

SMASH: Structured Matrix Approximation by Separation and Hierarchy

Difeng Cai¹, Edmond Chow², Yousef Saad³ and Yuanzhe Xi^{3*}

¹ Department of Mathematics, Purdue University, West Lafayette, USA

² School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, USA

³ Department of Computer Science & Engineering, University of Minnesota, Twin Cities, Minneapolis, MN 55455

SUMMARY

This paper presents an efficient method to perform Structured Matrix Approximation by Separation and Hierarchy (SMASH), when the original dense matrix is associated with a kernel function. Given points in a domain, a tree structure is first constructed based on an adaptive partitioning of the computational domain to facilitate subsequent approximation procedures. In contrast to existing schemes based on either analytic or purely algebraic approximations, SMASH takes advantage of both approaches and greatly improves the efficiency. The algorithm follows a bottom-up traversal of the tree and is able to perform the operations associated with each node on the same level in parallel. A strong rank-revealing factorization is applied to the initial analytic approximation in the *separation* regime so that a special structure is incorporated into the final nested bases. As a consequence, the storage is significantly reduced on one hand and a hierarchy of the original grid is constructed on the other hand. Due to this hierarchy, nested bases at upper levels can be computed in a similar way as the leaf level operations but on coarser grids. The main advantages of SMASH include its simplicity of implementation, its flexibility to construct various hierarchical rank structures and a low storage cost. Rigorous error analysis and complexity analysis are conducted to show that this scheme is fast and stable. The efficiency and robustness of SMASH are demonstrated through various test problems arising from integral equations, structured matrices, etc. Copyright © 2010 John Wiley & Sons, Ltd.

Received ...

KEY WORDS: hierarchical rank structure, nested basis, error analysis, integral equation, Cauchy-like matrix

1. INTRODUCTION

The invention of the Fast Multipole Method (FMM) [1, 2] opened a new chapter in scientific computing methodology by unraveling a set of effective techniques revolving around the powerful principle of divide-and-conquer. When sets of points are far apart from each other, the physical equations that couple them can be approximately expressed by means of a low rank matrix. Among the many variations to this elegant idea, a few schemes have been developed to gain further efficiency by building hierarchical bases in order to expand the various low-rank couplings. The resulting hierarchical rank structured matrices [3, 4, 5, 6], culminated in \mathcal{H}^2 matrices [7, 6], provide efficient solution techniques for structured linear systems (Toeplitz, Hankel, etc.) [8, 9, 10],

*Correspondence to: Yuanzhe Xi, Department of Computer Science & Engineering, University of Minnesota, Twin Cities, Minneapolis, MN 55455, U.S.A. E-mail: yxi@umn.edu

Contract/grant sponsor: Minnesota Supercomputing Institute

Contract/grant sponsor: NSF; contract/grant number: NSF/ACI-1306573, NSF/DMS-1521573

integral equations [3, 4, 11, 12, 13, 14, 15, 16], partial differential equations [4, 17, 6, 18], matrix equations [4, 19, 20] and eigenvalue problems [21, 22]. Though these methods come under various representations, they all start with a block partitioning of the coefficient matrix and approximate certain blocks with low-rank matrices. The refinements of these techniques embodied in the \mathcal{H}^2 [4, 6, 23] and HSS [5, 24, 25] matrix representations take advantage of the relationships between different (numerically) low-rank blocks and use *nested bases* [6] to minimize computational costs and storage requirements. What is often not emphasized in the literature, is that this additional efficiency in the solution phase is achieved at a rather high cost in the construction phase.

Both HSS and \mathcal{H}^2 matrices employ just two key ingredients: low-rank approximations and nested bases. The low-rank approximation, or compression, methods exploited in these techniques can be classified into three categories. The first category involves methods that rely on algebraic compression, such as the SVD and the rank-revealing QR (RRQR) [26] which are among the most common approaches. Utilizing these techniques to compress low-rank blocks [27, 23, 25] will result in nested bases that have orthogonal columns and an optimal rank. However, these methods will require the matrix entries to be explicitly available and usually lead to quadratic construction cost [23, 25]. Other compression techniques, such as adaptive cross approximation (ACA) [11, 28, 29], extract a low-rank factorization based only on part of the matrix entries and this leads to a nearly linear construction cost. However, ACA may fail for general kernel functions and complex geometries due to the heuristic nature of the method [30]. Other efficient low-rank approximation techniques include but are not limited to Pseudo-skeleton approximations [31, 32], Mosaic-skeleton approximations [33], Cross approximations [34] and their latest development [35, 36]. To the best of our knowledge, no algebraic approach is able to achieve linear cost for an \mathcal{H}^2 or HSS construction with guaranteed accuracy. The methods in the second category rely on information on the kernel to perform the compression. They include methods based on interpolation [4, 37], Taylor expansion [7] or multipole expansion (as in FMM [2, 1]), etc. Although these methods lead to a linear construction cost, they usually yield nested bases whose rank is much larger than the optimal one [38]. Moreover, since bases are stored as dense matrices, these methods suffer from high storage costs [3]. The methods in the third category either combine algebraic compression with the analytic kernel information to take advantage of both, or use other techniques like equivalent densities or randomized methods to obtain a low-rank factorization. For example, hybrid cross approximation (HCA) [30] technique improves the robustness of ACA by applying it only on a small matrix arising from the interpolation of the kernel function. The kernel independent fast multipole methods [39, 40] use equivalent densities to avoid explicit kernel expansions but it is only valid for certain kernels arising in potential theory. The randomized construction algorithms [41, 42, 43, 22] compute the hierarchical rank structured matrices by applying SVD/RRQR to the product of the original matrix and a random matrix and are effective when a fast matrix vector multiplication routine is available.

1.1. Contributions

The aim of this paper is to introduce an efficient and unified framework to construct an $n \times n$ \mathcal{H}^2 or HSS matrix based on structured matrix approximation by separation and hierarchy (SMASH). In terms of the three categories discussed above, SMASH belongs to the third category in that it starts with an initial analytic approximation in the *Separation* regime, then algebraic techniques are employed to postprocess the approximation in order to build up a *Hierarchy*. The main features of SMASH are as follows.

1. Fast and stable $O(n)$ construction SMASH starts with an adaptive partitioning of the computational domain and then constructs a tree structure to facilitate subsequent operations as in [44, 3, 4, 45]. The construction process follows a bottom-up traversal of the tree and is able to compute the bases associated with each node on the same level in parallel. In fact, the construction procedure is entirely local in the sense that the compression for a parent node only depends on the information passed from its children. By combining the analytic compression technique with strong RRQR [26], a special structure is incorporated into the final nested bases. In contrast to the

methods used in [42, 40], SMASH is able to set the approximation accuracy to any tolerance. In addition, the nested bases at each non-leaf level can be computed directly in a similar way as the leaf-level operations but on a coarser grid extracted from previous level of compression. Therefore, SMASH is also advantageous relative to one based on the HCA method [46] since it does not need to construct an \mathcal{H} matrix first and then use a recompression scheme to recover the nested bases. SMASH can be easily adapted to construct either an $n \times n$ \mathcal{H}^2 or HSS matrix depending on the properties of the underlying applications with $O(n)$ complexity. Equipped with a rigorous error analysis, the guaranteed accuracy/robustness of SMASH is justified by various test examples with complicated geometries (Section 6).

2. Low storage requirement Construction algorithms that use analytic approximations usually lead to high storage costs. SMASH alleviates this issue in several ways. First, instead of storing nested bases as dense matrices, only one vector and one smaller dense matrix need to be saved for each basis. Second, each coupling matrix [3] is a submatrix of the original matrix in this scheme. Therefore, it suffices to store row and column indices associated with the submatrix instead of the whole submatrix explicitly. Finally, the use of strong RRQR [26] can automatically reduce the rank of the nested bases if their columns obtained from the analytic approximation are not numerically linearly independent.

3. Simplicity and flexibility for approximation of variable order Unlike analytic approaches (e.g., FMM) in which farfield approximations and transfer matrices are obtained differently and extra information is needed to compute transfer matrices (cf.[47]), SMASH only requires a farfield approximation, which can be readily obtained for almost all kernels, for example, via interpolation [37]. Moreover, the approximation rank in the compression on upper levels is independent of the rank used in lower levels, which means that approximation rank can be chosen arbitrarily in the compression at any level while still maintaining the \mathcal{H}^2 or HSS structure. This is due to the fact that in each level of compression, SMASH produces transfer matrices directly, which is an advantage of algebraic approaches. For interpolation-based constructions, there are restrictions on the maximal degree of basis polynomials in each level in order to maintain the \mathcal{H}^2 structure. [23].

1.2. Outline and Notation

The paper is organized as follows. In Section 2 we review low-rank approximations ([48, 3, 6]) associated with some kernel functions. Section 3 introduces SMASH for the construction of hierarchical rank structured matrices with nested bases. The approximation error and complexity of SMASH are analyzed in Section 4 and Section 5, respectively. Numerical examples are provided in Section 6 and final concluding remarks are drawn in Section 7.

Throughout the paper, the following notation is used:

- A : a dense matrix associated with a kernel function κ ;
- \hat{A} : \mathcal{H}^2 or HSS approximation to A ;
- $i = 1 : n$ denotes an enumeration of index i from 1 to n ;
- $|\cdot|$ denotes the cardinality of a finite set if the input is a set;
- $\|\cdot\|$, $\|\cdot\|_F$ denote the L^2 norm, Frobenius norm, respectively, and $\|A\|_{\max}$ denotes the elementwise sup-norm of a matrix, i.e.,

$$\|A\|_{\max} := \max_{i,j} |a_{i,j}|, \quad A = [a_{i,j}]_{i,j};$$

- $\text{diag}(\dots)$ denotes a block diagonal matrix;
- Given a tree \mathcal{T} , $\text{children}(i)$ and $\text{lv}(i)$ represent the children and level of node i , respectively, where root node is at level 1. The location of a node i at level l is denoted as l_i when enumerated from left to right;

- Let X and Y be two nonempty finite sets of points and A be a matrix whose (i, j) th entry is determined by the i th point in X and j th point in Y . If \mathbf{i} denotes the index set corresponding to a subset $X_{\mathbf{i}}$ of X , then $A|_{\mathbf{i}}$ denotes the submatrix of A with rows determined by $X_{\mathbf{i}}$. Furthermore, if index set \mathbf{j} corresponds to a subset $Y_{\mathbf{j}}$ of Y , then $A|_{\mathbf{i} \times \mathbf{j}}$ denotes a submatrix of A whose rows and columns are determined by $X_{\mathbf{i}}$ and $Y_{\mathbf{j}}$, respectively.

