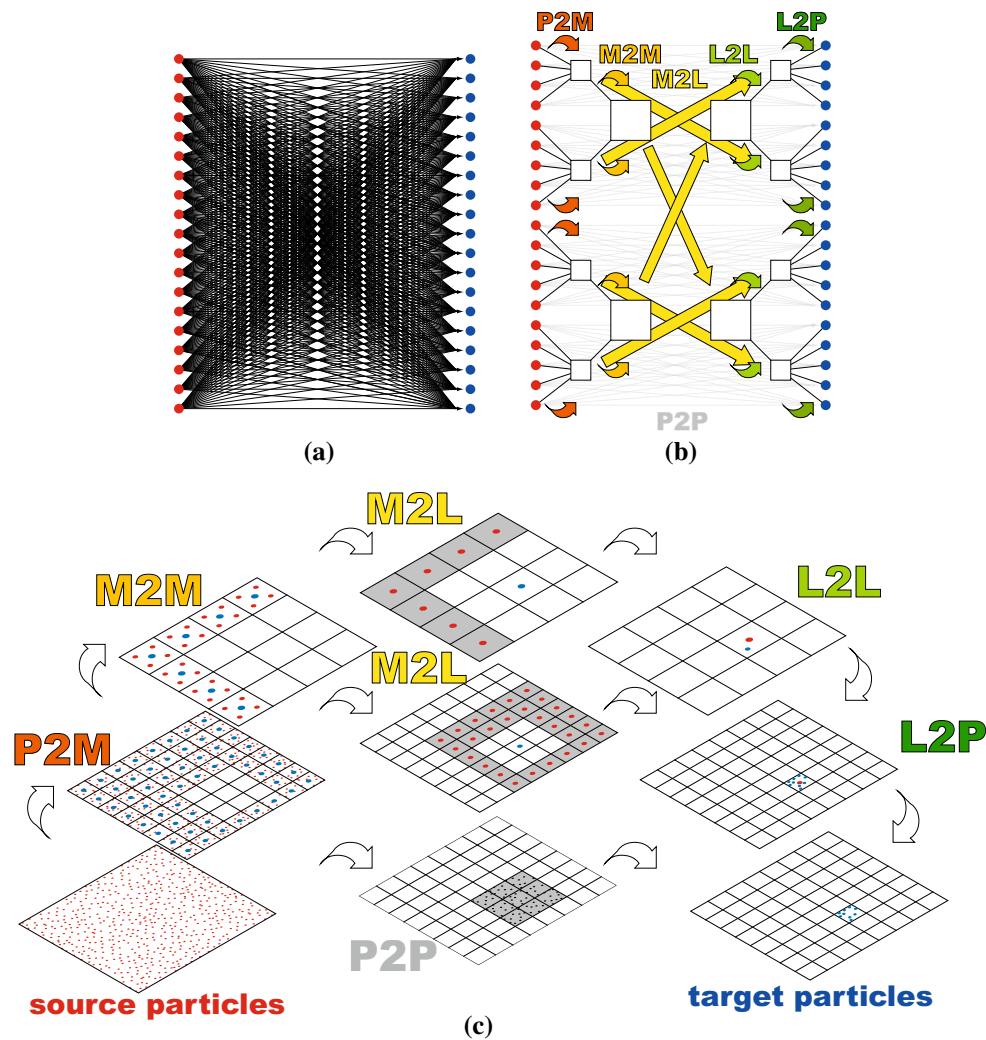


Fig. 2 Schematic of fast multipole method. **a** shows the interactions for a $\mathcal{O}(N^2)$ direct method. **b** shows the interactions for the $\mathcal{O}(N)$ FMM, describing the type of interaction between elements in the tree

data structure. **c** shows the same FMM kernels as in (b), but from a geometric point of view of the hierarchical domain decomposition

interaction of cells at different levels for an adaptive tree was introduced by Carrier et al. [15]. This scheme is used in many modern FMM codes, and is sometimes called the UVWX-list [50]. Another scheme to find well-separated pair of cells is to “simultaneously traverse the target and source tree while applying a multipole acceptance criterion,” as shown by Warren and Salmon [69]. Teng [64] showed that this dual tree traversal can produce interaction pairs that are almost identical to the adaptive interaction list by Carrier et al. [15]. A concise explanation and optimized implementation of the dual tree traversal is provided by Dehnen [21].

The dual tree traversal has many favorable properties compared to the explicit construction of interaction lists. First of all, the definition of well-separatedness can be defined quite flexibly. For example, if one were to construct explicit interaction lists by extending the definition of neighbors from

$3 \times 3 \times 3$ to $5 \times 5 \times 5$ using the traditional scheme, the M2L list size will increase rapidly from $6^3 - 3^3 = 189$ to $10^3 - 5^3 = 875$ in 3-D, which is never faster for any number of expansions. On the other hand, the dual tree traversal can adjust the definition of neighbors much more flexibly and the equivalent interaction list always has a spherical shape. (We say “equivalent interaction list” because there is no explicit interaction list construction in the dual tree traversal.) The cells no longer need to be cubic, since the cells themselves are not used to measure the proximity of cells. The cells can be any shape or size—even something derived from hierarchical K-means clustering. The explicit interaction list construction can be modified to include more flexibility, as well [37]. However, the resulting code becomes much more complicated than the dual tree traversal, which is literally a few lines of

code.² This simplicity is a large advantage on its own. Furthermore, the parallel version of the dual tree traversal simply traverses the local tree for the target with the local essential tree [68] for the sources, so the serial dual tree traversal code can be used once the local essential tree is assembled.

A possible (but unlikely) limitation of dual tree traversals is the loss of explicit parallelism—it has no loops. It would not be possible to simply use an OpenMP “parallel for” directive to parallelize the dual tree traversal. In contrast, the traditional schemes always have an outer loop over the target cells, which can be easily parallelized and dynamically load balanced with OpenMP directives. However, this is not an issue since task-based parallelization tools such as Intel Thread Building Blocks (TBB) can be used to parallelize the dual tree traversal. With the help of these tools, tasks are spawned as the tree is traversed and dispatched to idle threads dynamically. In doing so, we not only assure load-balance but also data-locality, so it may actually end up being a superior solution than parallelizing “for loops” with OpenMP, especially on NUMA architectures.

Considering the advantages mentioned above, we have decided to use the dual tree traversal in our current work. This allows us to perform low accuracy optimizations by adjusting the multipole acceptance criterion without increasing the order of expansions too much, which is the secret to our speed [72]. These low accuracy optimizations can give the FMM a performance boost when used as a preconditioner.

