

Fast Multipole Method as a Matrix-Free Hierarchical Low-Rank Approximation

Rio Yokota, Huda Ibeid, and David Keyes

Abstract There has been a large increase in the amount of work on hierarchical low-rank approximation methods, where the interest is shared by multiple communities that previously did not intersect. This objective of this article is two-fold; to provide a thorough review of the recent advancements in this field from both analytical and algebraic perspectives, and to present a comparative benchmark of two highly optimized implementations of contrasting methods for some simple yet representative test cases. The first half of this paper has the form of a survey paper, to achieve the former objective. We categorize the recent advances in this field from the perspective of compute-memory tradeoff, which has not been considered in much detail in this area. Benchmark tests reveal that there is a large difference in the memory consumption and performance between the different methods.

1 Introduction

The fast multipole method (FMM) was originally developed as an algorithm to bring down the $\mathcal{O}(N^2)$ complexity of the direct N -body problem to $\mathcal{O}(N)$ by approximating the hierarchically decomposed far field with multipole/local expansions. In its original form, the applicability of FMM is limited to problems that have a Green's function solution, for which the multipole/local expansions can be calculated analytically. Their function is also limited to matrix-vector multiplications, in contrast to the algebraic variants that can perform matrix-matrix multiplication and factorizations. However, these restrictions no longer apply to the FMM since the kernel independent FMM [103] does not require a Green's function, and inverse FMM [2] can be used as the inverse operator instead of the forward mat-vec. Therefore the FMM can be used for a wide range of scientific applications,

R. Yokota (✉)

Tokyo Institute of Technology, 2-12-1 O-okayama Meguro-ku, Tokyo, Japan
e-mail: rioyokota@gsic.titech.ac.jp

H. Ibeid • D. Keyes

King Abdullah University of Science and Technology, 4700 KAUST, Thuwal, Saudi Arabia
e-mail: huda.ibeid@kaust.edu.sa; david.keyes@kaust.edu.sa

© Springer International Publishing AG 2017

T. Sakurai et al. (eds.), *Eigenvalue Problems: Algorithms, Software and Applications in Petascale Computing*, Lecture Notes in Computational Science and Engineering 117, https://doi.org/10.1007/978-3-319-62426-6_17

267

which can be broadly classified into elliptic partial differential equations (PDE) and kernel summation. Integral form of elliptic PDEs can be further categorized into boundary integrals for homogeneous problems, discrete volume integrals, and continuous volume integrals.

Scientific applications of FMM for boundary integrals include acoustics [59, 97], biomolecular electrostatics [105], electromagnetics [33, 42], fluid dynamics for Euler [96] and Stokes [88] flows, geomechanics [92], and seismology [22, 95]. Application areas of FMM for discrete volume integrals are astrophysics [14], Brownian dynamics [75], classical molecular dynamics [84], density functional theory [90], vortex dynamics [106], and force directed graph layout [107]. FMM for continuous volume integrals have been used to solve Schrödinger [108] and Stokes [79] equations. More generalized forms of FMM can be used as fast kernel summation for Bayesian inversion [3], Kalman filtering [74], Machine learning [49, 72], and radial basis function interpolation [54].

All of these applications have in common the key feature that they are global problems where the calculation at every location depends on the values everywhere else. Elliptic PDEs that represent a state of equilibrium, many iterations with global inner products for their solution, dense matrices in boundary integral problems, all-to-all interaction in N -body problems, and kernel summations with global support are all different manifestations of the same source of global data dependency. Due to this global data dependency, their concurrent execution on future computer architectures with heterogeneous and deep memory hierarchy is one of the main challenges of exascale computing. For global problems that require uniform resolution, FFT is often the method of choice, despite its suboptimal communication costs. The methods we describe here have an advantage for global problems that require non-uniform resolution. For such non-uniform global problems multigrid methods are known to do quite well. Whether the reduced synchronization and increased arithmetic intensity of the FMM will become advantageous compared to multigrid on future architectures is something that is yet to be determined.

Many of the original FMM researchers have now moved on to develop algebraic variants of FMM, such as \mathcal{H} -matrix [55], \mathcal{H}^2 -matrix [57], hierarchically semi-separable (HSS) [26], hierarchically block-separable (HBS) [82], and hierarchically off-diagonal low-rank (HODLR) [1] matrices. The differences between these methods are concisely summarized by Ambikasaran and Darve [2]. These algebraic generalizations of the FMM can perform addition, multiplication, and even factorization of dense matrices with near linear complexity. This transition from analytic to algebraic did not happen suddenly, and semi-analytic variants were developed along the way [39, 103]. Optimization techniques for the FMM such as compressed translation operators and their precomputation, also fall somewhere between the analytic and algebraic extremes.

The spectrum that spans purely analytic and purely algebraic forms of these hierarchical low-rank approximation methods, represents the tradeoff between computation (Flops) and memory (Bytes). The purely analytic FMM is a matrix-free \mathcal{H}^2 -matrix-vector product, and due to its matrix-free nature it has very high arithmetic intensity (Flop/Byte) [9]. On the other end we have the purely algebraic

methods, which precompute and store the entire hierarchical matrix. This results in more storage and more data movement, both vertically and horizontally in the memory hierarchy. When the cost of data movement increases faster than arithmetic operations on future architectures, the methods that compute more to store/move less will become advantageous. Therefore, it is important to consider the whole spectrum of hierarchical low-rank approximation methods, and choose the appropriate method for a given pair of application and architecture.

There have been few attempts to quantitatively investigate the tradeoff between the analytic and algebraic hierarchical low-rank approximation methods. Previously, the applicability of the analytic variants were limited to problems with Green's functions, and could only be used for matrix-vector products but not to solve the matrix. With the advent of the kernel-independent FMM (KIFMM) [103] and inverse FMM (IFMM) [2], these restrictions no longer apply to the analytic variants. Furthermore, the common argument for using the algebraic variants because they can operate directly on the matrix without the need to pass geometric information is not very convincing. Major libraries like PETSc offer interfaces to insert one's own matrix free preconditioner as a function, and passing geometric information is something that users are willing to do if the result is increased performance. Therefore, there is no strong reason from the user's perspective to be monolithically inclined to use the algebraic variants. It is rather a matter of choosing the method with the right balance between its analytic (Flops) and algebraic (Bytes) features.