2. DEGENERATE AND LOW-RANK APPROXIMATIONS

Hierarchical rank structured matrices are often used to approximate matrices after a block partitioning such that most blocks display certain (numerical) low-rank characteristics. For matrices derived from kernel functions, a low-rank approximation can be determined when the kernel function can be locally approximated by *degenerate functions* [48]. In this section, we first review this property. For pedagogical reasons, we focus on the kernel function $1/(x - y)$ but more general kernel functions can be handled in a similar way as demonstrated in the numerical experiments section (Section 6).

2.1. Degenerate expansion

Consider the kernel function $\kappa(x, y)$ on $\mathbb{C} \times \mathbb{C}$ defined by

$$\kappa(x, y) = \begin{cases} \frac{1}{x-y}, & \text{if } x \neq y, \\ d_x, & \text{if } x = y, \end{cases} \quad (1)$$

where the number $d_x \in \mathbb{C}$ can be arbitrary. If x and y are far from each other (See Definition 2.1 below), then $\kappa(x, y)$ can be well approximated by a degenerate expansion

$$\kappa(x, y) \approx \sum_{k=0}^{r-1} \sum_{l=0}^k c_{k,l} \phi_k(x) \psi_l(y),$$

where ϕ_k and ψ_l are univariate functions. In fact, interpolation in the x variable yields the simplest, yet most general, way to obtain a degenerate approximation:

$$\kappa(x, y) \approx \sum_{k=1}^r p_k(x) \kappa(x_k, y), \quad (2)$$

where x_k 's are the interpolation points and the p_k 's are the associated Lagrange polynomials.

Several ways to quantify the distance between two sets of points that are away from each other have been defined [3, 6, 47]. One of these ([47]), given below, is often referred to. For a bounded nonempty set S of \mathbb{C} , let $\delta = \min_{c \in \mathbb{C}} \sup_{s \in S} |s - c|$. Then we refer to the minimizer c_* as the center of S and to the corresponding minimum value δ as its radius.

Definition 2.1

Let X and Y be two nonempty bounded sets in \mathbb{C} . Let $a \in \mathbb{C}$ and $\delta_a > 0$ be the center and radius of X with $|x - a| \leq \delta_a, \forall x \in X$. Analogously, let $b \in \mathbb{C}$ and $\delta_b > 0$ denote the center and radius of Y . Given a number $\tau \in (0, 1)$, we say that X and Y are *well-separated with separation ratio* τ if the following condition holds

$$\delta_a + \delta_b \leq \tau |a - b|. \quad (3)$$

Fig. 1 illustrates two well-separated intervals (centered at $a = 0.5, b = 2.5$, respectively) with separation ratio $\tau = 0.5$. Given two sets X and Y , if (3) only holds for $\tau \approx 1$, then this implies that X and Y are close to each other and we cannot regard X and Y as being *well-separated*. Hence we assume that τ is a given small or moderate constant (for example, $\tau \leq 0.7$) in the rest of this paper.

Consider the kernel function $1/(x - y)$ again. When X and Y are well-separated so that (3) holds, a degenerate expansion for the kernel function based on Taylor expansion takes the following form



Figure 1. Well-separated intervals X, Y (centered at $a = 0.5, b = 2.5$) with separation ratio $\tau = 0.5$.

[49]:

$$\kappa(x, y) = \sum_{k=0}^{r-1} \sum_{l=0}^k c_{k,l} \phi_{a,l}(x-a) \phi_{b,k-l}(y-b) + \epsilon_r, \quad \forall x \in X, y \in Y, \quad (4)$$

where

$$c_{k,l} := \begin{cases} -k!(b-a)^{-(k+1)} \eta_{a,l}^{-1} \eta_{b,k-l}^{-1} (-1)^{k-l} & \text{if } l \leq k, \\ 0 & \text{if } l > k, \end{cases} \quad (5)$$

$$\phi_{v,l}(x) := \eta_{v,l} \frac{x^l}{l!}, \quad \eta_{v,l} = \begin{cases} 1, & \text{if } l = 0, \\ \left(\frac{l}{e} (2\pi r)^{\frac{1}{2r}} \frac{1}{\delta_v} \right)^l & \text{if } l = 1, \dots, r-1, \end{cases}$$

and the approximation error ϵ_r satisfies

$$|\epsilon_r| \leq \frac{(1+\tau)\tau^r}{(1-\tau)} |\kappa(x, y)|, \quad \forall x \in X, y \in Y. \quad (6)$$

The above expansion will be used to illustrate the construction of hierarchical rank structured matrices and analyze the approximation error in the remaining sections.

2.2. Farfield and nearfield blocks

We now consider a dense matrix A defined by $A := [\kappa(x, y)]_{x \in X, y \in Y}$. The degenerate approximation (4) immediately indicates that certain blocks of A admit a low-rank approximation. In order to identify these low-rank blocks, it is necessary to exploit *nearfield* and *farfield* matrix blocks as they are defined in [3].

Definition 2.2

Given two sets of points X_i and Y_j , a submatrix $A|_{i \times j}$ is called a *farfield block* if X_i and Y_j are well-separated in the sense of Definition 2.1; otherwise, $A|_{i \times j}$ is called a *nearfield block*.

A major difference between farfield and nearfield blocks is that each farfield block can be approximated by low-rank matrices, as a consequence of (4). The following theorem restates (4) in matrix form.

Theorem 2.1

If X_i and Y_j are well-separated sets in \mathbb{C} in the sense of (3) with centers a_i and a_j , radii δ_i and δ_j , respectively, the farfield block $A|_{i \times j}$ admits a low-rank approximation of the form

$$A|_{i \times j} = \hat{U}_i \hat{B}_{i,j} \hat{V}_j^T + E_F|_{i \times j}, \quad (7)$$

where

$$\hat{U}_i = [\phi_{a_i,l}(x - a_i)]_{\substack{x \in X_i, \\ l=0:r-1}}, \quad \hat{V}_j = [\phi_{a_j,l}(y - a_j)]_{\substack{y \in Y_j, \\ l=0:r-1}}, \quad \hat{B}_{i,j} = [c_{k,l}]_{k,l=0:r-1}, \quad (8)$$

with $c_{k,l}, \phi_{v,l}(v = a_i, a_j)$ defined in (5), and

$$\|E_F|_{i \times j}\|_{\max} \leq \epsilon_{\text{far}} \|A|_{i \times j}\|_{\max} \quad \text{with} \quad \epsilon_{\text{far}} = \frac{(1+\tau)\tau^r}{(1-\tau)}. \quad (9)$$

Let $n_i = |X_i|$ and $n_j = |Y_j|$. If the points x of X_i are listed as columns and the various functions $\phi_{a_i,l}(x - a_i)$ are listed row-wise with $l = 0, \dots, r-1$ and similarly for y, Y_j , and $\phi_{a_j,l}(y - a_j)$ then the matrices $\hat{U}_i, \hat{B}_{i,j}$ and \hat{V}_j has dimensions $n_i \times r, r \times r$, and $n_j \times r$, respectively. The theorem is illustrated in Fig.2.




Figure 2. Illustration of Theorem 2.1.

2.3. Strong rank-revealing QR

Notice that in the approximation (7), \hat{U}_i only depends on the points in X_i , \hat{V}_j only depends the points in Y_j and $\hat{B}_{i,j}$ depends on both the centers of X_i and Y_j as well as their radii. This represents a standard expansion structure used in FMM [2, 1, 47]. As will be seen in the next section the construction of \mathcal{H}^2 and HSS matrices will be significantly simplified by further postprocessing \hat{U}_i and \hat{V}_j with a strong rank-revealing QR (SRRQR) factorization [26]. The following theorem summarizes Algorithm 4 in [26].

Theorem 2.2

([26, Algorithm 4]) Let M be an $m \times n$ matrix and $M \neq 0$. Given a real number $s > 1$ and a positive integer r ($r \leq \text{rank}(A)$), the strong rank-revealing QR algorithm computes a factorization of M in the form:

$$MP = Q \begin{bmatrix} R_{11} & R_{12} \\ & R_{22} \end{bmatrix}, \quad (10)$$

where P is a permutation matrix, $Q \in \mathbb{R}^{m \times m}$ is an orthogonal matrix, R_{11} is a $r \times r$ upper triangular matrix and R_{12} is a $r \times (n - r)$ dense matrix that satisfies the condition:

$$\|R_{11}^{-1}R_{12}\|_{\max} \leq s.$$

The complexity is $O(n^3 \log_s n)$ if $m \approx n$.

In all of our implementations, we set $s = 2$. SRRQR unravels a set of columns of A that nearly span the range of A – thus the term rank-revealing. Assume C is an $n \times r$ matrix with rank r . Applying SRRQR to C^T produces the following factorization:

$$C^T P = Q \begin{bmatrix} R_{11} & R_{12} \end{bmatrix},$$

where $Q \in \mathbb{R}^{r \times r}$ is an orthogonal matrix. A modification of the above equation leads to

$$C = P \begin{bmatrix} I \\ (R_{11}^{-1}R_{12})^T \end{bmatrix} \tilde{C},$$

where I is an identity matrix of order r and $\tilde{C} = (QR_{11})^T$. Note that the above relation implies that $\tilde{C} \in \mathbb{R}^{r \times r}$ is a submatrix of C consisting of the first r rows of the row-permuted matrix $P^T C$. From this perspective, the whole aim of the procedure is to extract a set of r rows from C that will nearly span its row space.