5.3 Multipole expansions

For the 2-D Laplace equation, the free space Green’s function, as noted in (10), has the form

$$G_{ij} = \frac{1}{2\pi} \log\left(\frac{1}{r_{ij}}\right), \tag{14}$$

where $r_{ij} = |\mathbf{x}_i - \mathbf{x}_j|$ is the distance between point i and point j . By using complex numbers to represent the two-dimensional coordinates $z = x + iy$, Eq. (13) can be written as

$$u_i = \sum_{j=1}^{N_\Omega} \frac{f_j}{2\pi} \Re\{-\log(z_{ij})\}, \tag{15}$$

where $\Re(z)$ represents the real part of z . Figure 3 shows the decomposition of vector \mathbf{x}_{ij} into five parts, $\mathbf{x}_{ij} = \mathbf{x}_{i\lambda} + \mathbf{x}_{\lambda\Lambda} + \mathbf{x}_{\Lambda M} + \mathbf{x}_{M\mu} + \mathbf{x}_{\mu j}$, where λ and Λ are the center of local expansions and μ and M are the center of multipole expansions. The lower case is used for the smaller cells and upper case is used for the larger cells. When assuming

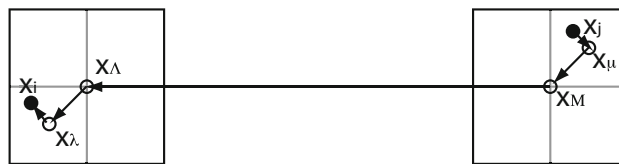


Fig. 3 Decomposition of the distance vector $\mathbf{x}_{ij} = \mathbf{x}_i - \mathbf{x}_j$ into five parts, that correspond to the five stages P2M, M2M, M2L, L2L, and L2P in the FMM

the relation $|\mathbf{x}_{\Lambda M}| > |\mathbf{x}_{i\lambda} + \mathbf{x}_{\lambda\Lambda}| + |\mathbf{x}_{M\mu} + \mathbf{x}_{\mu j}|$ the following FMM approximations are valid [15]. We denote the n th order multipole expansion coefficient at \mathbf{x} as $M_n(\mathbf{x})$, and the n th order local expansion coefficient as $L_n(\mathbf{x})$, where $n = 0, 1, \dots, p - 1$ for a p th order truncation of the series.

1. P2M from particle at \mathbf{x}_j to multipole expansion at \mathbf{x}_μ ,

$$M_0(\mathbf{x}_\mu) = \sum_{j=1}^N f_j, \tag{16}$$

$$M_n(\mathbf{x}_\mu) = \sum_{j=1}^N \frac{-f_j (-z_{\mu j})^n}{n}, n = \{1, 2, \dots, p - 1\}. \tag{17}$$

2. M2M from multipole expansion at \mathbf{x}_μ to multipole expansion at \mathbf{x}_M ,

$$M_0(\mathbf{x}_M) = M_0(\mathbf{x}_\mu), \tag{18}$$

$$M_n(\mathbf{x}_M) = -M_0(\mathbf{x}_\mu) \frac{(-z_{M\mu})^n}{n} + \sum_{k=1}^n M_k(\mathbf{x}_\mu) (-z_{M\mu})^{n-k} \binom{n-1}{k-1}. \tag{19}$$

3. M2L from multipole expansion at \mathbf{x}_M to local expansion at \mathbf{x}_Λ

$$L_0(\mathbf{x}_\Lambda) \approx M_0(\mathbf{x}_M) \log(z_{\Lambda M}) + \sum_{k=1}^{p-1} \frac{M_k(\mathbf{x}_M)}{z_{\Lambda M}^k}, \tag{20}$$

$$L_n(\mathbf{x}_\Lambda) \approx -\frac{M_0(\mathbf{x}_M)}{(-z_{\Lambda M})^n n} + \sum_{k=1}^{p-1} \frac{(-1)^n M_k(\mathbf{x}_M)}{z_{\Lambda M}^{n+k}} \binom{n+k-1}{k-1}. \tag{21}$$

4. L2L from local expansion at \mathbf{x}_Λ to local expansion at \mathbf{x}_λ ,

$$L_n(\mathbf{x}_\lambda) \approx \sum_{k=n}^{p-1} L_k(\mathbf{x}_\Lambda) z_{\lambda\Lambda}^{k-n} \binom{k}{n}. \tag{22}$$

² <https://github.com/exafmm/exafmm.git>.

5. L2P from local expansion at \mathbf{x}_λ to particle at \mathbf{x}_i ,

$$u_i \approx \Re \left(\sum_{n=0}^{p-1} L_n(\mathbf{x}_\lambda) z_{i\lambda}^n \right). \quad (23)$$

For the P2M, M2M, and M2L kernels, the first term requires special treatment. The expansions are truncated at order p , so the accuracy of the FMM can be controlled by adjusting p . When recurrence relations are used to calculate the powers of z and the combinations they can be calculated at the cost of one multiplication per inner loop (k loop) iteration. In our implementation, we do not construct any matrices during the calculation of these kernels. The P2P kernel is vectorized with the use of SIMD intrinsics, and the $\log()$ function is calculated using a polynomial fit for $\log_2(x)/(x-1)$ using SIMD.

6 Numerical results

In this section we demonstrate the potential of the FMM-based preconditioner by applying it to a number of test problems and comparing it with standard preconditioners. The primary aim is to assess the effectiveness of the preconditioner at reducing the number of Krylov subspace iterations that are required for convergence to a given tolerance. Additionally, we seek to ascertain whether mesh independence is achieved. We defer reporting on performance to Sect. 7. Accordingly, we choose problems that are small enough to enable solution by Matlab.

The Poisson problems are all two dimensional and include examples with homogeneous and inhomogeneous Dirichlet boundary conditions. We additionally present a two-dimensional Stokes flow problem and show that, as predicted, combining the FMM-based Poisson preconditioner with a block diagonal matrix gives an effective preconditioner for the saddle point problem (4).

Throughout, our stopping criterion is a decrease in the relative residual of six orders of magnitude. If such a decrease is not achieved after *maxit* iterations the computations are terminated; this is denoted by ‘—’ in the tables. This maximum number of iterations is stated for each problem below. For all problems and preconditioners the initial iterate is the zero vector.

6.1 The Poisson equation

We first test our preconditioner on three two-dimensional Poisson problems with a constant diffusion coefficient on the domain on $[-1, 1]^2$. Though these problems differ only in their inhomogeneities and appear redundant with respect to the operator, they verify different segments of code for the

Table 1 Preconditioned CG iterations for the relative residual to reduce by six orders of magnitude for the problem with $-\nabla^2 u = 1$ and homogeneous boundary conditions

h	GMG	AMG	FMM	IC
2^{-4}	6	5	5	10
2^{-5}	6	6	6	18
2^{-6}	7	6	6	—
2^{-7}	7	6	6	—
2^{-8}	7	6	7	—

PDE and boundary element discretization. We discretize the problems by Q_1 finite elements using IFISS [24, 63], with default settings. The fast multipole preconditioner is compared with the incomplete Cholesky (IC) factorization [54] with zero fill implemented in Matlab and the algebraic multigrid (AMG) and geometric multigrid (GMG) methods in IFISS. Within the GMG preconditioner we select point-damped Jacobi as a smoother instead of the default ILU, which is less amenable to parallelization. Otherwise, default settings for both multigrid methods are used. For all preconditioners, *maxit* = 20 and we apply preconditioned conjugate gradients.

Our first example is the first reference problem in Elman *et al.* [25, Section 1.1] for which

$$\nabla^2 u = 1 \text{ in } \Omega = [-1, 1]^2, \quad u = 0 \text{ on } \Gamma.$$

Table 1 lists the preconditioned CG iterations for each preconditioner applied. The FMM preconditioner, as well as GMG and AMG appear to give mesh-independent convergence, although the incomplete Cholesky factorization does not.