The topic of investigating the tradeoff between analytic and algebraic hierarchical low-rank approximation methods is too broad to cover in a page-constrained article. In the present work, we limit our investigation to the compute-memory tradeoff in a comparison between FMM and HSS for Laplace and Helmholtz kernels. We also investigate the use of FMM as a preconditioner for iterative solutions to the Laplace and Helmholtz problems with finite elements, for which we compare with geometric and algebraic multigrid methods.

2 Hierarchical Low-Rank Approximation: Analytic or Algebraic?

In this section we review the full spectrum of hierarchical low-rank approximations starting from the analytic side and proceeding to the algebraic side. The spectrum is depicted in Fig. 1, where various techniques like between the analytic and algebraic extremes. One can choose the appropriate method for a given architecture to achieve the best performance.




Fig. 1 The compute-memory tradeoff between the analytic and algebraic hierarchical low-rank approximation methods. Various techniques lie between the analytic and algebraic extremes

2.1 Analytic Low-Rank Approximation

On the analytic end of the spectrum, we have classical methods such as the Treecode [10], FMM [8, 50], and panel clustering methods [56]. These methods have extremely high arithmetic intensity (Flop/Byte) due to their matrix-free nature, and are compute-bound on most modern architectures. One important fact is that these are not brute force methods that do unnecessary Flops, but are (near) linear complexity methods that are only doing useful Flops, but they are still able to remain compute-bound. This is very different from achieving high Flops counts on dense matrix-matrix multiplication or LU decomposition that have $\mathcal{O}(N^3)$ complexity. The methods we describe in this section can approximate the same dense linear algebra calculation in $\mathcal{O}(N)$ or $\mathcal{O}(N \log N)$ time.

As an example of the absolute performance of the analytic variants, we refer to the Treecode implementation—`Bonsai`, which scales to the full node of Titan using 18,600 GPUs achieving 24.77 PFlops [14]. `Bonsai`'s performance comes not only from its matrix-free nature, but also from domain specific optimizations for hardcoded quadrupoles and an assumption that all charges are positive. Therefore, this kind of performance cannot be transferred to other applications that require higher accuracy. However, viewing these methods as a preconditioner instead of a direct solver significantly reduces the accuracy requirements [5, 67].

2.2 Fast Translation Operators

A large part of the calculation time of FMM is spent on the translation of multipole expansions to local expansions (or their equivalent charges). Therefore, much work has focused on developing fast translation operators to accelerate this part of the FMM. Rotation of spherical harmonics [94], Block FFT [37], Planewaves [51] are analytic options for fast translation operators.

These translation operators are applied to a pair of boxes in the FMM tree structure that satisfy a certain proximity threshold. This proximity is usually defined as the parent's neighbors' children that are non-neighbors. This produces a list of

boxes that are far enough that the multipole/local expansion converges, but are close enough that the expansion does not converge for the their parents. Such an interaction list can contain up to $6^3 - 3^3 = 189$ source boxes for each target box. Out of these 189 boxes, the ones that are further from the target box can perform the translation operation using their parent box as the source without loss of accuracy. There are a few variants for these techniques that reduce the interaction list size such as the level-skip M2L method [93] and 8, 4, 2-box method [95]. There are also methods that use the dual tree traversal along with the multipole acceptance criterion to construct optimal interaction lists [35], which automates the process of finding the optimal interaction list size.

Another technique to accelerate the translation operators is the use of variable expansion order, as proposed in the very fast multipole method (VFMM) [87], Gaussian VFMM [21], optimal parameter FMM [29], and error controlled FMM [32]. There are two main reasons why spatially varying the expansion order in the translation operators is beneficial. One is because not all boxes in the interaction list are of equal distance, and the boxes that are further from each other can afford to use lower expansion order, while retaining the accuracy. The other reason is because some parts of the domain may have smaller values, and the contribution from that part can afford to use lower expansion order without sacrificing the overall accuracy.

The translation operators can be stored as matrices that operate on the vector of expansion coefficients. Therefore, singular value decomposition (SVD) can be used to compress this matrix [43] and BLAS can be used to maximize the cache utilization [40]. Some methods use a combination of these techniques like Chebychev with SVD [39] and planewave with adaptive cross approximation (ACA) and SVD [61]. The use of SVD is a systematic and optimal way of achieving what the variable expansion order techniques in the previous paragraph were trying to do manually. Precomputing these translation matrices and storing them is a typical optimization technique in many FMM implementations [78].

One important connection to make here is that these matrices for the translation operators are precisely what \mathcal{H}^2 -matrices and *HSS* matrices store in the off-diagonal blocks after compression. One can think of FMM as a method that has the analytical form to generate these small matrices in the off-diagonal blocks, without relying on numerical low-rank approximation methods. To complete this analogy, we point out that the dense diagonal blocks in \mathcal{H}^2 -matrices and *HSS* matrices are simply storing the direct operator (Green's function) in FMM. Noticing this equivalence leads to many possibilities of hybridization among the analytic and algebraic variants. Possibly the most profound is the following. Those that are familiar with FMM know that translation operators for boxes with the same relative positioning are identical. This suggests that many of the entries in the off-diagonal blocks of \mathcal{H}^2 -matrices and *HSS* matrices are identical. For matrices that are generated from a mesh that has a regular structure even the diagonal blocks would be identical, which is what happens in FMMs for continuous volume integrals [78]. This leads to $\mathcal{O}(1)$ storage for the matrix entries at every level of the hierarchy, so the total storage cost of these hierarchical matrices could be reduced to $\mathcal{O}(\log N)$ if the identical entries are not stored redundantly. This aspect is currently

underutilized in the algebraic variants, but seems obvious from the analytic side. By making use of the translational invariance and rotational symmetry of the interaction list one can reduce the amount of storage even further [31, 34, 91]. This also results in blocking techniques for better cache utilization.