When \hat{U}_i and \hat{V}_j in (7) both have more rows than columns, applying SRRQR to \hat{U}_i^T and then to \hat{V}_j^T yields:

$$\hat{U}_i = P_i \begin{bmatrix} I \\ G_i \end{bmatrix} \hat{U}_i|_i, \quad \hat{V}_j = F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \hat{V}_j|_j. \quad (11)$$

Note that, as explained above for \tilde{C} , $\hat{U}_i|_{\hat{i}}$ denotes a matrix made up of selected rows of \hat{U}_i .

Substituting the above two equations into (7) leads to another form of the low-rank approximation to $A|_{i \times j}$:

$$A|_{i \times j} \approx P_i \begin{bmatrix} I \\ G_i \end{bmatrix} \hat{U}_i|_{\hat{i}} \hat{B}_{i,j} (\hat{V}_j|_{\hat{j}})^T \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T \quad (12)$$

$$= P_i \begin{bmatrix} I \\ G_i \end{bmatrix} (A|_{i \times \hat{j}} - E_{F|_{i \times \hat{j}}}) \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T \approx P_i \begin{bmatrix} I \\ G_i \end{bmatrix} A|_{i \times \hat{j}} \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T, \quad (13)$$

where \hat{i} and \hat{j} represent subsets of i and j , respectively, and (13) results from (7).

A major advantage of this form of approximation over (7) is a reduction in storage. Now only four index sets are needed to represent $(P_i, F_j, \hat{i}, \hat{j})$ and two smaller dense matrices (G_i, H_j) need to be stored rather than three dense matrices $(\hat{U}_i, \hat{B}_{i,j}, \hat{V}_j)$. There are other advantages that will be discussed in the next section.

The operations represented by (11) will be used extensively in the construction of hierarchical matrices to be seen in the next section. These will be denoted as follows:

$$[P_i, G_i, \hat{i}] = \text{compr}(\hat{U}_i, i) \quad \text{and} \quad [F_j, H_j, \hat{j}] = \text{compr}(\hat{V}_j, j). \quad (14)$$

Each of the above operations is also called a structure-preserving rank-revealing (SPRR) factorization [10] or an interpolative decomposition [50]. For recent developments on rank-revealing algorithms, see [51]. Notice that the matrices \hat{U}_i and \hat{V}_j in (7) serve as the approximate column and row bases for $A|_{i \times j}$, respectively. Taylor expansion (4) is used to illustrate the compression of low-rank blocks due to its simplicity. More advanced compression schemes such as weighted truncation techniques [52], the modified ACA method [29] and the fast algorithm combining Green's representation formula with quadrature [53], can also be exploited to compute these bases.

3. CONSTRUCTION OF HIERARCHICAL RANK STRUCTURED MATRICES WITH NESTED BASES

This section presents SMASH, an algorithm to construct either an \mathcal{H}^2 or an HSS matrix approximation to an $n \times n$ matrix $A := [\kappa(x, y)]_{x \in X, y \in Y}$, where κ is a given kernel function and X and Y are two finite sets of points. Although the discussion focuses on square matrices, SMASH can be extended to rectangular ones [8] without any difficulty.


3.1. Adaptive partitioning

SMASH starts with a hierarchical partitioning of the computational domain Ω and then builds a tree structure \mathcal{T} to facilitate subsequent operations. In order to deal with the case when X and Y are non-uniformly distributed, an adaptive partitioning scheme is necessary.

Without loss of generality, assume both X and Y are contained in a unit box $\Omega = [0, 1]^d$ in \mathbb{R}^d ($d = 1, 2, 3$). The basic idea of this partitioning algorithm is to recursively subdivide the computational domain Ω into several subdomains until the number of points included in each resulting subdomain is less than a prescribed constant ν_0 (usually much smaller than the number of points in the domain). Specifically, at level 1, Ω is not partitioned. Starting from level l ($l \geq 2$), each box obtained at level $l - 1$ that contains more than ν_0 points is bisected along each of the d dimensions.

For convenience we assume that the number of points from X and Y in each partitioning is the same. If a box is empty, it is discarded. Let L be the maximum level where the recursion stops. Then the information about the partitioning can be represented by a tree \mathcal{T} with L levels, where the root node is at level 1 and corresponds to the domain Ω and each nonroot node i corresponds to a partitioned subdomain Ω_i . See Fig. 3 for a 1D example. The adaptive partitioning guarantees that each subdomain corresponding to a leaf node contains a small number of points less than the prescribed constant ν_0 .

Figure 3. Illustration of an adaptive partitioning for the case $X = Y = \{x_1, x_2, \dots, x_8\}$. Left: the computational domain Ω is recursively bisected until the number of points in each sub-interval Ω_i centered at a_i is less than 4 (circled dots represent the points x_i). Right: the corresponding postordered binary tree \mathcal{T} with indices of points stored at each node.



3.2. Review of \mathcal{H}^2 and HSS matrices

The low-rank property of a block $A|_{i \times j}$ associated with a node pair (i, j) is related to the *strong (or standard) admissibility condition* employed to define \mathcal{H} and \mathcal{H}^2 matrices ([4],[38]):

for a fixed $\tau \in (0, 1)$, the node pair (i, j) in \mathcal{T} is admissible if X_i and Y_j are well-separated in the sense of Definition 2.1.

Hierarchical matrices are often defined in terms of the above condition, which, in essence, spells out when a given block in the matrix can be compressed. A matrix \hat{A} (associated with a tree \mathcal{T}) is called an \mathcal{H} matrix ([37]) of rank r if there exists a positive integer r such that

$$\text{rank}(\hat{A}|_{i \times j}) \leq r, \quad \text{whenever } (i, j) \text{ is admissible.}$$

Furthermore, \hat{A} is called a *uniform \mathcal{H} matrix* ([37]) if there exist a column basis set $\{U_i\}_{i \in \mathcal{T}}$ and a row basis set $\{V_i\}_{i \in \mathcal{T}}$ associated with \mathcal{T} , such that when (i, j) is admissible, $\hat{A}|_{i \times j}$ admits a low-rank factorization:

$$\hat{A}|_{i \times j} = U_i B_{i,j} V_j^T, \quad \text{for some matrix } B_{i,j}.$$

This factorization is referred to as an *AKB representation* in [6], where $B_{i,j}$ is called a *coefficient matrix*. In [3], $B_{i,j}$ is termed a *coupling matrix* and we will follow this terminology here.

The class of \mathcal{H}^2 matrices [37] is a subset of uniform \mathcal{H} matrices with a more refined structure. That is, \hat{A} is an \mathcal{H}^2 matrix if it is a uniform \mathcal{H} matrix with nested bases in the sense that one can readily express a basis at one level from that of its children (see (16)). What is exploited here is that admissible blocks are low-rank and in addition their factors (or generators) can be expressed from lower levels.




Figure 4. A parent node p with children c_1, \dots, c_4 and the corresponding partition of the matrix block $A|_{p \times \mathbf{Q}}$, where \mathbf{Q} is the collection of indices associated with all nodes q such that (p, q) is admissible. In the context of HSS matrices there are at most 2 children since the trees are binary. For \mathcal{H}^2 matrices the trees are more general.

Assume we have a situation illustrated in Fig. 4 where the parent node p has children c_1, \dots, c_4 . According to the interpolation in (2), the column basis U_i associated with the set \mathbf{i} (for any nonroot

node i) can be chosen as

$$U_i = \left[p_k^{(i)}(x) \right]_{\substack{x \in X_i \\ k=1:r}},$$

where $p_k^{(i)}(k = 1, \dots, r)$ are Lagrange basis polynomials corresponding to interpolation points x_{i_1}, \dots, x_{i_r} . If i is a child of node p , we can write (cf.[37])

$$p_k^{(p)}(x) = \sum_{l=1:r} p_k^{(p)}(x_{i_l}) p_l^{(i)}(x). \quad (15)$$

The matrix version of (15) then leads to the so-called *nested basis*:

$$U_p = \begin{bmatrix} U_{c_1} R_{c_1} \\ \vdots \\ U_{c_4} R_{c_4} \end{bmatrix}, \quad \text{with} \quad R_i = \begin{bmatrix} p_1^{(p)}(x_{i_1}) & \dots & p_r^{(p)}(x_{i_1}) \\ \vdots & \ddots & \vdots \\ p_1^{(p)}(x_{i_r}) & \dots & p_r^{(p)}(x_{i_r}) \end{bmatrix}.$$

The nested basis can also be obtained through algebraic compressions based on a bottom-up procedure. Let $A|_{\mathbf{p} \times \mathbf{Q}}$ denote the entire (numerically) low rank block row associated with node p , i.e., \mathbf{Q} is the union of all indices \mathbf{q} such that (p, \mathbf{q}) is admissible. As illustrated in Fig. 4, assuming that the column basis U_{c_i} has been obtained from a rank- r factorization of the submatrix $A|_{c_i \times \mathbf{Q}} \approx U_{c_i} \tilde{T}_{c_i}$, we then derive

$$A|_{\mathbf{p} \times \mathbf{Q}} \approx \begin{bmatrix} U_{c_1} & & & \\ & U_{c_2} & & \\ & & U_{c_3} & \\ & & & U_{c_4} \end{bmatrix} \begin{bmatrix} \tilde{T}_{c_1} \\ \tilde{T}_{c_2} \\ \tilde{T}_{c_3} \\ \tilde{T}_{c_4} \end{bmatrix}.$$