In Table 2 we list the eigenvalues of the FMM preconditioned stiffness matrix for $h = 2^{-4}, 2^{-5}$ and 2^{-6} . It is clear that the smallest eigenvalue of A decreases as the mesh is refined leading to an increase in the condition number; this is particularly problematic for Krylov subspace methods, such as CG iteration, whose iteration count can grow as the square root of the condition number. However, the eigenvalues of the FMM-preconditioned matrix are bounded away from the origin in a small interval that does not increase in size as the mesh is refined. This hints at spectral equivalence between the FMM-based preconditioner and the stiffness matrix which is unsurprising given that FMM is derived from the exact inverse of the continuous problem. The condition number appears to be bounded, which explains the mesh-independent convergence observed.

Our second example is the third reference problem from Elman *et al.* [25, Section 1.1] posed on $[-1, 1]^2$ which is characterized by inhomogeneous Dirichlet boundary conditions and the analytic solution

Table 2 Smallest (λ_{min}) and largest (λ_{max}) eigenvalues and condition number (κ) of the stiffness matrix A and FMM-preconditioned matrix $P^{-1}A$ for the problem with $-\nabla^2 u = 1$ and homogeneous boundary conditions

h	$\lambda_{min}(A)$	$\lambda_{max}(A)$	$\kappa(A)$	$\lambda_{min}(P^{-1}A)$	$\lambda_{max}(P^{-1}A)$	$\kappa(P^{-1}A)$
2^{-4}	0.076	3.94	52	0.73	1.29	1.77
2^{-5}	0.019	3.99	207	0.72	1.29	1.79
2^{-6}	0.005	4.00	830	0.72	1.30	1.80

Table 3 Preconditioned CG iterations for the relative residual to reduce by six orders of magnitude for the problem with $-\nabla^2 u = 0$ and inhomogeneous boundary conditions

h	GMG	AMG	FMM	IC
2^{-4}	5	5	5	11
2^{-5}	5	5	5	19
2^{-6}	5	5	5	—
2^{-7}	5	5	5	—
2^{-8}	5	5	5	—

Table 4 Preconditioned CG iterations for the relative residual to reduce by six orders of magnitude for the problem with $-\nabla^2 u = -4$ and inhomogeneous boundary conditions

h	GMG	AMG	FMM	IC
2^{-4}	5	5	5	10
2^{-5}	5	5	5	18
2^{-6}	5	5	5	—
2^{-7}	5	5	5	—
2^{-8}	5	5	5	—

$$u(x, y) = \frac{2(1 + y)}{(3 + x)^2 + (1 + y)^2}.$$

From Table 3 we find that, similarly to the previous problem, the FMM preconditioner and both multigrid preconditioners are mesh independent but the Cholesky preconditioner is not. The FMM preconditioner is also competitive with the multigrid methods. Thus, on systems on which applying the FMM preconditioner is significantly faster than applying the multigrid preconditioners, we will achieve a faster time-to-solution with the former. We note that the eigenvalues and condition numbers obtained for the FMM preconditioned stiffness matrix are the same as those computed for the previous example.

The final problem we consider in this section is the Poisson problem with solution

$$u(x, y) = x^2 + y^2$$

on $[-1, 1]^2$, which has forcing term $f \equiv -4$ in the domain and inhomogeneous Dirichlet boundary conditions. The convergence results for this problem, given in Table 4, are similar to those for the previous problems. They show that the FMM preconditioner gives mesh independent convergence and is competitive with AMG and GMG. We also obtain the same eigenvalue results as for the previous examples.

6.2 Effect of FMM precision on convergence

For the results shown above, the FMM precision was set to preserve six significant digits. However, the FMM can be sped up by reducing precision. Since we are using the FMM as a preconditioner, the accuracy requirements are somewhat lower than that of general applications of FMM. Although

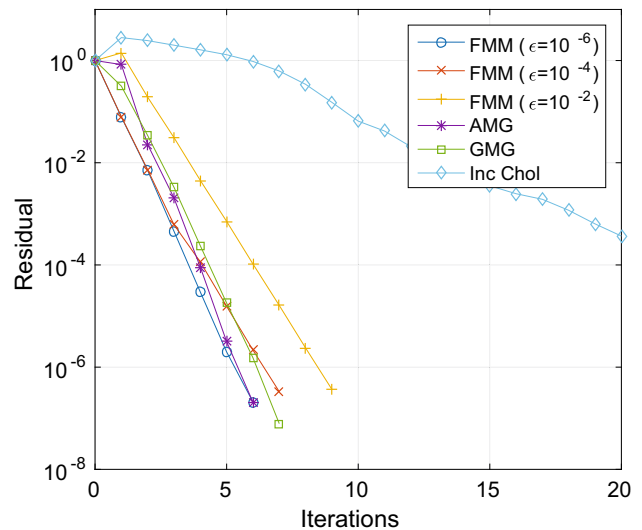


Fig. 4 Convergence rate of the FMM preconditioner with different precision, plotted along with algebraic multigrid, geometric multigrid, and incomplete Cholesky preconditioners. The ϵ represents the precision of the FMM, where $\epsilon = 10^{-6}$ corresponds to six significant digits of accuracy

this balance between the accuracy and speed of FMM is a critical factor for evaluating the usefulness of FMM as a preconditioner, the relation between the FMM precision and convergence rate has not been studied previously.

In Fig. 4 the relative residual at each CG iteration is plotted against the number of iterations for FMM, AMG, GMG, and IC. The problem is the same as in Table 1. Three cases of FMM are used with six, four, and two significant digits of accuracy, respectively. The $\epsilon = 10^{-6}$ case corresponds to the condition for the tests in Tables 1, 2, 3 and 4. Decreasing the FMM accuracy to four digits has little effect during the first few iterations, but slows down the convergence near the end.

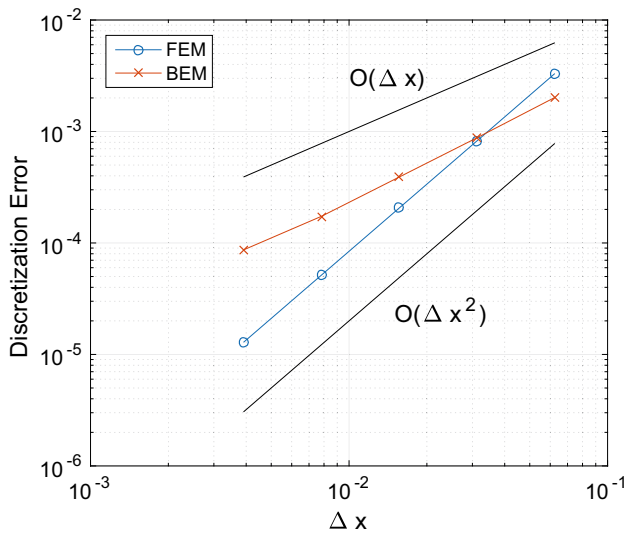


Fig. 5 Convergence of spatial discretization error for the FEM and BEM. The relative L^2 norm of the difference between the analytical solution is plotted against the grid spacing Δx

Decreasing the FMM accuracy further to two digits slows down the convergence somewhat, but is still much better than the incomplete Cholesky.