2.3 *Semi-analytical FMM*

The methods described in the previous subsection all require the existence of an analytical form of the multipole/local translation operator, which is kernel dependent. There are a class of methods that remove this restriction by using equivalent charges instead of multipole expansions [7, 15, 77]. A well known implementation of this method is the kernel independent FMM (KIFMM) code [103]. There are also variants that use Chebychev polynomials [36], and a representative implementation of this is the Black-box FMM [39]. As the name of these codes suggest, these variants of the FMM have reduced requirements for the information that has to be provided by the user. The translation operators are kernel-independent, which frees the user from the most difficult task of having to provide an analytical form of the translation operators. For example, if one wants to calculate the Matérn function for covariance matrices, or multiquadrics for radial basis function interpolation, one simply needs to provide these functions and the location of the points and the FMM will handle the rest. It is important to note that these methods are not entirely kernel independent or black-box because the user still needs to provide the kernel dependent analytic form of the original equation they wish to calculate. Using the vocabulary of the algebraic variants, one could say that these analytical expressions for the hierarchical matrices are kernel independent only for the off-diagonal blocks, and for the diagonal blocks the analytical form is kernel dependent.

FMM for continuous volume integrals [38] also has important features when considering the analytic-algebraic tradeoff. The volume integrals are often combined with boundary integrals, as well [104]. One can think of these methods as an FMM that includes the discretization process [70]. Unlike the FMM for discrete particles, these methods have the ability to impose regular underlying geometry. This enables the use of precomputation of the direct interaction matrix in the analytic variants [78], and reduces the storage requirements of the dense diagonal blocks in the algebraic variants.

2.4 *Algebraic Low-Rank Approximation*

There are many variants of algebraic low-rank approximation methods. They can be categorized based on whether they are hierarchical, whether they use weak admissibility, or if the basis is nested, as shown in Table 1. For the definition of admissibility see [45]. Starting from the top, \mathcal{H} -matrices [12, 55] are hierarchical, usually use

Table 1 Categorization of algebraic low-rank approximation methods

Method	Hierarchical	Weak admissibility	Nested basis
\mathcal{H} -matrix [55]	Yes	Maybe	No
\mathcal{H}^2 -matrix [57]	Yes	Maybe	Yes
HODLR [1]	Yes	Yes	No
HSS [26]/HBS [82]	Yes	Yes	Yes
BLR [4]	No	Yes	No

standard or strong admissibility, and no nested basis. The analytic counterpart of the \mathcal{H} -matrix is the Treecode. The \mathcal{H}^2 -matrices [16, 57] are also hierarchical and use standard or strong admissibility, but unlike \mathcal{H} -matrices use a nested basis. This brings the complexity down from $\mathcal{O}(N \log N)$ to $\mathcal{O}(N)$. The analytic counterpart of the \mathcal{H}^2 -matrix is the FMM. The next three entries in Table 1 do not have analytic counterparts because analytic low-rank approximations do not converge under weak admissibility conditions. Hierarchical off-diagonal low-rank (HODLR) matrices [1, 6], are basically \mathcal{H} -matrices with weak admissibility conditions. Similarly, hierarchically semi-separable (HSS) [26, 101], and hierarchically block-separable (HBS) [82] matrices are \mathcal{H}^2 -matrices with weak admissibility conditions. The block low-rank (BLR) matrices [4] are a non-hierarchical version of the HODLR, with just the bottom level. A summary of implementations and their characteristics are presented in [89].

For methods that do not have weak admissibility, it is common to use geometrical information to calculate the standard/strong admissibility condition. This dependence on the geometry of the algebraic variants is not ideal. There have been various proposals for algebraic clustering methods [46, 71, 85]. This problem requires even more advanced solutions for high dimension problems [80]. Stronger admissibility is also problem for parallelization since it results in more communication. There have been studies on how to partition hierarchical matrices on distributed memory [68]. There are also methods to reduce the amount of memory consumption during the construction of HSS matrices [73].

The categorization in Table 1 is for the hierarchical matrix structure, and any low-rank approximation method can be used with each of them during the compression phase. The singular value decomposition is the most naïve and expensive way to calculate a low-rank approximation. QR or LU decompositions can be used to find the numerical rank by using appropriate pivoting. Rank-revealing QR [24] has been proposed along with efficient pivoting strategies [25, 53, 64]. Rank-revealing LU [23] also requires efficient pivoting strategies [65, 66, 83]. Rank-revealing LU is typically faster than rank-revealing QR [86]. There are other methods like the pseudo-skeletal method [44] and adaptive cross approximation (ACA) [11, 13], which do not yield the optimal low-rank factorizations but have a much lower cost. ACA has a better pivoting strategy than pseudo-skeletal methods, but can still fail because of bad pivots [18]. The hybrid cross approximation (HCA) [17] has the same proven convergence as standard interpolation but also the same

efficiency as ACA. Yet another class of low-rank approximation is the interpolative decomposition (ID) [28, 82], where a few of its columns are used to form a well-conditioned basis for the remaining columns. ID can be combined with randomized methods [76], which has much lower complexity. For a nice review on these randomized methods see [58].

3 Low-Rank Approximation for Factorization

3.1 Sparse Matrix Factorization

Hierarchical low-rank approximation methods can be used as direct solvers with controllable accuracy. This makes them useful as preconditioners within a Krylov subspace method, which in turn reduces the accuracy requirements of the low-rank approximation. High accuracy and completely algebraic methods are demanding in terms of memory consumption and amount of communication, so they are unlikely to be the optimal choice unless they are the only solution to that problem.

There are two ways to use hierarchical low-rank approximations for factorization of a sparse matrix. The first way is to perform the LU decomposition on the sparse matrix, and use hierarchical low-rank approximations for the dense blocks that appear during the process [98, 100, 101]. The other way is to represent the sparse matrix with a hierarchical low-rank approximation and perform an LU decomposition on it [46–48]. The main difference is whether you view the base method as the nested dissection and the additional component as HLRA or vice versa. The former has the advantage of being able to leverage the existing highly optimized sparse direct solvers, whereas the latter has the advantage of handling both sparse and dense matrices with the same infrastructure.

There are various ways to minimize the fill-in and compress the dense blocks during factorization. These dense blocks (Schur complements) are an algebraic form of the Green's function [99], and have the same low-rank properties [27] stemming from the fact that some of the boundary points in the underlying geometry are distant from each other. Formulating a boundary integral equation is the analytical way of arriving to the same dense matrix. From an algebraic point of view, the sparse matrix for the volume turns into a dense matrix for the boundary, through the process of trying to minimize fill-in. Considering the minimization of fill-in and the compression of the dense matrices in separate phases leads to methods like HSS + multifrontal [98, 100, 101].