Applying a rank- r factorization to the transpose of $[\tilde{T}_{c_1}^T \quad \tilde{T}_{c_2}^T \quad \tilde{T}_{c_3}^T \quad \tilde{T}_{c_4}^T]$ leads to

$$\begin{bmatrix} \tilde{T}_{c_1} \\ \tilde{T}_{c_2} \\ \tilde{T}_{c_3} \\ \tilde{T}_{c_4} \end{bmatrix} \approx \begin{bmatrix} R_{c_1} \\ R_{c_2} \\ R_{c_3} \\ R_{c_4} \end{bmatrix} T_p \longrightarrow A|_{\mathbf{p} \times \mathbf{Q}} \approx U_p T_p \quad \text{with} \quad U_p = \begin{bmatrix} U_{c_1} R_{c_1} \\ U_{c_2} R_{c_2} \\ U_{c_3} R_{c_3} \\ U_{c_4} R_{c_4} \end{bmatrix}.$$

Thus, we can get the basis U_p for the parent node from the children's U_{c_i} 's and the matrices R_{c_i} from both analytic and algebraic compression schemes. The R_i 's are called transfer matrices. Clearly, a similar process can be applied to obtain a row-basis V_p and so, more generally, we can write

$$U_p = \begin{bmatrix} U_{c_1} R_{c_1} \\ \vdots \\ U_{c_k} R_{c_k} \end{bmatrix}, \quad V_p = \begin{bmatrix} V_{c_1} W_{c_1} \\ \vdots \\ V_{c_k} W_{c_k} \end{bmatrix}. \quad (16)$$

Hence only the matrices U_i, V_i for all leaf nodes i need to be stored. Matrices U_p, V_p for a non-leaf node p can be obtained via transfer matrices which require much less storage. This is at the origin of the improvement from an $O(n \log n)$ cost for the early method in this class developed by Barnes and Hut [44] (\mathcal{H} structure) into an $O(n)$ cost method obtained later by the FMM [2] (\mathcal{H}^2 structure) for computing matrix-vector multiplications for some kernel matrices ([47]).

Note that as they are described in the literature \mathcal{H} and \mathcal{H}^2 matrices are associated with more general trees than those traditionally used for HSS matrices [54, 5] which are binary trees, according to the partitioning algorithm described in Section 3.1. In fact HSS matrices can be viewed as a special class of \mathcal{H}^2 matrices in which the strong admissibility condition is replaced by the *weak admissibility condition*[38]:

the node pair (i, j) in \mathcal{T} is admissible if $i \neq j$.

The above weak admissibility condition implies that, if \hat{A} is an HSS matrix, and i, j are two children of the root node, then the matrix block $\hat{A}|_{i \times j}$ should admit a low-rank factorization.

In the context of integral equations, this requirement means that the HSS structure will face difficulties in situations when couplings between nearfield blocks require a relatively high rank representation. Approximation by HSS matrices will work well for integral equations defined on a curve where kernel functions are discretized. In other cases the numerical rank of $A|_{i \times j}$ may not necessarily be small even when a non-oscillatory kernel function is discretized on a surface or in a volume [3, 9].

The construction of \mathcal{H}^2 and HSS matrices involves computing the basis matrices U, V at the leaf level, along with the transfer matrices R, W , and the coupling matrices B associated with a tree \mathcal{T} . In particular, each leaf node i is assigned four matrices $\{U_i, V_i, R_i, W_i\}$ and each nonleaf node i at level $l \geq 3$ is assigned two matrices $\{R_i, W_i\}$.

There are two types of $B_{i,j}$ matrices, those corresponding to the nearfield blocks at the leaf level, and those corresponding to the coupling matrix where the product $U_i B_{i,j} V_j^T$ approximates block $A|_{i \times j}$ for *certain* admissible (i, j) . In general, the computation of $B_{i,j}$ is more complicated because one has to carefully specify the set of admissible node pairs (i, j) to be used for the efficient approximation of A . If the distribution of points is uniform, the corresponding node pairs (i, j) are related to what is called *interaction list* in FMM [2, 47]. In more general settings where points can be arbitrarily distributed, they are called *admissible leaves* [3]. The set of admissible leaves corresponding to the minimal admissible partitioning [6] can be defined as follows:

$$\begin{aligned} \mathcal{L} = & \{(i, j) : i, j \in \mathcal{T} \text{ are nodes at the same level such that } (i, j) \text{ is admissible} \\ & \text{but } (p_i, p_j) \text{ is not admissible, where } p_i, p_j \text{ are parents of } i, j, \text{ respectively}\} \\ & \cup \{(i, j) : i \in \mathcal{T} \text{ is a leaf node and } j \in \mathcal{T} \text{ with } \text{lv}(j) > \text{lv}(i) \text{ such that} \\ & (i, j) \text{ is admissible but } (i, p_j) \text{ is not admissible with } p_j \text{ the parent of } j\} \\ & \cup \{(i, j) : j \in \mathcal{T} \text{ is a leaf node and } i \in \mathcal{T} \text{ with } \text{lv}(i) > \text{lv}(j) \text{ such that} \\ & (i, j) \text{ is admissible but } (p_i, j) \text{ is not admissible with } p_i \text{ the parent of } i\}. \end{aligned} \quad (17)$$

The node pairs (i, j) corresponding to blocks $B_{i,j}$ that can not be compressed or partitioned, can be identified through inadmissible leaves as defined below (cf.[3]):

$$\mathcal{L}^- := \{(i, j) : i, j \in \mathcal{T} \text{ are leaf nodes and } (i, j) \text{ is not admissible}\}. \quad (18)$$

In particular, for HSS matrices, it can be seen that \mathcal{L} and \mathcal{L}^- have the following simple expression:

$$\mathcal{L} = \{(i, j) : i, j \in \mathcal{T} \text{ and } j = \text{sibling of } i\}, \quad \mathcal{L}^- = \{(i, i) : i \in \mathcal{T} \text{ is a leaf node}\}. \quad (19)$$

This special feature will be used later (Section 3.3.1) to simplify the notation associated with HSS matrices. The U, V, R, W, B matrices are called \mathcal{H}^2 , or HSS, *generators* in the remaining sections.

3.3. Levelwise construction

Although the HSS structure may appear to be simpler than the \mathcal{H}^2 structure, based on their algebraic definitions the HSS construction procedure is actually more complex. This is because HSS matrices require the compression of both nearfield and farfield blocks while \mathcal{H}^2 matrices only require the compression on farfield blocks. For example, if two sets X_i and Y_j are almost adjacent to each other ($\tau \approx 1$ in (3)), then the analytic approximation will not produce a *low rank*, i.e., to get an accurate approximation, r has to be large in (6). In this case, the \mathcal{H}^2 matrix will form this block explicitly as a dense matrix while the HSS matrix still requires the block to be factorized. In what follows, we will first discuss SMASH for the HSS construction in detail and then present the \mathcal{H}^2 construction with an emphasis on their differences.

3.3.1. HSS construction Due to the simple structure of $\mathcal{L}, \mathcal{L}^-$ in (19), the notation denoting the coupling matrices and nearfield blocks can be simplified in the HSS representation. Specifically, for

$(i, j) \in \mathcal{L}$, $B_{i,j}$ can be represented as B_i with the second index j dropped because $j = \text{sibling of } i$ is uniquely determined in a binary tree. An additional symbol D_i is introduced to represent diagonal blocks $B_{i,i}$ because $(i, j) \in \mathcal{L}^-$ implies $j = i$.

The basic idea of SMASH for the HSS construction is to first apply a truncated SVD to obtain a basis for nearfield blocks, use interpolation or expansion to obtain a basis for farfield blocks and then apply SRRQR to the combination of these two bases to obtain the U, V, R and W matrices. The D and B matrices are submatrices of the original kernel matrix and their indices are readily available after the computation of U, V, R, W matrices. In order to distinguish between column and row indices associated with a node i , we use superscript *row* to mark its row indices and *col* to mark its column indices. For example, \mathbf{i}^{row} and \mathbf{i}^{col} denote the indices of points from X and Y contained in Ω_i , respectively.

Assuming the HSS tree \mathcal{T} has L levels, the HSS construction algorithm traverses \mathcal{T} through level $l = L, L-1, \dots, 2$. Before the construction, two intermediate sets $\bar{\mathbf{i}}^{row}$ and $\bar{\mathbf{i}}^{col}$ are initialized as follows for each node i :

$$\bar{\mathbf{i}}^{row} = \begin{cases} \mathbf{i}^{row} & \text{if } i \text{ is a leaf,} \\ \emptyset & \text{otherwise,} \end{cases} \quad \bar{\mathbf{i}}^{col} = \begin{cases} \mathbf{i}^{col} & \text{if } i \text{ is a leaf,} \\ \emptyset & \text{otherwise,} \end{cases} \quad (20)$$

where the index sets \mathbf{i}^{row} and \mathbf{i}^{col} have been saved for each leaf node after the partitioning of Ω .

Let \mathcal{N}_i be the set of blocks that are nearfield to node i . This set is used for HSS matrices only and is defined in the appendix. For each node i at level l , the construction algorithm first applies a truncated SVD to compute an approximate column basis for the nearfield block rows in terms of $X_{\bar{\mathbf{i}}^{row}}$:

$$A_i^- := [A]_{\bar{\mathbf{i}}^{row} \times \bar{\mathbf{j}}^{col}}]_{j \in \mathcal{N}_i} = S_i \Sigma_i^- [\tilde{S}_j]_{j \in \mathcal{N}_i} + [E_\Sigma^-]_{\bar{\mathbf{i}}^{row} \times \bar{\mathbf{j}}^{col}}]_{j \in \mathcal{N}_i}, \quad (21)$$

where the columns of S_i and $[\tilde{S}_j]_{j \in \mathcal{N}_i}$ are the left/right singular vectors of A_i^- and Σ_i^- is a diagonal matrix composed of corresponding singular values of A_i^- such that the following estimate holds

$$\|E_\Sigma^-|_{\bar{\mathbf{i}}^{row} \times \bar{\mathbf{j}}^{col}}\|_F \leq \sqrt{|\bar{\mathbf{i}}^{row}| |\bar{\mathbf{j}}^{col}|} \epsilon_{\text{SVD}} \|A_i^-\|_2, \quad \forall j \in \mathcal{N}_i. \quad (22)$$

Here, ϵ_{SVD} is the relative tolerance used in the truncated SVD. This estimate will be used to analyze the overall construction error in Section 4. The matrix S_i is then taken as an approximate column basis for the nearfield block rows A_i^- .