Increasing the precision of the FMM past six digits does not result in any noticeable improvement because truncation error begins to dominate. We are preconditioning a matrix resulting from a FEM discretization by using an integral equation with Green’s function kernels. Each has its own error, below which algebraic error need not be reduced. We show in Fig. 5 the convergence of spatial discretization error for the FEM and BEM approaches. We use the same reference problem as in Table 1, which has an analytical solution. The discretization error is measured by taking the relative L^2 norm of the difference between the analytical solution and the individual numerical solutions. We see that the FEM is second order and BEM is first order. The five different values of Δx correspond to $h = \{2^{-4}, 2^{-5}, 2^{-6}, 2^{-7}, 2^{-8}\}$, which were used in the previous experiments. For the current range of grid spacing, the discrepancies between the FEM and BEM truncation error is in the range of 10^{-3} – 10^{-4} .

6.3 Stokes problem

Next, we examine convergence for a two-dimensional Stokes flow. The leaky cavity problem [25, Example 5.1.3] on $[-1, 1]$ is discretized by $Q_1 - P_0$ elements in Matlab using IFISS with default settings. As described in Sect. 3, by combining a stiffness matrix preconditioner P_A with the diagonal of the pressure mass matrix P_S , an effective preconditioner (5) for the saddle point system (4) is obtained. Here, we are interested in using the FMM preconditioner for P_A , and we compare its performance with AMG and GMG.

Table 5 Preconditioned MINRES iterations for the relative residual to reduce by six orders of magnitude for the Stokes problem

h	GMG	AMG	FMM
2^{-4}	32	31	35
2^{-5}	32	32	35
2^{-6}	33	32	33
2^{-7}	31	31	33
2^{-8}	31	31	31

We do not consider the incomplete Cholesky factorization of A because of its poor performance on the stiffness matrix (see Tables 1, 3, 4). We set $maxit = 50$ and apply preconditioned MINRES to the saddle point system (Table 5).

As for the Poisson problem, the FMM-based preconditioner provides a mesh-independent preconditioner that is comparable to algebraic and geometric multigrid. Although for coarser grids two or three more iterations are required by the FMM preconditioner than the AMG preconditioner, if each iteration is faster the time-to-solution may be lower.

6.4 Variable coefficient Poisson equation

For sufficiently smooth diffusion coefficient variation, we can precondition the variable coefficient problem with the constant coefficient problem, since they are spectrally equivalent. We test this approach on the variable coefficient Poisson equation of the form

$$-\nabla \cdot (a \nabla u) = 1 \text{ in } \Omega, \\ u = 0 \text{ on } \Gamma,$$

Table 6 Preconditioned CG iterations for the relative residual to reduce by six orders of magnitude ($h = 2^{-6}$, $m = 1$ and $n = 1$)

μ	AMG	FMM
2^{-16}	6	6
2^{-8}	6	6
2^{-4}	6	7
2^{-2}	6	8

Table 7 Preconditioned CG iterations for the relative residual to reduce by six orders of magnitude ($h = 2^{-6}$, $\mu = 2^{-4}$)

m	n	AMG	FMM
1	1	6	7
2	2	6	7
4	4	6	7
8	8	6	7
16	16	6	7

Table 8 Preconditioned CG iterations for the relative residual to reduce by six orders of magnitude ($m = 4$, $h = 2^{-6}$, $\mu = 2^{-4}$)

n	AMG	FMM
1	6	7
2	6	7
4	6	7
8	6	7
16	6	7

Table 9 Smallest (λ_{min}) and largest (λ_{max}) eigenvalues and condition number (κ) of the stiffness matrix A and FMM-preconditioned matrix $P^{-1}A$ with $\mu = 2^{-4}$

h	m	n	$\lambda_{min}(A)$	$\lambda_{max}(A)$	$\kappa(A)$	$\lambda_{min}(P^{-1}A)$	$\lambda_{max}(P^{-1}A)$	$\kappa(P^{-1}A)$
2^{-4}	3	3	0.0759	3.9923	53	0.4423	1.0000	2.26
2^{-5}	6	6	0.0192	4.0199	209	0.4371	1.0040	2.29
2^{-6}	12	12	0.0048	4.0280	839	0.4360	1.0061	2.31

where

$$a = 1 + \mu(\sin(m\pi x) \sin(n\pi y)).$$

Tables 6, 7, and 8 show that the FMM preconditioner and both multigrid preconditioners achieve mesh independent convergence for different amplitudes μ and frequencies m and n . Also, the FMM preconditioner is competitive with the algebraic multigrid method requiring comparable number of iterations.

Similar to Table 2, Table 9 shows that the eigenvalues of the FMM-preconditioned matrix are bounded away from the origin.

7 Performance analysis

In this section we evaluate the performance of the FMM-based preconditioner by comparing its time-to-solution to an algebraic multigrid code BoomerAMG. We have implemented our FMM-preconditioner in PETSc [5,6] via PetIGA [55]. PetIGA is a software layer that sits on top of PETSc that facilitates NURBS-based Galerkin finite element analysis. For our present analysis, we simply use PetIGA to reproduce the same finite element discretization as the tests shown in Sect. 6, but in a high performance computing environment. We select the first problem in Sect. 6.1 with $-\nabla^2 u = 1$ and homogeneous Dirichlet boundary conditions for the following performance evaluation.

All codes that were used for the current study are publicly available. A branch of PetIGA that includes the FMM preconditioner is hosted on bitbucket.³

All calculations were performed on the TACC Stampede system without using the coprocessors. Stampede has 6400 nodes, each with two Xeon E5-2680 processors and one Intel Xeon Phi SE10P coprocessor and 32GB of memory. We used the Intel compiler (version 13.1.0.146) and configured PETSc with “COPTFLAGS=-O3 FOPTFLAGS=-O3 -with-clanguage=cxx -download-fblaslapack -download-hypre -download-metis -download-parmetis -download-superlu_dist -with-debugging=0”.

³ <https://bitbucket.org/rioyokota/petiga-fmm>.

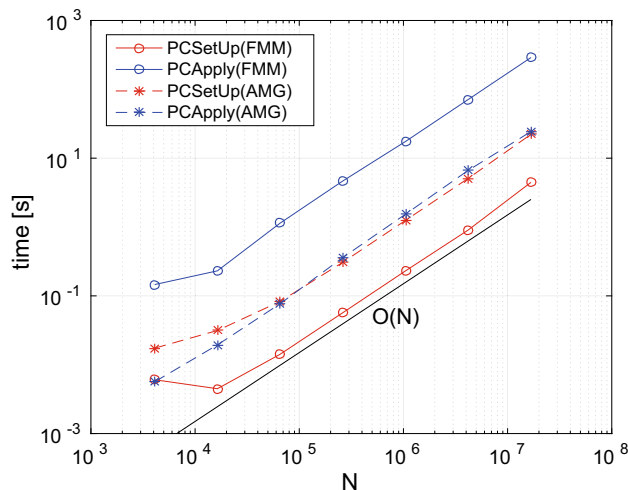


Fig. 6 Time-to-solution for different problem sizes of the FMM and AMG preconditioners on a single core of a Xeon E5-2680

7.1 Serial results

We first evaluate the serial performance of our method using the same two-dimensional Poisson problem used in Sect. 2. We confirmed that the iteration counts shown in Table 1 did not change for the PETSc version of our code. Then, we measured the time-to-solution for different problem sizes. Since the domain size is $[-1, 1]$, the grid spacing of $h = \{2^{-4}, 2^{-5}, 2^{-6}, 2^{-7}, 2^{-8}\}$ in Sect. 6 correspond to a grid size of $N = \{32^2, 64^2, 128^2, 256^2, 512^2\}$. In the PETSc version, the time-to-solution improves significantly so we tested for larger problem sizes of $N = \{64^2, 128^2, 256^2, 512^2, 1024^2, 2048^2, 4096^2\}$.