Ultimately, minimizing fill-in and minimizing off-diagonal rank should not be conflicting objectives. The former depends on the connectivity and the latter depends on the distance in the underlying geometry. In most applications, the closer points are connected (or interact) more densely, so reordering according to the distance should produce near optimal ordering for the connectivity as well. The same can be said about minimizing communication for the parallel implementation

of these methods. Mapping the 3-D connectivity/distance to a 1-D locality in the memory space (or matrix column/row) is what we are ultimately trying to achieve.

3.2 *Dense Matrix Factorization*

The methods in the previous subsection are direct solvers/preconditioners for sparse matrices. As we have mentioned, there is an analogy between minimizing fill-in in sparse matrices by looking at the connectivity, and minimizing the rank of off-diagonal blocks of dense matrices by looking at the distance. Using this analogy, the same concept as nested dissection for sparse matrices can be applied to dense matrices. This leads to methods like the recursive skeletonization [62], or hierarchical Poincare-Steklov (HPS) [41, 81]. HPS is like a bottom-up version of what nested dissection and recursive skeletonization do top-down. For high contrast coefficient problems, it makes sense to construct the domain dissection bottom-up, to align the bisectors with the coefficient jumps. There are also other methods that rely on a similar concept [19, 52, 69, 102]. Furthermore, since many of these methods use weak admissibility with growing ranks for 3-D problems, it is useful to have nested hierarchical decompositions, which is like a nested dimension reduction. In this respect, the recursive skeletonization has been extended to hierarchical interpolative factorization (HIF) [63], the HSS has been extended to HSS2D [99]. There is also a combination of HSS and Skeletonization [30]. There are methods that use this nested dimension reduction concept without the low-rank approximation [60] in the context of domain decomposition for incomplete LU factorization. One method that does not use weak admissibility is the inverse FMM [2], which makes it applicable to 3-D problems in $\mathcal{O}(N)$ without nested dimension reduction.

4 Experimental Results

4.1 *FMM vs. HSS*

There have been few comparisons between the analytic and algebraic hierarchical low-rank approximation methods [20]. From a high performance computing perspective, the practical performance of highly optimized implementations of these various methods is of great interest. There have been many efforts to develop new methods in this area, which has resulted in a large amount of similar methods with different names without a clear overall picture of their relative performance on modern HPC architectures. The trend in architecture where arithmetic operations are becoming cheap compared to data movement, is something that must be considered

carefully when predicting which method will perform better on computers of the future.

We acknowledge that the comparisons we present here are far from complete, and much more comparisons between all the different methods are needed in order to achieve our long term objective. The limitation actually comes from the lack of highly optimized implementations of these methods that are openly available to us at the moment.

In the present work we start by comparing exaFMM—a highly optimized implementation of FMM, with STRUMPACK—a highly optimized implementation of HSS. We select the 2D and 3D Laplace equation on uniform lattices as test cases. For HSS we directly construct the compressed matrix by calling the Green’s function in the randomized low-rank approximation routine. We perform the matrix-vector multiplication using the FMM and HSS, and measure the time for the compression/precalculation and application of the matrix-vector multiplication. We also measure the peak memory consumption of both methods.

The elapsed time for the FMM and HSS for different problem sizes is shown in Fig. 2. In order to isolate the effect of the thread scalability of the two methods, these runs are performed on a single core of a 12-core Ivy Bridge (E5-2695 v2). For the 2D Laplace equation, the FMM shows some overhead for small N , but is about 3 orders of magnitude faster than HSS for larger problems. For the 3D Laplace equation, the FMM is about 2 orders of magnitude faster than HSS for smaller N , but HSS exhibits non-optimal behavior for large N because the rank keeps growing.

The large difference in the computational time is actually coming from the heavy computation in the sampling phase and compression phase of the HSS. In Fig. 3,




Fig. 2 Elapsed time for the matrix-vector multiplication using FMM and HSS for different problem sizes




Fig. 3 Percentage of the computation time of HSS for different problem sizes




Fig. 4 Peak memory usage of FMM and HSS for the 3D Laplace equation

we show the percentage of the computation time of HSS for different problem sizes N . “Sample” is the sampling time, “Compress” is the compression time, and “Mat-Vec” is the matrix-vector multiplication time. We can see that the sampling is taking longer and longer as the problem size increases. This is because the rank k increases with the problem size N , and both sampling and compression time increase with the k and N .

The peak memory usage of FMM and HSS is shown in Fig. 4 for the 3D Laplace equation. We see that the FMM has strictly $\mathcal{O}(N)$ storage requirements, but since

the rank in the HSS grows for 3D kernels it does not show the ideal $\mathcal{O}(N \log N)$ behavior. The disadvantage of HSS is two-fold. First of all, its algebraic nature requires it to store the compressed matrix, where as the FMM is analytic and therefore matrix-free. Secondly, the weak admissibility causes the rank to grow for 3D problems, and with that the memory consumption grows at a suboptimal complexity.

4.2 FMM vs. Multigrid

If we are to use the FMM as a matrix-free $\mathcal{O}(N)$ preconditioner based on hierarchical low-rank approximation, the natural question to ask is “How does it compare against multigrid?”, which is a much more popular matrix-free $\mathcal{O}(N)$ preconditioner for solving elliptic PDEs. We perform a benchmark test similar to the one in the previous subsection, for the Laplace equation and Helmholtz equation on a 3D cubic lattice $[-1, 1]^3$, but for this case we impose Dirichlet boundary conditions at the faces of the domain. The preconditioners are used inside a Krylov subspace solver. The runs were performed on Matlab using a finite element package IFISS. Our fast multipole preconditioner is compared with the incomplete Cholesky (IC) factorization with zero fill implemented in Matlab and the algebraic multigrid (AMG) and geometric multigrid (GMG) methods in IFISS. The FMM code is written in C and called as a MEX function.

The convergence rate of the FMM and Multigrid preconditioners for the Laplace equation is shown in Fig. 5, for a grid spacing of $h = 2^{-5}$. “AMG” is algebraic multigrid, “GMG” is geometric multigrid, “Inc Chol” is incomplete Cholesky. The ϵ value represents the accuracy of the FMM. We see that the FMM preconditioner has comparable convergence to the algebraic and geometric multigrid method. Even for a very low-accuracy FMM with $\epsilon = 10^{-2}$, the convergence rate is much better than the incomplete Cholesky. We refer to the work by Ibeid et al. [67] for more detailed comparisons between FMM and Multigrid.