For farfield blocks with respect to $X_{\bar{\mathbf{i}}^{row}}$, a column basis \hat{U}_i can be easily obtained through interpolation (2) or Taylor expansion (8) that only rely on $X_{\bar{\mathbf{i}}^{row}}$ and Ω_i . Next, we apply SRRQR to the combined basis $[\hat{U}_i, S_i]$ as shown below

$$[P_i, G_i, \hat{\mathbf{i}}^{row}] = \text{compr}([\hat{U}_i, S_i], \bar{\mathbf{i}}^{row}). \quad (23)$$

From these outputs, we set

$$U_i := P_i \begin{bmatrix} I \\ G_i \end{bmatrix} \quad \text{if } i \text{ is a leaf node,} \\ \begin{bmatrix} R_{c_1} \\ R_{c_2} \end{bmatrix} := P_i \begin{bmatrix} I \\ G_i \end{bmatrix} \quad \text{if } i \text{ is a parent with children } c_1, c_2. \quad (24)$$

Similarly, in order to compute V, W generators, a truncated SVD is first applied to the nearfield block columns (transposed) in terms of $Y_{\bar{\mathbf{i}}^{col}}$:

$$A_i^| := \left[(A_{\bar{\mathbf{j}}^{row} \times \bar{\mathbf{i}}^{col}})^T \right]_{j \in \mathcal{N}_i} = T_i \Sigma_i^| [\tilde{T}_j]_{j \in \mathcal{N}_i} + \left[(E_\Sigma^|)_{\bar{\mathbf{j}}^{row} \times \bar{\mathbf{i}}^{col}} \right]_{j \in \mathcal{N}_i}^T, \quad (25)$$

where the truncation error satisfies

$$\|E_\Sigma^|]_{\bar{\mathbf{j}}^{row} \times \bar{\mathbf{i}}^{col}}\|_F \leq \sqrt{|\bar{\mathbf{i}}^{col}| |\bar{\mathbf{j}}^{row}|} \epsilon_{\text{SVD}} \|A_i^|\|_2, \quad \forall j \in \mathcal{N}_i. \quad (26)$$

In the next step we compute a row basis \hat{V}_i for the farfield blocks with respect to $Y_{\hat{\mathbf{i}}^{col}}$ based on (2) or (8) and apply SRRQR to $[\hat{V}_i, T_i]$:

$$[F_i, H_i, \hat{\mathbf{i}}^{col}] = \text{compr}([\hat{V}_i, T_i], \bar{\mathbf{i}}^{col}).$$

Then we set

$$\begin{aligned} V_i &:= F_i \begin{bmatrix} I \\ H_i \end{bmatrix} && \text{if } i \text{ is a leaf node,} \\ \begin{bmatrix} W_{c_1} \\ W_{c_2} \end{bmatrix} &:= F_i \begin{bmatrix} I \\ H_i \end{bmatrix} && \text{if } i \text{ is a parent with children } c_1, c_2. \end{aligned} \quad (27)$$

Once the compressions for children nodes (at level l) are complete, we update the intermediate index set associated with the parent node (at level $l - 1$) as shown below :

$$\bar{\mathbf{p}}^{row} = \hat{\mathbf{c}}_1^{row} \cup \hat{\mathbf{c}}_2^{row}, \quad \bar{\mathbf{p}}^{col} = \hat{\mathbf{c}}_1^{col} \cup \hat{\mathbf{c}}_2^{col}, \quad (28)$$

where c_1, c_2 are the children of p .

After the bottom-up traversal of \mathcal{T} and hence the computation of U, V, R, W matrices, the B and D matrices can be extracted as follows:

$$B_i := A|_{\hat{\mathbf{i}}^{row} \times \hat{\mathbf{j}}^{col}}, \quad j = \text{sibling of } i, \quad \text{and} \quad D_i := A|_{\hat{\mathbf{i}}^{row} \times \hat{\mathbf{i}}^{col}}, \quad i = \text{leaf node.} \quad (29)$$

3.3.2. \mathcal{H}^2 construction As mentioned at the beginning of Section 3.3, the \mathcal{H}^2 construction is simpler because nearfield blocks will not be factorized, and the only complication is that an \mathcal{H}^2 matrix may be associated with a more general tree structure where a parent can have more than two children.

SMASH for the \mathcal{H}^2 construction also follows a bottom-up levelwise traversal of \mathcal{T} through level $l = L, L - 1, \dots, 2$. For each node i at level l , a column/row basis \hat{U}_i/\hat{V}_i corresponding to a farfield block row/column with index $\bar{\mathbf{i}}^{row}/\bar{\mathbf{i}}^{col}$ can be obtained via either interpolation (2) or Taylor expansion (8). The bases \hat{U}_i and \hat{V}_i are then passed into SRRQR

$$[P_i, G_i, \hat{\mathbf{i}}^{row}] = \text{compr}(\hat{U}_i, \bar{\mathbf{i}}^{row}) \quad \text{and} \quad [F_i, H_i, \hat{\mathbf{i}}^{col}] = \text{compr}(\hat{V}_i, \bar{\mathbf{i}}^{col}). \quad (30)$$

The \mathcal{H}^2 generators U, R, V, W are then set as

$$\begin{aligned} U_i &:= P_i \begin{bmatrix} I \\ G_i \end{bmatrix} && \text{if } i \text{ is a leaf node,} \\ \begin{bmatrix} R_{c_1} \\ \vdots \\ R_{c_k} \end{bmatrix} &:= P_i \begin{bmatrix} I \\ G_i \end{bmatrix} && \text{if } i \text{ is a parent with children } c_1, \dots, c_k, \\ V_i &:= F_i \begin{bmatrix} I \\ H_i \end{bmatrix} && \text{if } i \text{ is a leaf node,} \\ \begin{bmatrix} W_{c_1} \\ \vdots \\ W_{c_k} \end{bmatrix} &:= F_i \begin{bmatrix} I \\ H_i \end{bmatrix} && \text{if } i \text{ is a parent with children } c_1, \dots, c_k. \end{aligned} \quad (31)$$

Again, once the compressions for children nodes (at level l) are complete, the intermediate index set associated with the parent node (at level $l - 1$) can be updated as in (28). Namely,

$$\bar{\mathbf{p}}^{row} = \hat{\mathbf{c}}_1^{row} \cup \dots \cup \hat{\mathbf{c}}_k^{row}, \quad \bar{\mathbf{p}}^{col} = \hat{\mathbf{c}}_1^{col} \cup \dots \cup \hat{\mathbf{c}}_k^{col}, \quad \text{with } \text{children}(p) = \{c_1, \dots, c_k\}. \quad (32)$$

Finally, analogous to (29), the coupling matrices associated with admissible leaves are extracted based on index sets $\hat{\mathbf{i}}^{row}$ and $\hat{\mathbf{j}}^{col}$ as

$$B_{i,j} := A|_{\hat{\mathbf{i}}^{row} \times \hat{\mathbf{j}}^{col}}, \quad \forall (i, j) \in \mathcal{L}, \quad (33)$$

and the nearfield blocks associated with inadmissible leaves are formed by

$$B_{i,j} := A|_{i^{row} \times j^{col}}, \quad \forall (i,j) \in \mathcal{L}^-. \quad (34)$$

Compared with standard \mathcal{H}^2 constructions based on either expansion or interpolation, SMASH is more efficient and easier to implement. First, in order to complete the \mathcal{H}^2 construction procedure, it suffices to provide only the column/row basis for each farfield block, which can be easily obtained based on interpolation (2), for example, and the coupling matrices $B_{i,j}$ can be simply extracted from the original matrix without resorting to any other formulas. Second, no information is required about the translation to compute transfer matrices because the computation of R/W is essentially the same as that of U/V at leaf level. For all the children i of a node p , R_i/W_i are calculated jointly based on a subset of points located inside Ω_p (i.e., $X_{\bar{p}^{row}}/Y_{\bar{p}^{col}}$). Therefore, SMASH essentially builds a hierarchy of grids and computes the bases at each level of the tree by repeating the same operations (30) on each coarse grid (this can be more clearly seen in (41) in view of a perfect tree). In addition, the use of SRRQR guarantees that each entry of the U, V, R, W matrices is bounded by a user-specified constant, which ensures the numerical stability of the construction procedure. Note that the special structures in the nested bases (31) result in not only a reduced storage but also in faster matrix operations such as matrix-vector multiplications, linear system solutions, etc. Finally, since the computation of the basis matrices only relies on the information local to each node, as can be seen from (23) and (30), this construction algorithm is inherently suitable for a parallel computing environment.

3.4. Matrix-vector multiplication

Among various hierarchical rank structured matrix operations, the matrix-vector multiplication is the most widely used, as indicated by the popularity of tree code [44] (for \mathcal{H} matrices) and fast multipole method [2, 47] (for \mathcal{H}^2 matrices).

The matrix-vector multiplication for an \mathcal{H}^2 matrix A follows first a bottom-up and then a top-down traversal of \mathcal{T} [3, 6], which is a succinct algebraic generalization of the fast multipole method (cf.[47]). Suppose \mathcal{T} has L levels, the node-wise version of this algorithm to evaluate $z = Aq$ can be summarized as follows:

1. from level $l = L$ to level $l = 2$, for each node i at level l , compute $\hat{q}_i := V_i^T q|_{i^{col}}$ if $l = L$; otherwise, compute $\hat{q}_i := \sum_{c \in \text{children}(i)} W_c^T \hat{q}_c$;
2. for each nonroot node $i \in \mathcal{T}$, compute $\hat{z}_i = \sum_{j: (i,j) \in \mathcal{L}} B_{i,j} \hat{q}_j$;
3. from level $l = 2$ to level $l = L$, for each node i at level l , if $l < L$, for each child c of i , compute $\hat{z}_c = \hat{z}_c + R_c \hat{z}_i$; otherwise, compute $z|_{i^{row}} = U_i \hat{z}_i + \sum_{j: (i,j) \in \mathcal{L}^-} B_{i,j} q|_{j^{col}}$.