The time-to-solution is plotted against the problem size N in Fig. 6. Since we are using PETSc, it is trivial to change the preconditioner to AMG by passing the option “-pc_type hypre” during runtime. Therefore, the time-to-solution of BoomerAMG is shown as a reference in the same figure. For BoomerAMG we compared different relaxation, coarsening, and interpolation methods and found that “-pc_hypre_boomeramg_relax_type_all backward-SOR/Jacobi -pc_hypre_boomeramg_coarsen_type modifiedRuge-Stueben -pc_hypre_bommeramg_interp_type classical” gives the best performance.

Both FMM and AMG runs are serial, where we used a single MPI process and a single thread. The majority of the time goes into the setup of the preconditioner “PCSetUp” and the actual preconditioning “PCApply”, so only these events are shown in the legend. The “PCSetUp” is called only once for the entire run, while “PCApply” is called every iteration. For the present runs, both FMM and AMG required six iterations for the relative residual to drop six digits, so all runs are calling “PCApply” six times. The order of expansion for the FMM is set to $p = 6$ and $\theta = 0.4$, which gives about six significant digits of accuracy. With this accuracy for the FMM, we are still able to converge in six iterations. The P2P kernel in the FMM code is performed in single precision using SIMD intrinsics, but this does not prevent us from reaching the required accuracy of six significant digits because we use Kahan’s summation technique [47] for the reduction.

Figure 6 shows that both the FMM and AMG possess $\mathcal{O}(N)$ asymptotic behavior. The FMM seems to have a slower preconditioning time, but a much faster setup time compared to AMG. The FMM also has a constant overhead which becomes evident when N is small. In summary, the time-to-solution of the FMM is approximately an order of magnitude larger than that of AMG for the serial runs. This is consistent with the intuition that FMM is not the preconditioner of choice for solving small problems on a single core. We will show in the following section that the FMM becomes competitive when scalability comes into the picture.

7.2 Parallel results

Using the same Poisson problem, we now compare the performance of FMM and AMG for parallel runs on Stampede. We also compare with a sparse direct solver MUMPS by invoking at runtime “-ksp_type preonly -pc_type lu -pc_factor_mat_solver_package mumps”.

The strong scaling of FMM, AMG, and MUMPS are shown in Fig. 7. We use the largest grid size in the previous runs $N = 4096^2$. Stampede has 16 cores per node so all runs first parallelize among these cores and then among the nodes after the 16 cores are filled. The FMM strong scales quite well up to 1024 cores, while the parallel efficiency of AMG starts to decrease after 128 cores. The sparse direct solver has a much larger time-to-solution even on a single core, and is much less scalable than the other two hierarchical preconditioners. It is worth mentioning that the setup cost of the direct solver is dominant and so if several linear systems are solved with the same coefficient matrix then this cost is amortized. For this particular Poisson problem on this particular machine using this particular FMM code we see an advantage over BoomerAMG past 512 cores.

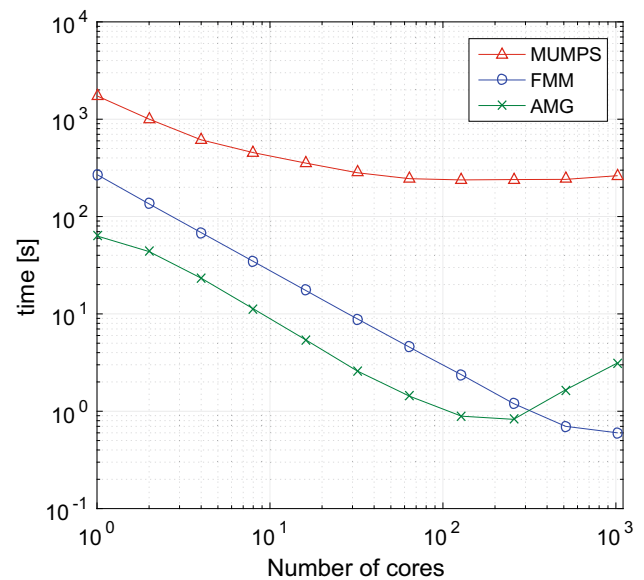


Fig. 7 Strong scaling of the 2-D FMM and AMG preconditioners

7.3 Extension to 3-D

The results above are all two-dimensional. A natural question that arises is whether the extension to 3-D is straightforward, and whether FMM will still be competitive as a preconditioner or not. Our results showed that a dominant part of the calculation time for the FMM preconditioner is the “PCApply” stage, which is the dual tree traversal for calculation of M2L and P2P kernels. For 3-D kernels, the M2L operation is much more complicated so the calculation time of the FMM will increase, even for the same number of unknowns N . In the traditional use of FMM as a solver, the M2L operation has a computational complexity of $\mathcal{O}(p^2)$ in 2-D and at least $\mathcal{O}(p^2 \log p)$ in 3-D, $\mathcal{O}(p^3)$ in our 3-D implementation. However, FMM may exhibit lower order asymptotic behavior when used as a preconditioner in the low accuracy regime making the difference in computational complexity smaller [72].

Figure 8 shows the calculation time of our 2-D FMM and 3-D FMM, both for the Laplace kernel with four significant digits of accuracy on a single core of a Xeon E5-2680, 2.7 GHz CPU. The problem size N varies from 10^5 to 10^7 . We see that the 3-D FMM is about an order of magnitude slower than the 2-D FMM for the same problem size. Nevertheless, Fig. 9 shows that the three-dimensional FMM preconditioner strong scales quite well up to 1024 cores for $N = 256^3$ when compared to BoomerAMG with these configurations “-pc_hyre_boomeramg_coarsen_type hmis -pc_hyre_boomeramg_interp_type ext+i -pc_hyre_boomeramg_p_max 4 -pc_hyre_boomeramg_agg_n1 1”, and GMG with “-da_grid_x 5 -da_grid_y 5 -pc_mg_levels 5”.