A similar plot is shown for the Helmholtz equation with grid spacing of $h = 2^{-5}$ and wave number $\kappa = 7$ in Fig. 6. The nomenclature of the legend is identical to that of Fig. 5. In this case, we see a larger difference between the convergence rate of FMM and Multigrid. Even the FMM with the worst accuracy does better than the multigrid. We have also confirmed that the FMM preconditioner has a convergence rate that is independent of the problem size, up to moderate wave numbers of κ .

The strong scaling of FMM and AMG are shown in Fig. 7, which includes the setup phase and all iterations it took to converge. 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




Fig. 5 Convergence rate of the FMM and Multigrid preconditioners for the Laplace equation on a $[-1, 1]^3$ lattice with spacing $h = 2^{-5}$




Fig. 6 Convergence rate of the FMM and Multigrid preconditioners for the Helmholtz equation on a $[-1, 1]^3$ lattice with spacing $h = 2^{-5}$ and wave number $\kappa = 7$




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

```
-download-f-blas-lapack -download-hypr
-downloa-d-metis -downloa-d-parmetis
```

-download-superlu_dist -with-debugging=0". For BoomerAMG we compared different relaxation, coarsening, and interpolation methods and found that

```
"-pc_hypr_boomeramg_relax_type_all
backward-SOR/Jacobi
-pc_hypr_boomeramg_coarsen_type
modifiedRuge-Stueben
```

-pc_hypr_bommeramg_interp_type classical" gives the best performance. We use a grid size of $N = 4096^2$ and run from 1 to 1024 cores using up to 16 cores per node on Stampede. For this particular Poisson problem on this particular machine using this particular FMM code we see an advantage over BoomerAMG past 512 cores.

5 Conclusions and Outlook

We have shown the contrast between the analytical and algebraic hierarchical low-rank approximations, by reviewing the contributions over the years and placing them along the analytical-algebraic spectrum. The relation between Treecode, FMM, KIFMM, black-box FMM, \mathcal{H} -matrix, \mathcal{H}^2 -matrix, HODLR, HSS, HBS, and BLR were explained from the perspective of compute-memory tradeoff. This birds-eye

view of the entire hierarchical low-rank approximation landscape from analytical to algebraic, allows us to place ideas like precomputation of FMM translation matrices and relate that to storage reduction techniques for the algebraic variants.

Some important findings from this cross-disciplinary literature review are:

- Translational invariance of the FMM operators suggest that \mathcal{H}^2 -matrices (and the like) have mostly duplicate entries, which many are redundantly storing at the moment.
- The analytical variants can now perform factorization and are kernel independent, so the decision to use the algebraic variants at the cost of consuming more memory should be made carefully.
- The kernel-independent variants of FMM can be used as a matrix-free $\mathcal{O}(N)$ compression technique.
- The use of SVD to compress the FMM translation matrices, makes the work on variable expansion order and its error optimized variants redundant.
- The hierarchical compression should not be applied directly to the inverse or factorizations of sparse matrices just because they fill-in. One must first try to minimize fill-in, and then compress only the dense blocks that cannot be avoided.

The comparison benchmarks between FMM and HSS are still preliminary tests for a very simple case. However, they clearly demonstrate the magnitude of the difference that lies between the various hierarchical low-rank approximation methods. The comparison between FMM and multigrid is also a very simple test case, but it reveals the previously unquantified convergence properties of low-accuracy FMM as a preconditioner. Of course, for such simple problems the FMM can give the exact solution in finite arithmetic and therefore solve the problem in a single iteration. The interesting point here is not the fact that it can be used as a preconditioner, but the practical performance of the low-accuracy FMM being significantly faster than the high accuracy FMM, even if it requires a few iterations.

There is much more that can be done if all of these complicated hierarchical low-rank approximation methods could somehow be made easier to code. We believe a modular view of these methods will help the developers through separation of concerns. Instead of everyone coding a slightly different version of the whole thing, we could each choose a module to focus on that fits our research interests, and contribute to a larger and more sustainable ecosystem. A few ideas to facilitate the transition to such a community effort are:

1. Create a common benchmark (mini app) for each of the modules.
2. Gradually propagate standards in the community, starting from the major codes.
3. Develop a common interface between the hierarchical structure and inner kernels.
4. Do not try to unify code, just have a standard with a common API (like MPI).

Acknowledgements We thank François-Henry Rouet, Pieter Ghysels, and Xiaoye, S. Li for providing the STRUMPACK interface for our comparisons between FMM and HSS. This work was supported by JSPS KAKENHI Grant-in-Aid for Research Activity Start-up Grant Number 15H06196. 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.