When X and Y are uniformly distributed in $[0, 1]^d$, the resulting tree \mathcal{T} is a perfect 2^d -tree (each parent has 2^d children and all leaf nodes are at the same level). If, in addition we assume the ordering of points to be consistent with the postordering of tree \mathcal{T} , i.e., for two siblings $i, j \in \mathcal{T}$, if $i < j$ then the index of any point in box i must be smaller than point in box j , then an \mathcal{H}^2 matrix A has a telescoping representation:

$$A = B^{(L)} + U^{(L)} \left(U^{(L-1)} \left(\dots \left(U^{(2)} B^{(1)} \left(V^{(2)} \right)^T + B^{(2)} \right) \dots \right) \left(V^{(L-1)} \right)^T + B^{(L-1)} \right) \left(V^{(L)} \right)^T, \quad (35)$$

where $U^{(l)}, V^{(l)}$ are block diagonal matrices:

$$U^{(l)} = \begin{cases} \text{diag}(U_i)_{|v(i)=l} & \text{if } l = L, \\ \text{diag} \left(\begin{bmatrix} R_{c_1} \\ \vdots \\ R_{c_k} \end{bmatrix} \right)_{|v(i)=l} & \text{if } l < L, \end{cases} \quad V^{(l)} = \begin{cases} \text{diag}(V_i)_{|v(i)=l} & \text{if } l = L, \\ \text{diag} \left(\begin{bmatrix} W_{c_1} \\ \vdots \\ W_{c_k} \end{bmatrix} \right)_{|v(i)=l} & \text{if } l < L, \end{cases} \quad (36)$$

and $B^{(l)}$ has a block structure. $B^{(L)}$ has $\#\{i \in \mathcal{T} : \text{lv}(i) = L\} \times \#\{i \in \mathcal{T} : \text{lv}(i) = L\}$ blocks where each nonzero block corresponds to a nearfield block, while for $l < L$, there are $\#\{i \in \mathcal{T} : \text{lv}(i) = l + 1\} \times \#\{i \in \mathcal{T} : \text{lv}(i) = l + 1\}$ blocks in $B^{(l)}$ and each nonzero block corresponds to a coupling matrix. That is, block (l_i, l_j) is equal to $B_{i,j}$ if $\text{lv}(i) = \text{lv}(j) = l$ such that $(i, j) \in \mathcal{L}^-$ when $l = L$ or $(i, j) \in \mathcal{L}$ when $l \leq L$. Here l_i denotes the location of node i at level l enumerated from left to right. If A is an HSS matrix associated with a perfect binary tree \mathcal{T} , the structures of $U^{(l)}$ and $V^{(l)}$ are identical to those in (36) with $k = 2$ but $B^{(l)}$ has a much simpler block diagonal structure:

$$B^{(l)} = \begin{cases} \text{diag}(D_i)_{i \text{ is a leaf node}} & \text{if } l = L, \\ \text{diag} \left(\begin{bmatrix} 0 & B_{c_1} \\ B_{c_2} & 0 \end{bmatrix} \right)_{\text{lv}(i)=l} & \text{if } l < L, \end{cases} \quad (37)$$

where c_1 and c_2 are the children of node i .

Based on the explicit representation (35) of an \mathcal{H}^2 matrix associated with a perfect 2^d -tree, we can write down a levelwise version of the matrix-vector multiplication:

1. at level $l(2 \leq l \leq L)$, compute

$$\hat{q}^{(l)} = \left(V^{(l)}\right)^T \dots \left(V^{(L)}\right)^T q; \quad (38)$$

2. at level $l(2 \leq l \leq L)$, compute

$$\hat{z}^{(l)} = B^{(l-1)} \hat{q}^{(l)}; \quad (39)$$

- 3.

$$z = B^{(L)} q + U^{(L)} \left(\dots \left(U^{(3)} (U^{(2)} \hat{z}^{(2)} + \hat{z}^{(3)}) + \hat{z}^{(4)} \right) + \dots \hat{z}^{(L)} \right). \quad (40)$$

Notice that when X, Y are not uniformly distributed, \mathcal{T} is not necessarily a perfect tree. Under this condition, the nodes i, j corresponding to a coupling matrix $B_{i,j}$ may not be at the same level of \mathcal{T} and the telescoping expansion (35) does not exist.

As for linear system solutions, \mathcal{H}^2 and HSS matrices take completely different approaches to directly solve the resulting system. Linear complexity \mathcal{H}^2 matrix solvers ([3, 6]) heavily depend on recursion to reduce the number of floating point operations while HSS matrices could benefit from highly parallelizable ULV-type algorithms (cf.[54], [5]) due to the special structure of HSS. However, as mentioned before, since the requirement of an HSS structure is too strong, the application of HSS matrices is limited as compared to \mathcal{H}^2 matrices.

4. ERROR ANALYSIS

In this section, we analyze the approximation error of SMASH. Since the HSS construction is more complicated than \mathcal{H}^2 construction due to the factorization of nearfield blocks, the corresponding error analysis is more involved. Here we first present error analysis for the HSS construction associated with a perfect binary tree and then an error bound for the \mathcal{H}^2 construction associated with a perfect 2^d -tree can be easily derived.

For a perfect 2^d -tree, it is natural and easy to interpret the levelwise construction in Section 3.3 in the following recursive manner:

$$A^{(l)} = U^{(l)} A^{(l-1)} \left(V^{(l)}\right)^T + B^{(l)} + E^{(l)}, \quad \forall l \geq 3, \quad (41)$$

where $A^{(L)} = A$ and $A^{(l)}$ ($l < L$) is a submatrix of A with the following block structure:

$$(l_i, l_j)\text{-block of } A^{(l)} = \begin{cases} A|_{\bar{i} \times \bar{j}} & \text{if } \text{lv}(i) = \text{lv}(j) = l \text{ and } (i, j) \text{ is admissible,} \\ 0 & \text{if } \text{lv}(i) = \text{lv}(j) = l \text{ and } (i, j) \text{ is not admissible,} \end{cases} \quad (42)$$

$U^{(l)}, V^{(l)}, B^{(l)}$ follow the definition in Section 3.4 and $E^{(l)}$ denotes the factorization error at level l . The superscripts *row* and *col* used in (20) are dropped in (42) in order to simplify the notation used in the proof. Here, we assume that the items involving the index i refer to rows and the items involving the index j refer to columns. In addition, we introduce the notation

$$\mathbf{I}^{(l)} := \cup_{\text{lv}(i)=l} \bar{\mathbf{i}} \quad \text{and} \quad \mathbf{J}^{(l)} := \cup_{\text{lv}(j)=l} \bar{\mathbf{j}}.$$

Then the size of $A^{(l)}$ is equal to $|\mathbf{I}^{(l)}| \times |\mathbf{J}^{(l)}|$.

We also assume there exists a constant $r^{(l)}$ associated with each level such that

$$|\hat{\mathbf{i}}| \leq r^{(l)}, \quad |\hat{\mathbf{j}}| \leq r^{(l)}, \quad i, j \text{ at level } l, \text{ and } r^{(l+1)} \leq r^{(l)}. \quad (43)$$

This constant $r^{(l)}$ is actually an upper bound for the numerical ranks in the admissible blocks at level l .

Expanding the recursion (41) leads to

$$\begin{aligned} A = & U^{(L)} \left(U^{(L-1)} \left(\dots \left(U^{(2)} B^{(1)} \left(V^{(2)} \right)^T + B^{(2)} + E^{(2)} \right) \dots \right) \left(V^{(L-1)} \right)^T \right. \\ & \left. + B^{(L-1)} + E^{(L-1)} \right) \left(V^{(L)} \right)^T + B^{(L)} + E^{(L)}. \end{aligned} \quad (44)$$

Now we focus on the analysis of HSS approximation error. To estimate the norm of each diagonal block in $U^{(l)}$ and $V^{(l)}$, the following lemma is needed, which is a simple consequence of SRRQR in Theorem 2.2 and thus holds for both \mathcal{H}^2 and HSS construction.

Lemma 4.1

Let $s > 1$ be the prescribed elementwise bound in SRRQR in Theorem 2.2, then the norms of the \mathcal{H}^2 generators in (31) or of the HSS generators in (24) and (27) satisfy the following estimate

$$\|P_i \begin{bmatrix} I \\ G_i \end{bmatrix}\|_F \leq s\sqrt{|\hat{\mathbf{i}}|r^{(l)}} \quad \text{and} \quad \|F_j \begin{bmatrix} I \\ H_j \end{bmatrix}\|_F \leq s\sqrt{|\hat{\mathbf{j}}|r^{(l)}}, \quad (45)$$

for any nodes i, j at level l .