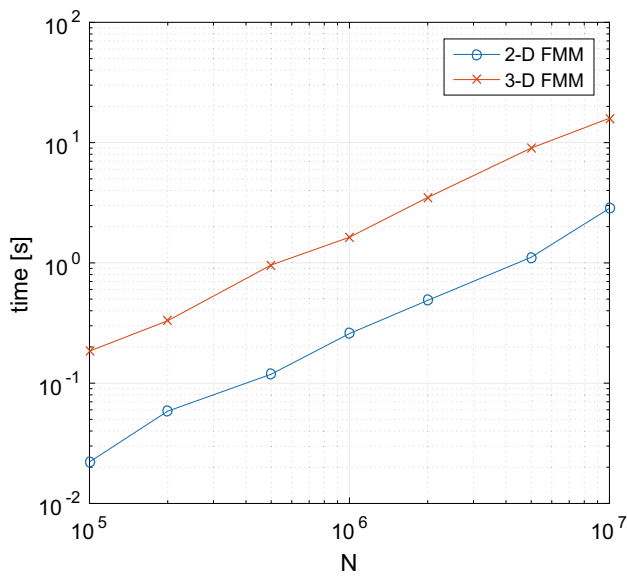


Fig. 8 Calculation time of 2-D and 3-D FMM for the same problem size

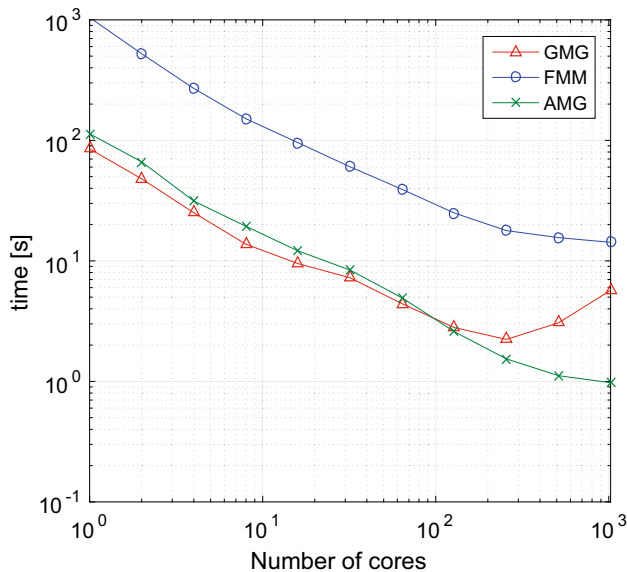
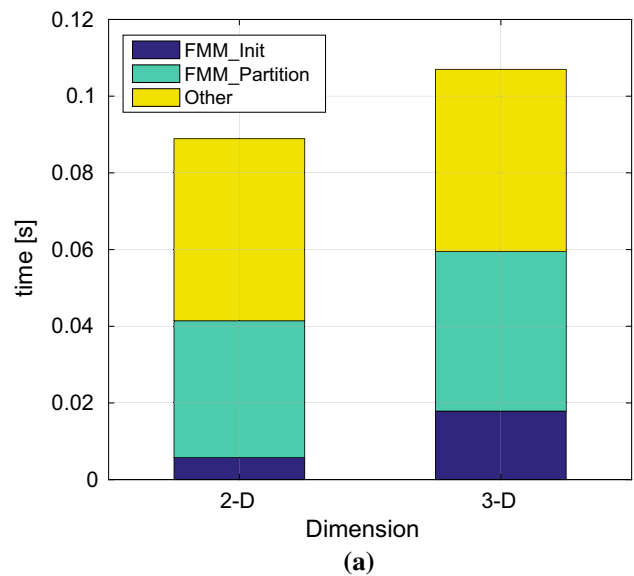


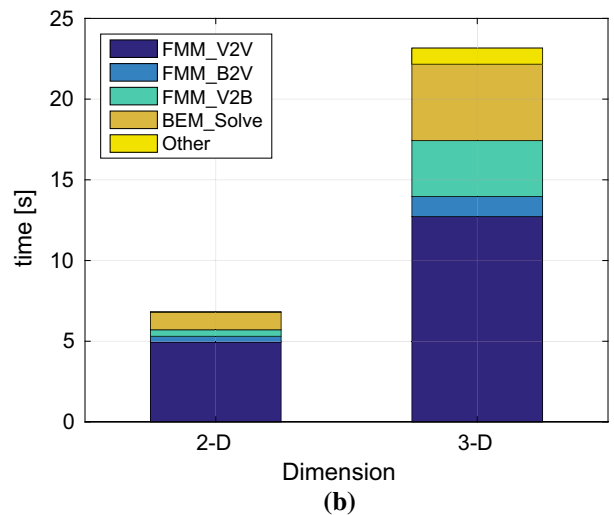
Fig. 9 Strong scaling of the 3-D FMM and AMG preconditioners

These runs were performed on Shaheen II which is a Cray XC40 with 6174 compute nodes, each with two 16-core Intel Haswell CPUs (Intel®Xeon®E5-2698 v3). The nodes are connected by a dragonfly network using the Aries interconnect.

Figure 10 shows the timing breakdown for “PCSetUp” and “PCApply” of the 2-D and 3-D FMM preconditioners for the same problem size on a single core of a Xeon E5-2698. Two routines are shown in Fig. 10a for “PCSetUp”, “FMM_Init” which initializes the FMM, and “FMM_Partition” which performs Orthogonal Recursive Bisection (ORB) of the domain. For the “PCApply” stage, Fig. 10b, the FMM functions



(a)



(b)

Fig. 10 Time breakdown for “PCSetUp” and “PCApply” of the 2-D and 3-D FMM preconditioners for the same problem size on a single core ($N = 512^2$ in 2-D and $N = 64^3$ in 3-D)

“FMM_V2B”, “FMM_B2V”, and “FMM_V2V” perform a volume integral for every point on the boundary, a boundary integral for every point in the volume, and a volume integral for every point in the volume, respectively. “BEM_Solve” is the BEM solve used to obtain u in (8).

8 Conclusions

The Fast Multipole Method, originally developed as a free-standing solver, can be effectively combined with Krylov iteration as a scalable and highly performant preconditioner for traditional low-order finite discretizations of elliptic boundary value problems. In model problems it performs similarly

to algebraic multigrid in convergence rate, while excelling in scalings where AMG becomes memory bandwidth-bound locally and/or syn-chronization-bound globally. The focus of this study is on motivating and describing the fundamentals of the FMM-based preconditioner. For this purpose, we choose simple Poisson problems where standard multigrid methods are known to perform well. Additional algorithmic development and additional testing of implementations on emerging architectures are necessary to more fully define the niche in which FMM is the preconditioner of choice. An example is the indefinite Helmholtz equation, where multigrid methods have severe convergence problems [65] while FMM-based preconditioners maintain convergence independent of mesh and wavenumber, over a moderate range of wavenumbers [44].

No preconditioner considered in isolation can address the fundamental architectural challenges of Krylov methods for sparse linear systems, which are being simultaneously adapted to less synchronization tolerant computational environments through pipelining, but it is important to address the bottlenecks of preconditioning this most popular class of solvers by making a wide variety of tunable preconditioners available and better integrating them into the overall solver. Fast multipole-based preconditioners are demonstrably ready to play an important role in the migration of sparse iterative solvers to the exascale.