References

1. Ambikasaran, S., Darve, E.: An $O(N \log N)$ fast direct solver for partial hierarchically semi-separable matrices. *J. Sci. Comput.* **57**, 477–501 (2013)
2. Ambikasaran, S., Darve, E.: The inverse fast multipole method. arXiv:1407.1572v1 (2014)
3. Ambikasaran, S., Li, J.-Y., Kitanidis, P.K., Darve, E.: Large-scale stochastic linear inversion using hierarchical matrices. *Comput. Geosci.* **17**(6), 913–927 (2013)
4. Amestoy, P., Ashcraft, C., Boiteau, O., Buttari, A., L'Excellent, J.-Y., Weisbecker, C.: Improving multifrontal methods by means of block low-rank representations. *SIAM J. Sci. Comput.* **37**(3), A1451–A1474 (2015)
5. Aminfar, A., Darve, E.: A fast, memory efficient and robust sparse preconditioner based on a multifrontal approach with applications to finite-element matrices. *Int. J. Numer. Methods Eng.* **107**, 520–540 (2016)
6. Aminfar, A., Ambikasaran, S., Darve, E.: A fast block low-rank dense solver with applications to finite-element matrices. *J. Comput. Phys.* **304**, 170–188 (2016)
7. Anderson, C.R.: An implementation of the fast multipole method without multipoles. *SIAM J. Sci. Stat. Comput.* **13**(4), 923–947 (1992)
8. Appel, A.W.: An efficient program for many-body simulation. *SIAM J. Sci. Stat. Comput.* **6**(1), 85–103 (1985)
9. Barba, L.A., Yokota, R.: How will the fast multipole method fare in the exascale era? *SIAM News* **46**(6), 1–3 (2013)
10. Barnes, J., Hut, P.: $O(N \log N)$ force-calculation algorithm. *Nature* **324**, 446–449 (1986)
11. Bebendorf, M.: Approximation of boundary element matrices. *Numer. Math.* **86**, 565–589 (2000)
12. Bebendorf, M.: *Hierarchical Matrices. Lecture Notes in Computational Science and Engineering*, vol. 63. Springer, Berlin (2008)
13. Bebendorf, M., Rjasanow, S.: Adaptive low-rank approximation of collocation matrices. *Computing* **70**, 1–24 (2003)
14. Bédorf, J., Gaburov, E., Fujii, M.S., Nitadori, K., Ishiyama, T., Portegies Zwart, S.: 24.77 Pflops on a gravitational tree-code to simulate the milky way galaxy with 18600 GPUs. In: *Proceedings of the 2014 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–12 (2014)
15. Berman, C.L.: Grid-multipole calculations. *SIAM J. Sci. Comput.* **16**(5), 1082–1091 (1995)
16. Börm, S.: Construction of data-sparse h^2 -matrices by hierarchical compression. *SIAM J. Sci. Comput.* **31**(3), 1820–1839 (2009)
17. Börm, S., Grasedyck, L.: Hybrid cross approximation of integral operators. *Numer. Math.* **101**, 221–249 (2005)
18. Börm, S., Grasedyck, L., Hackbusch, W.: Introduction to hierarchical matrices with applications. *Eng. Anal. Bound. Elem.* **27**, 405–422 (2003)
19. Bremer, J.: A fast direct solver for the integral equations of scattering theory on planar curves with corners. *J. Comput. Phys.* **231**, 1879–1899 (2012)
20. Brunner, D., Junge, M., Rapp, P., Bebendorf, M., Gaul, L.: Comparison of the Fast Multipole Method with Hierarchical Matrices for the Helmholtz-BEM. *Comput. Model. Eng. Sci.* **58**(2), 131–160 (2010)
21. Burant, J.C., Strain, M.C., Scuseria, G.E., Frisch, M.J.: Analytic energy gradients for the Gaussian very fast multipole method (GvFMM). *Chem. Phys. Lett.* **248**, 43–49 (1996)
22. Chaillat, S., Bonnet, M., Semblat, J.-F.: A multi-level fast multipole BEM for 3-D elastodynamics in the frequency domain. *Comput. Methods Appl. Mech. Eng.* **197**, 4233–4249 (2008)

23. Chan, T.F.: On the existence and computation of LU-factorizations with small pivots. *Math. Comput.* **42**(166), 535–547 (1984)
24. Chan, T.F.: Rank revealing QR factorizations. *Linear Algebra Appl.* **88/89**, 67–82 (1987)
25. Chandrasekaran, S., Ipsen, I.C.F.: On rank-revealing factorizations. *SIAM J. Matrix Anal. Appl.* **15**(2), 592–622 (1994)
26. Chandrasekaran, S., Dewilde, P., Gu, M., Lyons, W., Pals, T.: A fast solver for HSS representations via sparse matrices. *SIAM J. Matrix Anal. Appl.* **29**(1), 67–81 (2006)
27. Chandrasekaran, S., Dewilde, P., Gu, M., Somasunderam, N.: On the numerical rank of the off-diagonal blocks of Schur complements of discretized elliptic PDEs. *SIAM J. Matrix Anal. Appl.* **31**(5), 2261–2290 (2010)
28. Cheng, H., Gimbutas, Z., Martinsson, P.G., Rokhlin, V.: On the compression of low rank matrices. *SIAM J. Sci. Comput.* **26**(4), 1389–1404 (2005)
29. Choi, C.H., Ruedenberg, K., Gordon, M.S.: New parallel optimal-parameter fast multipole method (OPFMM). *J. Comput. Chem.* **22**(13), 1484–1501 (2001)
30. Corona, E., Martinsson, P.G., Zorin, D.: An $O(N)$ direct solver for integral equations on the plane. *Appl. Comput. Harmon. Anal.* **38**, 284–317 (2015)
31. Coulaud, O., Fortin, P., Roman, J.: High performance BLAS formulation of the multipole-to-local operator in the fast multipole method. *J. Comput. Phys.* **227**, 1836–1862 (2008)
32. Dachselt, H.: Corrected article: “an error-controlled fast multipole method”. *J. Chem. Phys.* **132**, 119901 (2010)
33. Darve, E., Havé, P.: A fast multipole method for Maxwell equations stable at all frequencies. *Philos. Trans. R. Soc. Lond. A* **362**, 603–628 (2004)
34. Darve, E., Cecka, C., Takahashi, T.: The fast multipole method on parallel clusters, multicore processors, and graphics processing units. *C.R. Mec.* **339**, 185–193 (2011)
35. Dehnen, W.: A hierarchical $O(N)$ force calculation algorithm. *J. Comput. Phys.* **179**(1), 27–42 (2002)
36. Dutt, A., Gu, M., Rokhlin, V.: Fast algorithms for polynomial interpolation, integration, and differentiation. *SIAM J. Numer. Anal.* **33**(5), 1689–1711 (1996)
37. Elliott, W.D., Board, J.A.: Fast Fourier transform accelerated fast multipole algorithm. *SIAM J. Sci. Comput.* **17**(2), 398–415 (1996)
38. Ethridge, F., Greengard, L.: A new fast-multipole accelerated Poisson solver in two dimensions. *SIAM J. Sci. Comput.* **23**(3), 741–760 (2001)
39. Fong, W., Darve, E.: The black-box fast multipole method. *J. Comput. Phys.* **228**, 8712–8725 (2009)
40. Fortin, P.: Multipole-to-local operator in the fast multipole method: Comparison of FFT, rotations and BLAS improvements. Technical Report RR-5752, *Rapports de recherche, et theses de l’Inria* (2005)
41. Gillman, A., Barnett, A., Martinsson, P.G.: A spectrally accurate direct solution technique for frequency-domain scattering problems with variable media. *BIT Numer. Math.* **55**, 141–170 (2015)
42. Gimbutas, Z., Greengard, L.: Fast multi-particle scattering: a hybrid solver for the Maxwell equations in microstructured materials. *J. Comput. Phys.* **232**, 22–32 (2013)
43. Gimbutas, Z., Rokhlin, V.: A generalized fast multipole method for nonoscillatory kernels. *SIAM J. Sci. Comput.* **24**(3), 796–817 (2002)
44. Goreinov, S.A., Tyrtyshnikov, E.E., Zamarashkin, N.L.: A theory of pseudoskeleton approximations. *Linear Algebra Appl.* **261**(1–3), 1–21 (1997)
45. Grasedyck, L., Hackbusch, W.: Construction and arithmetics of H-matrices. *Computing* **70**, 295–334 (2003)
46. Grasedyck, L., Kriemann, R., Le Borne, S.: Parallel black box H-LU preconditioning for elliptic boundary value problems. *Comput. Vis. Sci.* **11**, 273–291 (2008)
47. Grasedyck, L., Hackbusch, W., Kriemann, R.: Performance of H-LU preconditioning for sparse matrices. *Comput. Methods Appl. Math.* **8**(4), 336–349 (2008)
48. Grasedyck, L., Kriemann, R., Le Borne, S.: Domain decomposition based H-LU preconditioning. *Numer. Math.* **112**, 565–600 (2009)