Proof

Under the assumption (43), we know that the column size of $\begin{bmatrix} I \\ G_i \end{bmatrix}$ is $|\hat{\mathbf{i}}| \leq r^{(l)}$. Since the row size of $\begin{bmatrix} I \\ G_i \end{bmatrix}$ is equal to $|\bar{\mathbf{i}}|$ and the entries of G_i are bounded by s , we get

$$\|P_i \begin{bmatrix} I \\ G_i \end{bmatrix}\|_F \leq s\sqrt{|\hat{\mathbf{i}}| + (|\bar{\mathbf{i}}| - |\hat{\mathbf{i}}|)|\hat{\mathbf{i}}|} \leq s\sqrt{|\bar{\mathbf{i}}|r^{(l)}}.$$

The same argument applies to $\|F_j \begin{bmatrix} I \\ H_j \end{bmatrix}\|_F$. □

In the following two lemmas, Lemma 4.2 and Lemma 4.3, we investigate the local error generated in farfield approximation and nearfield approximation, respectively. Lemma 4.2 applies to both \mathcal{H}^2 and HSS construction, while Lemma 4.3 is only necessary for HSS construction because there is no nearfield approximation in \mathcal{H}^2 construction. Referring to Section 3.1, in the following we will denote by Sep the set of all well-separated pairs of subdomains corresponding to a given partitioning:

$$\text{Sep} := \{(i, j) : \Omega_i, \Omega_j \text{ are well-separated}\}. \quad (46)$$

Lemma 4.2

Suppose i and j are two nodes at level l and $(i, j) \in \text{Sep}$. Denote the farfield approximation tolerance by ϵ_{far} as in (9). Then the \mathcal{H}^2 generators computed in (31) or the HSS generators computed in (24)

and (27) produce the following factorization

$$A|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}} = P_i \begin{bmatrix} I \\ G_i \end{bmatrix} A|_{\hat{\mathbf{i}} \times \hat{\mathbf{j}}} \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T + E_{(i,j)},$$

where the approximation error $E_{(i,j)}$ satisfies

$$\|E_{(i,j)}\|_F \leq 2s^2 \sqrt{|\bar{\mathbf{i}}| |\bar{\mathbf{j}}|} r^{(l)} \epsilon_{\text{far}} \|A|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}}\|_F.$$

Proof

According to (7) and (11), we know that for each $(i, j) \in \text{Sep}$

$$\begin{aligned} A|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}} &= \hat{U}_{\bar{\mathbf{i}}} \hat{B}_{i,j} \hat{V}_{\bar{\mathbf{j}}}^T + E_F|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}} = P_i \begin{bmatrix} I \\ G_i \end{bmatrix} \hat{U}_{\hat{\mathbf{i}}} \hat{B}_{i,j} \hat{V}_{\hat{\mathbf{j}}}^T \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T + E_F|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}} \\ &= P_i \begin{bmatrix} I \\ G_i \end{bmatrix} A|_{\hat{\mathbf{i}} \times \hat{\mathbf{j}}} \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T + E_{(i,j)}, \end{aligned}$$

where

$$E_{(i,j)} := -P_i \begin{bmatrix} I \\ G_i \end{bmatrix} E_F|_{\hat{\mathbf{i}} \times \hat{\mathbf{j}}} \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T + E_F|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}}. \quad (47)$$

Based on (9) and Lemma 4.1, we deduce that

$$\begin{aligned} \|E_{(i,j)}\|_F &\leq \|P_i \begin{bmatrix} I \\ G_i \end{bmatrix}\|_F \|E_F|_{\hat{\mathbf{i}} \times \hat{\mathbf{j}}}\|_F \|F_j \begin{bmatrix} I \\ H_j \end{bmatrix}\|_F + \|E_F|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}}\|_F \\ &\leq 2s \sqrt{|\bar{\mathbf{i}}|} r^{(l)} \epsilon_{\text{far}} \|A|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}}\|_F s \sqrt{|\bar{\mathbf{j}}|} r^{(l)} = 2s^2 \sqrt{|\bar{\mathbf{i}}| |\bar{\mathbf{j}}|} r^{(l)} \epsilon_{\text{far}} \|A|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}}\|_F. \end{aligned}$$

□

Lemma 4.3

Suppose i and j are two nodes at level l of the HSS tree and $(i, j) \notin \text{Sep}$. Then in the HSS construction, the application of the truncated SVD with relative tolerance ϵ_{SVD} in (21) and (25) produces the following factorization

$$A|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}} = S_i \Sigma_i^- \tilde{S}_j + E_{(i,j)},$$

where

$$\|E_{(i,j)}\|_F \leq s^2 |\bar{\mathbf{j}}| \sqrt{|\bar{\mathbf{i}}|} r^{(l)} r^{(l)} \epsilon_{\text{SVD}} \|A_j\|_F + s |\bar{\mathbf{i}}| \sqrt{|\bar{\mathbf{j}}|} r^{(l)} \epsilon_{\text{SVD}} \|A_i^-\|_F.$$

Proof

According to (23), we know that

$$S_i = P_i \begin{bmatrix} I \\ G_i \end{bmatrix} S_i|_{\hat{\mathbf{i}}}.$$

Therefore,

$$\begin{aligned} S_i \Sigma_i^- \tilde{S}_j &= P_i \begin{bmatrix} I \\ G_i \end{bmatrix} S_i|_{\hat{\mathbf{i}}} \Sigma_i^- \tilde{S}_j \\ &= P_i \begin{bmatrix} I \\ G_i \end{bmatrix} \left(A|_{\hat{\mathbf{i}} \times \bar{\mathbf{j}}} - E_{\Sigma}^-|_{\hat{\mathbf{i}} \times \bar{\mathbf{j}}} \right). \end{aligned}$$

Based on (25), we further have:

$$A|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}} = \tilde{T}_i^T |_{\hat{\mathbf{i}}} (\Sigma_j^|)^T T_j^T + E_{\Sigma}^-|_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}}.$$

Since

$$T_j^T = (T_j|_{\hat{j}})^T \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T,$$

we obtain

$$A|_{\hat{i} \times \bar{j}} = \left(A|_{\hat{i} \times \hat{j}} - E_{\Sigma}^{|} |_{\hat{i} \times \hat{j}} \right) \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T + E_{\Sigma}^{|} |_{\hat{i} \times \bar{j}}.$$

Therefore, we get

$$A|_{\hat{i} \times \bar{j}} = P_i \begin{bmatrix} I \\ G_i \end{bmatrix} A|_{\hat{i} \times \hat{j}} \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T + E_{(i,j)},$$

where

$$E_{(i,j)} := -P_i \begin{bmatrix} I \\ G_i \end{bmatrix} E_{\Sigma}^{|} |_{\hat{i} \times \hat{j}} \left(F_j \begin{bmatrix} I \\ H_j \end{bmatrix} \right)^T + P_i \begin{bmatrix} I \\ G_i \end{bmatrix} \left(E_{\Sigma}^{|} |_{\hat{i} \times \bar{j}} - E_{\Sigma}^{-} |_{\hat{i} \times \bar{j}} \right) + E_{\Sigma}^{-} |_{\hat{i} \times \bar{j}}. \quad (48)$$

Introduce the notation

$$\check{i} := \bar{i} \setminus \hat{i} \quad \text{and} \quad \check{j} := \bar{j} \setminus \hat{j},$$

we then have

$$E_{\Sigma}^{|} |_{\hat{i} \times \bar{j}} = E_{\Sigma}^{|} |_{\hat{i} \times \hat{j}} \left(F_j \begin{bmatrix} I \\ 0 \end{bmatrix} \right) + E_{\Sigma}^{|} |_{\hat{i} \times \check{j}} \left(F_j \begin{bmatrix} 0 \\ I \end{bmatrix} \right) \quad \text{and} \quad E_{\Sigma}^{-} |_{\hat{i} \times \bar{j}} = P_i \begin{bmatrix} I \\ 0 \end{bmatrix} E_{\Sigma}^{-} |_{\hat{i} \times \bar{j}} + P_i \begin{bmatrix} 0 \\ I \end{bmatrix} E_{\Sigma}^{-} |_{\hat{i} \times \check{j}}.$$

Substituting the above identities into (48), we obtain

$$\begin{aligned} E_{(i,j)} &= -P_i \begin{bmatrix} I \\ G_i \end{bmatrix} E_{\Sigma}^{|} |_{\hat{i} \times \hat{j}} \left(F_j \begin{bmatrix} 0 \\ H_j \end{bmatrix} \right)^T + P_i \begin{bmatrix} I \\ G_i \end{bmatrix} E_{\Sigma}^{|} |_{\hat{i} \times \check{j}} \left(F_j \begin{bmatrix} 0 \\ I \end{bmatrix} \right)^T \\ &\quad - P_i \begin{bmatrix} 0 \\ G_i \end{bmatrix} E_{\Sigma}^{-} |_{\hat{i} \times \bar{j}} + P_i \begin{bmatrix} 0 \\ I \end{bmatrix} E_{\Sigma}^{-} |_{\hat{i} \times \check{j}}. \end{aligned} \quad (49)$$

Based on (22), (26) and Lemma 4.1, we have the estimate

$$\begin{aligned} \|E_{(i,j)}\|_F &\leq \left\| \begin{bmatrix} I \\ G_i \end{bmatrix} \right\|_F \left(\|H_j\|_F \|E_{\Sigma}^{|} |_{\hat{i} \times \hat{j}}\|_F + \|E_{\Sigma}^{|} |_{\hat{i} \times \check{j}}\|_F \right) + \|G_i\|_F \|E_{\Sigma}^{-} |_{\hat{i} \times \bar{j}}\|_F + \|E_{\Sigma}^{-} |_{\hat{i} \times \check{j}}\|_F \\ &\leq \left\| \begin{bmatrix} I \\ G_i \end{bmatrix} \right\|_F \sqrt{\|H_j\|_F^2 + 1} \|E_{\Sigma}^{|} |_{\hat{i} \times \check{j}}\|_F + \sqrt{\|G_i\|_F^2 + 1} \|E_{\Sigma}^{-} |_{\hat{i} \times \check{j}}\|_F \\ &\leq s^2 |\bar{j}| \sqrt{|\hat{i}|} r^{(l)} \epsilon_{\text{SVD}} \|A_j^{|}\|_F + s |\hat{i}| \sqrt{|\bar{j}|} r^{(l)} \epsilon_{\text{SVD}} \|A_i^{-}\|_F. \end{aligned}$$

□

Based on Lemmas 4.2–4.3, we can estimate the total approximation error at level l in HSS construction.

Lemma 4.4

Assume ϵ_{SVD} and ϵ_{far} are the approximation tolerances used in the approximation of nearfield and farfield blocks in the HSS construction. Then the approximation error $E^{(l)}$ in (41) satisfies the following bound

$$\|E^{(l)}\|_F \leq s^2 2^{l/2+2} \left(r^{(l+1)} \right)^{3/2} \left(r^{(l)} \right)^{3/2} \epsilon_{\text{SVD}} \|A^{(l)}\|_F + 4s^2 r^{(l)} r^{(l+1)} \epsilon_{\text{far}} \|A^{(l)}\|_F,$$

where we have set $r^{(L+1)} := r^{(L)}$.