Acknowledgements The authors would like to acknowledge the open source software packages that made this work possible: PETSc [5,6], PetIGA [55] and IFISS [24,63]. We thank Dave Hewett, David May, Andy Wathen, and Ulrike Yang for helpful discussions and comments and are indebted to Nathan Collier for his help with the PetIGA framework. We thank Lorena Barba for the support in developing ExaFMM. This publication was based on work supported in part by Award No KUK-C1-013-04, made by King Abdullah University of Science and Technology (KAUST). This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575, and the resources of the KAUST Supercomputing Laboratory.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Aarseth, S.J.: Dynamical evolution of clusters of galaxies, I. *Mon. Not. R. Astron. Soc.* **126**, 233 (1963)
- Adams, M.F., Brown, J., Knepley, M., Samtaney, R.: A multigrid technique for data locality. *SIAM J. Sci. Comput.* **38**, C426–C440 (2016)
- Aminfar, A., Ambikasaran, S., Darve, E.: A fast block low-rank dense solver with applications to finite-element matrices. *J. Comput. Phys.* **304**(C), 170–188 (2016)
- Baker, A.H., Falgout, R.D., Gamblin, T., Kolev, T.V., Schulz, M., Yang, U.M.: Scaling algebraic multigrid solvers: on the road to exascale. In: Bischof, C., Hegering, H.-G., Nagel, W.E., Wittum, G. (eds.) *Competence in High Performance Computing*, pp. 215–226. Springer, Berlin (2012)
- Balay, S., Brown, J., Buschelman, K., Eijkhout, V., Gropp, W.D., Kaushik, D., Knepley, M.G., McInnes, L.C., Smith, B.F., Zhang, H.: *PETSc Users Manual*. Technical Report ANL-95/11—Revision 3.4. Argonne National Laboratory, Argonne (2013)
- Balay, S., Brown, J., Buschelman, K., Gropp, W.D., Kaushik, D., Knepley, M.G., McInnes, L.C., Smith, B.F., Zhang, H.: *PETSc Web Page* (2013). <http://www.mcs.anl.gov/petsc>
- Banerjee, P. K.: Nonlinear problems of potential flow. In: Banerjee, P. K., Butterfield, R. (eds) *Developments in boundary element Methods-I*, pp. 21–30. Applied Science Publishers, London (1979)
- Barba, L.A., Yokota, R.: How will the fast multipole method fare in the exascale era. *SIAM News* **46**(6), 1–3 (2013)
- Benzi, M., Golub, G.H., Liesen, J.: Numerical solution of saddle point problems. *Acta Numer.* **14**, 1–137 (2005)
- Börm, S.: Hybrid cross approximation of integral operators. *Numer. Math.* **101**, 221–249 (2005)
- Börm, S., Hackbusch, W.: Hierarchical quadrature for singular integrals. *Computing* **74**(2), 75–100 (2005)
- Brandt, A.: Multi-level adaptive techniques (MLAT) for partial differential equations: ideas and software. In: Rice, J.R. (ed.) *Mathematical Software*, vol. III, pp. 277–318. Academic Press, Cambridge (1977)
- Brunton, I.: *Solving Variable Coefficient Partial Differential Equations Using the Boundary Element Method*. Ph.D. thesis. University of Auckland (1996)
- Cahouet, J., Chabard, J.-P.: Some fast 3D finite element solvers for the generalized Stokes problem. *Int. J. Numer. Methods Fluids* **8**, 869–895 (1988)
- Carrier, J., Greengard, L., Rokhlin, V.: A fast adaptive multipole algorithm for particle simulations. *SIAM J. Sci. Stat. Comput.* **9**(4), 669–686 (1988)
- Chandrasekaran, S., Gu, M., Pals, T.: A fast ULV decomposition solver for hierarchically semiseparable representations. *SIAM J. Matrix Anal. Appl.* **28**(3), 603–622 (2006)
- Cheng, A.H.-D.: Darcy's flow with variable permeability: a boundary integral solution. *Water Resour. Res.* **20**, 980–984 (1984)
- Clements, D.L.: A boundary integral equation method for the numerical solution of a second order elliptic equation with variable coefficients. *J. Aust. Math. Soc.* **22**(Series B), 218–228 (1980)
- Concus, P., Golub, G.H.: Use of fast direct methods for the efficient numerical solution of nonseparable elliptic equations. *SIAM J. Numer. Anal.* **10**, 1103–1120 (1973)
- Czechowski, K., McClanahan, C., Battaglini, C., Iyer, K., Yeung, P.-K., Vuduc, R.: On the communication complexity of 3D FFTs and its implications for exascale. In: *Proceedings of the 26th ACM International Conference on Supercomputing*, pp. 205–214 (2012)
- Dehnen, W.: A hierarchical O(N) force calculation algorithm. *J. Comput. Phys.* **179**(1), 27–42 (2002)
- Demmel, J., Grigori, L., Hoemmen, M.F., Langou, J.: Communication-avoiding parallel and sequential QR factorizations. Technical Report EECS-2008-74. UC Berkeley (2008)
- Dutt, A., Gu, M., Rokhlin, V.: Fast algorithms for polynomial interpolation integration and differentiation. *SIAM J. Numer. Anal.* **33**(5), 1689–1711 (1996)
- Elman, H., Ramage, A., Silvester, D.: Algorithm 866: IFISS, a Matlab toolbox for modelling incompressible flow. *ACM Trans. Math. Softw.* **33**, 2–14 (2007)
- Elman, H.C., Silvester, D.J., Wathen, A.J.: *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*. Oxford University Press, Oxford (2005)