49. Gray, A.G., Moore, A.W.: N-body problems in statistical learning. In: Leen, T.K., Dietterich, T.G., Tresp, V. (eds.) *Advances in Neural Information Processing Systems*, vol. 13, pp. 521–527. MIT Press, Cambridge (2001)
50. Greengard, L., Rokhlin, V.: A fast algorithm for particle simulations. *J. Comput. Phys.* **73**(2), 325–348 (1987)
51. Greengard, L., Rokhlin, V.: A new version of the fast multipole method for the Laplace equation in three dimensions. *Acta Numer.* **6**, 229–269 (1997)
52. 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)
53. Gu, M., Eisenstat, S.C.: Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM J. Sci. Comput.* **17**(4), 848–869 (1996)
54. Gumerov, N.A., Duraiswami, R.: Fast radial basis function interpolation via preconditioned Krylov iteration. *SIAM J. Sci. Comput.* **29**(5), 1876–1899 (2007)
55. Hackbusch, W.: A sparse matrix arithmetic based on H-matrices, part I: Introduction to H-matrices. *Computing* **62**, 89–108 (1999)
56. Hackbusch, W., Nowak, Z.P.: On the fast matrix multiplication in the boundary element method by panel clustering. *Numer. Math.* **54**, 463–491 (1989)
57. Hackbusch, W., Khoromskij, B., Sauter, S.A.: On h^2 -matrices. In: Bungartz, H., Hoppe, R., Zenger, C. (eds.) *Lectures on Applied Mathematics*. Springer, Berlin (2000)
58. Halko, N., Martinsson, P.G., Tropp, J.A.: Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.* **53**(2), 217–288 (2011)
59. Hao, S., Martinsson, P.G., Young, P.: An efficient and highly accurate solver for multi-body acoustic scattering problems involving rotationally symmetric scatterers. *Comput. Math. Appl.* **69**, 304–318 (2015)
60. Hénon, P., Saad, Y.: A parallel multistage ILU factorization based on a hierarchical graph decomposition. *SIAM J. Sci. Comput.* **28**(6), 2266–2293 (2006)
61. Hesford, A.J., Waag, R.C.: Reduced-rank approximations to the far-field transform in the gridded fast multipole method. *J. Comput. Phys.* **230**, 3656–3667 (2011)
62. 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)
63. Ho, K.L., Ying, L.: Hierarchical interpolative factorization for elliptic operators: Integral equations. *arXiv:1307.2666* (2015)
64. Hong, Y.P., Pan, C.T.: Rank-revealing QR factorizations and the singular value decomposition. *Math. Comput.* **58**(197), 213–232 (1992)
65. Hwang, T.-M., Lin, W.-W., Yang, E.K.: Rank revealing LU factorizations. *Linear Algebra Appl.* **175**, 115–141 (1992)
66. Hwang, T.-M., Lin, W.-W., Pierce, D.: Improved bound for rank revealing LU factorizations. *Linear Algebra Appl.* **261**(1), 173–186 (1997)
67. Ibeid, H., Yokota, R., Pestana, J., Keyes, D.: Fast multipole preconditioners for sparse matrices arising from elliptic equations. *arXiv:1308.3339* (2016)
68. Izadi, M.: *Hierarchical Matrix Techniques on Massively Parallel Computers*. Ph.D. thesis, Universitat Leipzig (2012)
69. Kong, W.Y., Bremer, J., Rokhlin, V.: An adaptive fast direct solver for boundary integral equations in two dimensions. *Appl. Comput. Harmon. Anal.* **31**, 346–369 (2011)
70. Langston, H., Greengard, L., Zorin, D.: A free-space adaptive FMM-based PDE solver in three dimensions. *Commun. Appl. Math. Comput. Sci.* **6**(1), 79–122 (2011)
71. Le Borne, S.: Multilevel hierarchical matrices. *SIAM J. Matrix Anal. Appl.* **28**(3), 871–889 (2006)
72. Lee, D., Vuduc, R., Gray, A.G.: A distributed kernel summation framework for general-dimension machine learning. In: *Proceedings of the 2012 SIAM International Conference on Data Mining* (2012)