26. Gahvari, H., Baker, A.H., Schulz, M., Yang, U.M., Jordan, K.E., Gropp, W.: Modeling the performance of an algebraic multigrid cycle on HPC platforms. In: Proceedings of the International Conference on Supercomputing, pp. 172–181 (2011)
27. Gholami, A., Malhotra, D., Sundar, H., Biros, G.: FFT, FMM, or multigrid? A comparative study of state-of-the-art poisson solvers for uniform and nonuniform grids in the unit cube. *SIAM J. Sci. Comput.* **38**(3), C280–C306 (2016)
28. Golub, G.H., van Loan, C.F.: *Matrix Computations*. The John Hopkins University Press, Baltimore (1996)
29. Golub, G.H., Varga, R.S.: Chebyshev semi-iterative methods, successive overrelaxation iterative methods and second order Richardson iterative methods. Part I. *Numer. Math.* **3**, 147–156 (1961)
30. Graham, I.G., Lechner, P.O., Scheichl, R.: Domain decomposition for multiscale PDEs. *Numer. Math.* **106**, 589–626 (2007)
31. Grasedyck, L., Hackbusch, W.: Construction and arithmetics of H-matrices. *Computing* **70**, 295–334 (2003)
32. Grasedyck, L., Hackbusch, W., Kriemann, R.: Performance of H-LU preconditioning for sparse matrices. *Comput. Methods Appl. Math.* **8**(4), 336–349 (2008)
33. Greenbaum, A.: *Iterative Methods for Solving Linear Systems*. SIAM, Philadelphia (1997)
34. Greengard, L., Gueyffier, D., Martinsson, P.G., Rokhlin, V.: Fast direct solvers for integral equations in complex three dimensional domains. *Acta Numer.* **18**, 243–275 (2009)
35. Greengard, L., Rokhlin, V.: A fast algorithm for particle simulations. *J. Comput. Phys.* **73**(2), 325–348 (1987)
36. Gu, M., Eisenstat, S.C.: Efficient algorithms for computing a strong rank-revealing qr factorization. *SIAM J. Sci. Comput.* **17**(4), 848–869 (1996)
37. Gumerov, N.A., Duraiswami, R.: Fast multipole methods on graphics processors. *J. Comput. Phys.* **227**, 8290–8313 (2008)
38. Hackbusch, W.: A sparse matrix arithmetic based on H-matrices. Part I: introduction to H-matrices. *Computing* **62**, 89–108 (1999)
39. Hackbusch, W., Khoromskij, B., Sauter, S.: On H²-matrices. In: Bungartz, H.-J., Hoppe, R., Zenger, C. (eds.) *Lectures on Applied Mathematics*, pp. 9–29. Springer, Berlin (2000)
40. Henson, V.E., Yang, U.M.: BoomerAMG: a parallel algebraic multigrid solver and preconditioner. *Appl. Numer. Math.* **41**, 155–177 (2002)
41. Hestenes, M.R., Stiefel, E.: Methods of conjugate gradients for solving linear systems. *J. Res. Natl. Bur. Stand.* **49**, 409–436 (1952)
42. Ho, K.L., Greengard, L.: A fast direct solver for structured linear systems by recursive skeletonization. *SIAM J. Sci. Comput.* **34**(5), A2507–A2532 (2012)
43. Hundsdorfer, W., Verwer, J.G.: *Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations*. Springer, Berlin (2003)
44. Ibeid, H., Yokota, R., Keyes, D.: A Matrix-free Preconditioner for the Helmholtz Equation Based on the Fast Multipole Method (2016). [arXiv:1608.02461](https://arxiv.org/abs/1608.02461)
45. Ibeid, H., Yokota, R., Keyes, D.: A performance model for the communication in fast multipole methods on high-performance computing platforms. *Int. J. High Perform. Comput. Appl.* (2016). <https://doi.org/10.1177/1094342016634819>
46. Jung, J., Kobayashi, C., Imamura, T., Sugita, Y.: Parallel implementation of 3D FFT with volumetric decomposition schemes for efficient molecular dynamics simulations. *Comput. Phys. Commun.* **200**, 57–65 (2016)
47. Kahan, W.: Pracniques: further remarks on reducing truncation errors. *Commun. ACM* **8**(1), 40 (1965)
48. Kong, W.Y., Bremer, J., Rokhlin, V.: An adaptive fast direct solver for boundary integral equations in two dimensions. *Appl. Comput. Harmonic Anal.* **31**, 346–369 (2011)
49. Langer, U., Of, G., Steinbach, O., Zulehner, W.: Inexact fast multipole boundary element tearing and interconnecting methods. In: Widlund, O.B., Keyes, D.E. (eds.) *Domain Decomposition Methods in Science and Engineering XVI*. Springer, Berlin (2007)
50. Lashuk, I., Chandramowlishwaran, A., Langston, H., Nguyen, T.-A., Sampath, R., Shringarpure, A., Vuduc, R., Ying, L., Zorin, D., Biros, G.: A massively parallel adaptive fast multipole method on heterogeneous architectures. *Commun. ACM* **55**(5), 101–109 (2012)
51. Liberty, E., Woolfe, F., Martinsson, P.G., Rokhlin, V., Tygert, M.: Randomized algorithms for the low-rank approximation of matrices. *PNAS* **104**(51), 20167–20172 (2007)
52. Malhotra, D., Biros, G.: PVFMM: a parallel kernel independent FMM for particle and volume potentials. *Commun. Comput. Phys.* **18**, 808–830 (2015)
53. Martin, R.M.: *Electronic Structure: Basic Theory and Practical Methods*. Cambridge University Press, Cambridge (2004)
54. Meijerink, J.A., van der Vorst, H.A.: An iterative solution method for linear systems of which the coefficient matrix is a symmetric *M*-matrix. *Math. Comput.* **31**, 148–162 (1977)
55. Collier, V.C.N., Dalcin, L.: PetIGA: High-Performance Isogeometric Analysis. [arxiv \(1305.4452\)](https://arxiv.org/abs/1305.4452) (2013). [arXiv:1305.4452](https://arxiv.org/abs/1305.4452)
56. Paige, C.C., Saunders, M.A.: Solution of sparse indefinite systems of linear equations. *SIAM J. Numer. Anal.* **12**, 617–629 (1975)
57. Pan, C.T.: On the existence and computation of rank-revealing LU factorizations. *Linear Algebra Appl.* **316**, 199–222 (2000)
58. Rjasanow, S.: Adaptive cross approximation of dense matrices. In: *International Association for Boundary Element Methods*, UT Austin, TX, USA, 28–30 May (2002)
59. Saad, Y.: *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia (2003)
60. Saad, Y., Schultz, M.H.: GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.* **7**, 856–869 (1986)
61. Sambavaram, S.R., Sarin, V., Sameh, A., Grama, A.: Multipole-based preconditioners for large sparse linear systems. *Parallel Comput.* **29**, 1261–1273 (2003)
62. Sauter, S.A., Schwab, Ch.: *Boundary Element Methods*. Springer, Heidelberg (2011)
63. Silvester, D., Elman, H., Ramage, A.: *Incompressible Flow and Iterative Solver Software (IFISS) Version 3.2* (2012). <http://www.manchester.ac.uk/ifiss>
64. Teng, S.-H.: Provably good partitioning and load balancing algorithms for parallel adaptive N-body simulation. *SIAM J. Sci. Comput.* **19**(2), 635–656 (1998)
65. Trottenberg, U., Oosterlee, C., Schüller, A.: *Multigrid*. Academic Press, London (2001)
66. Vassilevski, P.S., Yang, U.M.: Reducing communication in algebraic multigrid using additive variants. *Numer. Linear Algebra Appl.* **21**, 275–296 (2014)
67. Wardle, L.J., Crotty, J.M.: Two Dimensional boundary integral equation analysis for non-homogeneous mining applications. In: Brebbia, C.A. (ed) *Recent advances in boundary element methods*. Pentech Press, Plymouth, pp. 233–251 (1978)
68. Warren, M.S., Salmon, J.K.: Astrophysical N-body simulation using hierarchical tree data structures. In: *Proceedings of the 1992 ACM/IEEE Conference on Supercomputing*, pp. 570–576 (1992)
69. Warren, M.S., Salmon, J.K.: A portable parallel particle program. *Comput. Phys. Commun.* **87**, 266–290 (1995)
70. Wathen, A., Rees, T.: Chebyshev semi-iteration in preconditioning for problems including the mass matrix. *Electron. Trans. Numer. Anal.* **34**, 125–135 (2009)
71. Wathen, A.J.: Realistic eigenvalue bounds for the Galerkin mass matrix. *IMA J. Numer. Anal.* **7**, 449–457 (1987)
72. Yokota, R.: An FMM based on dual tree traversal for many-core architectures. *J. Algorithms Comput. Technol.* **7**(3), 301–324 (2013)