73. Lessel, K., Hartman, M., Chandrasekaran, S.: A fast memory efficient construction algorithm for hierarchically semi-separable representations. <http://scg.ece.ucs.edu/publications/MemoryEfficientHSS.pdf> (2015)
74. Li, J.-Y., Ambikasaran, S., Darve, E.F., Kitanidis, P.K.: A Kalman filter powered by h^2 -matrices for quasi-continuous data assimilation problems. *Water Resour. Res.* **50**, 3734–3749 (2014)
75. Liang, Z., Gimbutas, Z., Greengard, L., Huang, J., Jiang, S.: A fast multipole method for the Rotne-Prager-Yamakawa tensor and its applications. *J. Comput. Phys.* **234**, 133–139 (2013)
76. Liberty, E., Woolfe, F., Martinsson, P.G., Rokhlin, V., Tytgert, M.: Randomized algorithms for the low-rank approximation of matrices. *Proc. Natl. Acad. Sci. U.S.A.* **104**(51), 20167–20172 (2007)
77. Makino, J.: Yet another fast multipole method without multipoles – Pseudoparticle multipole method. *J. Comput. Phys.* **151**(2), 910–920 (1999)
78. Malhotra, D., Biros, G.: PVFMM: a parallel kernel independent FMM for particle and volume potentials. *Commun. Comput. Phys.* **18**(3), 808–830 (2015)
79. Malhotra, D., Gholami, A., Biros, G.: A volume integral equation stokes solver for problems with variable coefficients. In: *Proceedings of the 2014 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–11 (2014)
80. March, W.B., Xiao, B., Biros, G.: ASKIT: approximate skeletonization kernel-independent treecode in high dimensions. *SIAM J. Sci. Comput.* **37**(2), A1089–A1110 (2015)
81. Martinsson, P.G.: The hierarchical Poincaré-Steklov (HPS) solver for elliptic PDEs: a tutorial. *arXiv:1506.01308* (2015)
82. Martinsson, P.G., Rokhlin, V.: A fast direct solver for boundary integral equations in two dimensions. *J. Comput. Phys.* **205**, 1–23 (2005)
83. Miranian, L., Gu, M.: Strong rank revealing LU factorizations. *Linear Algebra Appl.* **367**, 1–16 (2003)
84. Ohno, Y., Yokota, R., Koyama, H., Morimoto, G., Hasegawa, A., Masumoto, G., Okimoto, N., Hirano, Y., Ibeid, H., Narumi, T., Taiji, M.: Petascale molecular dynamics simulation using the fast multipole method on k computer. *Comput. Phys. Commun.* **185**, 2575–2585 (2014)
85. Oliveira, S., Yang F.: An algebraic approach for H-matrix preconditioners. *Computing* **80**, 169–188 (2007)
86. Pan, C.T.: On the existence and computation of rank-revealing LU factorizations. *Linear Algebra Appl.* **316**, 199–222 (2000)
87. Petersen, H.G., Soelvason, D., Perram, J.W., Smith, E.R.: The very fast multipole method. *J. Chem. Phys.* **101**(10), 8870–8876 (1994)
88. Rahimian, A., Lashuk, I., Veerapaneni, K., Chandramowlishwaran, A., Malhotra, D., Moon, L., Sampath, R., Shringarpure, A., Vetter, J., Vuduc, R., Zorin, D., Biros, G.: Petascale direct numerical simulation of blood flow on 200k cores and heterogeneous architectures. In: *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '10* (2010)
89. Rouet, F.-H., Li, X.-S., Ghysels, P., Napov, A.: A distributed-memory package for dense hierarchically semi-separable matrix computations using randomization. *arXiv:1503.05464* (2015)
90. Shao, Y., White, C.A., Head-Gordon, M.: Efficient evaluation of the Coulomb force in density-functional theory calculations. *J. Chem. Phys.* **114**(15), 6572–6577 (2001)
91. Takahashi, T., Cecka, C., Fong, W., Darve, E.: Optimizing the multipole-to-local operator in the fast multipole method for graphical processing units. *Int. J. Numer. Methods Eng.* **89**, 105–133 (2012)
92. Verde, A., Ghassemi, A.: Fast multipole displacement discontinuity method (FM-DDM) for geomechanics reservoir simulations. *Int. J. Numer. Anal. Methods Geomech.* **39**(18), 1953–1974 (2015)
93. Wang, Y., Wang, Q., Deng, X., Xia, Z., Yan, J., Xu, H.: Graphics processing unit (GPU) accelerated fast multipole BEM with level-skip M2L for 3D elasticity problems. *Adv. Eng. Softw.* **82**, 105–118 (2015)

94. White, C.A., Head-Gordon, M.: Rotating around the quartic angular momentum barrier in fast multipole method calculations. *J. Chem. Phys.* **105**(12), 5061–5067 (1996)
95. Wilkes, D.R., Duncan, A.J.: A low frequency elastodynamic fast multipole boundary element method in three dimensions. *Comput. Mech.* **56**, 829–848 (2015)
96. Willis, D., Peraire, J., White, J.: FastAero – a precorrected FFT-fast multipole tree steady and unsteady potential flow solver. <http://hdl.handle.net/1721.1/7378> (2005)
97. Wolf, W.R., Lele, S.K.: Aeroacoustic integrals accelerated by fast multipole method. *AIAA J.* **49**(7), 1466–1477 (2011)
98. Xia, J.: Randomized sparse direct solvers. *SIAM J. Matrix Anal. Appl.* **34**(1), 197–227 (2013)
99. Xia, J.: O(N) complexity randomized 3D direct solver with HSS2D structure. In: Proceedings of the Project Review, Geo-Mathematical Imaging Group, Purdue University, pp. 317–325 (2014)
100. Xia, J., Chandrasekaran, S., Gu, M., Li, X.S.: Superfast multifrontal method for large structured linear systems of equations. *SIAM J. Matrix Anal. Appl.* **31**(3), 1382–1411 (2009)
101. Xia, J., Chandrasekaran, S., Gu, M., Li, X.S.: Fast algorithms for hierarchically semiseparable matrices. *Numer. Linear Algebra Appl.* **17**, 953–976 (2010)
102. Yarvin, N., Rokhlin, V.: An improved fast multipole algorithm for potential fields on the line. *SIAM J. Numer. Anal.* **36**(2), 629–666 (1999)
103. Ying, L., Biros, G., Zorin, D.: A kernel-independent adaptive fast multipole algorithm in two and three dimensions. *J. Comput. Phys.* **196**(2), 591–626 (2004)
104. Ying, L., Biros, G., Zorin, D.: A high-order 3D boundary integral equation solver for elliptic PDEs in smooth domains. *J. Comput. Phys.* **219**, 247–275 (2006)
105. Yokota, R., Bardhan, J.P., Knepley, M.G., Barba, L.A., Hamada, T.: Biomolecular electrostatics using a fast multipole BEM on up to 512 GPUs and a billion unknowns. *Comput. Phys. Commun.* **182**, 1272–1283 (2011)
106. Yokota, R., Narumi, T., Yasuoka, K., Barba, L.A.: Petascale turbulence simulation using a highly parallel fast multipole method on GPUs. *Comput. Phys. Commun.* **184**, 445–455 (2013)
107. Yunis, E., Yokota, R., Ahmadi, A.: Scalable force directed graph layout algorithms using fast multipole methods. In: The 11th International Symposium on Parallel and Distributed Computing, Munich, June 2012
108. Zhao, Z., Kovvali, N., Lin, W., Ahn, C.-H., Couchman, L., Carin, L.: Volumetric fast multipole method for modeling Schrödinger's equation. *J. Comput. Phys.* **224**, 941–955 (2007)