Proof

Based on Lemma 4.3, we know that the approximation error from the nearfield compression at level l can be estimated as follows:

$$\begin{aligned} \sum_{(i,j) \notin \text{Sep}} \|E_{(i,j)}\|_F^2 &\leq 2 \sum_{(i,j) \notin \text{Sep}} \left(s^4 |\bar{\mathbf{j}}|^2 |\bar{\mathbf{i}}| \left(r^{(l)} \right)^3 \epsilon_{\text{SVD}}^2 \|A_j^l\|_F^2 + s^2 |\bar{\mathbf{i}}|^2 |\bar{\mathbf{j}}| r^{(l)} \epsilon_{\text{SVD}}^2 \|A_i^-\|_F^2 \right) \\ &\leq 2 \sum_{i,j \text{ at level } l} \left(s^4 |\bar{\mathbf{j}}|^2 |\bar{\mathbf{i}}| \left(r^{(l)} \right)^3 \epsilon_{\text{SVD}}^2 \|A_j^l\|_F^2 + s^2 |\bar{\mathbf{i}}|^2 |\bar{\mathbf{j}}| r^{(l)} \epsilon_{\text{SVD}}^2 \|A_i^-\|_F^2 \right). \end{aligned}$$

Since $|\bar{\mathbf{i}}| \leq 2r^{(l+1)}$, $|\bar{\mathbf{j}}| \leq 2r^{(l+1)}$ for nodes i, j at level l and there are 2^{l-1} nodes at this level, we further have

$$\begin{aligned} \sum_{(i,j) \notin \text{Sep}} \|E_{(i,j)}\|_F^2 &\leq \sum_{i,j \text{ at level } l} \left(2s^4 \left(2r^{(l+1)} \right)^3 \left(r^{(l)} \right)^3 \epsilon_{\text{SVD}}^2 \|A_j^l\|_F^2 + 2s^2 \left(2r^{(l+1)} \right)^3 r^{(l)} \epsilon_{\text{SVD}}^2 \|A_i^-\|_F^2 \right) \\ &= \sum_{lv(i)=l} \left(\sum_{lv(j)=l} 2s^4 \left(2r^{(l+1)} \right)^3 \left(r^{(l)} \right)^3 \epsilon_{\text{SVD}}^2 \|A_j^l\|_F^2 \right) \\ &\quad + \sum_{lv(j)=l} \left(\sum_{lv(i)=l} 2s^2 \left(2r^{(l+1)} \right)^3 r^{(l)} \epsilon_{\text{SVD}}^2 \|A_i^-\|_F^2 \right) \\ &\leq \sum_{lv(i)=l} 2s^4 \left(2r^{(l+1)} \right)^3 \left(r^{(l)} \right)^3 \epsilon_{\text{SVD}}^2 \|A^{(l)}\|_F^2 \\ &\quad + \sum_{lv(j)=l} 2s^2 \left(2r^{(l+1)} \right)^3 r^{(l)} \epsilon_{\text{SVD}}^2 \|A^{(l)}\|_F^2 \\ &\leq 2^l \left(s^4 \left(2r^{(l+1)} \right)^3 \left(r^{(l)} \right)^3 + s^2 \left(2r^{(l+1)} \right)^3 r^{(l)} \right) \epsilon_{\text{SVD}}^2 \|A^{(l)}\|_F^2 \\ &\leq s^4 2^{l+4} \left(r^{(l+1)} \right)^3 \left(r^{(l)} \right)^3 \epsilon_{\text{SVD}}^2 \|A^{(l)}\|_F^2. \end{aligned}$$

Based on Lemma 4.2, we know that the approximation error for the farfield compression at level l satisfies:

$$\sum_{(i,j) \in \text{Sep}} \|E_{(i,j)}\|_F^2 \leq \sum_{(i,j) \in \text{Sep}} \left(2s^2 \sqrt{|\bar{\mathbf{i}}| |\bar{\mathbf{j}}|} r^{(l)} \right)^2 \epsilon_{\text{far}}^2 \|A_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}}\|_F^2.$$

Following the construction procedure (28) and the assumption (43), we have

$$|\bar{\mathbf{i}}| \leq 2r^{(l+1)}, \quad |\bar{\mathbf{j}}| \leq 2r^{(l+1)}. \quad (50)$$

Thus, we obtain the following estimation:

$$\begin{aligned} \sum_{(i,j) \in \text{Sep}} \|E_{(i,j)}\|_F^2 &\leq \sum_{(i,j) \in \text{Sep}} \left(4s^2 r^{(l+1)} r^{(l)} \right)^2 \epsilon_{\text{far}}^2 \|A_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}}\|_F^2 \\ &= 16s^4 \left(r^{(l+1)} r^{(l)} \right)^2 \epsilon_{\text{far}}^2 \sum_{(i,j) \in \text{Sep}} \|A_{\bar{\mathbf{i}} \times \bar{\mathbf{j}}}\|_F^2 \\ &\leq 16s^4 \left(r^{(l+1)} r^{(l)} \right)^2 \epsilon_{\text{far}}^2 \|A^{(l)}\|_F^2. \end{aligned}$$

To sum up both farfield and nearfield approximation errors, we obtain the estimate for the overall approximation error introduced at level l :

$$\begin{aligned} \|E^{(l)}\|_F &= \left(\sum_{(i,j) \notin \text{Sep}} \|E_{(i,j)}\|_F^2 + \sum_{(i,j) \in \text{Sep}} \|E_{(i,j)}\|_F^2 \right)^{\frac{1}{2}} \\ &\leq s^2 2^{l/2+2} \left(r^{(l+1)} \right)^{3/2} \left(r^{(l)} \right)^{3/2} \epsilon_{\text{SVD}} \|A^{(l)}\|_F + 4s^2 r^{(l)} r^{(l+1)} \epsilon_{\text{far}} \|A^{(l)}\|_F. \end{aligned}$$

□

The overall HSS approximation error in the Frobenius norm can then be derived in the following theorem.

Theorem 4.1

Suppose the HSS tree has L levels. With the assumptions in Lemma 4.4, SMASH produces the following factorization

$$\begin{aligned} A = U^{(L)} &\left(U^{(L-1)} \left(\dots \left(U^{(2)} B^{(1)} \left(V^{(2)} \right)^T + B^{(2)} \right) \dots \right) \left(V^{(L-1)} \right)^T \right. \\ &\left. + B^{(L-1)} \right) \left(V^{(L)} \right)^T + B^{(L)} + E, \end{aligned} \quad (51)$$

where the approximation error E satisfies the estimate

$$\|E\|_F \leq C_1 \epsilon_{\text{SVD}} \|A\|_F + C_2 \epsilon_{\text{far}} \|A\|_F,$$

with

$$\begin{aligned} C_1 &= \sum_{l=2}^{L-1} 2^{L+l/2+2} s^{2L-2l+2} \left(r^{(l+1)} \dots r^{(L)} \right)^2 \left(r^{(l+1)} \right)^{3/2} \left(r^{(l)} \right)^{5/2} \\ C_2 &= \sum_{l=2}^L 2^{L+2} s^{2L-2l+2} \left(r^{(l)} r^{(l+1)} \dots r^{(L)} \right)^2 r^{(l+1)}. \end{aligned}$$

Proof

According to (44), we know that the overall HSS approximation error E has the expression

$$E = \left(U^{(L)} \dots U^{(3)} \right) E^{(2)} \left(V^{(L)} \dots V^{(3)} \right)^T + \dots + U^{(L)} E^{(L-1)} \left(V^{(L)} \right)^T + E^{(L)}. \quad (52)$$

Note that the column size of $\left(U^{(L)} \dots U^{(l+1)} \right)$ is bounded by $r^{(l)} 2^l$. Thus we have

$$\|U^{(L)} \dots U^{(l+1)}\|_F \leq \sqrt{r^{(l)} 2^l} \|U^{(L)} \dots U^{(l+1)}\|_2 \leq 2^{L/2} s^{L-l} \sqrt{r^{(l)} r^{(l+1)}} \dots r^{(L)}, \quad (53)$$

and same upper bound holds for $\|V^{(L)} \dots V^{(l+1)}\|_F$. It follows from (52) that

$$\begin{aligned} \|E\|_F &\leq \sum_{l=2}^{L-1} \|U^{(L)} \dots U^{(l+1)}\|_F \|E^{(l)}\|_F \|V^{(L)} \dots V^{(l+1)}\|_F + \|E^{(L)}\|_F \\ &\leq \sum_{l=2}^{L-1} 2^L s^{2L-2l} r^{(l)} \left(r^{(l+1)} \dots r^{(L)} \right)^2 s^2 2^{l/2+2} \left(r^{(l+1)} \right)^{3/2} \left(r^{(l)} \right)^{3/2} \epsilon_{\text{SVD}} \|A^{(l)}\|_F \\ &+ \sum_{l=2}^{L-1} 2^L s^{2L-2l} r^{(l)} \left(r^{(l+1)} \dots r^{(L)} \right)^2 4s^2 r^{(l+1)} r^{(l)} \epsilon_{\text{far}} \|A^{(l)}\|_F + \|E^{(L)}\|_F \\ &\leq \sum_{l=2}^{L-1} 2^{L+l/2+2} s^{2L-2l+2} \left(r^{(l+1)} \dots r^{(L)} \right)^2 \left(r^{(l+1)} \right)^{3/2} \left(r^{(l)} \right)^{5/2} \epsilon_{\text{SVD}} \|A\|_F \\ &+ \sum_{l=2}^L 2^{L+2} s^{2L-2l+2} \left(r^{(l)} r^{(l+1)} \dots r^{(L)} \right)^2 r^{(l+1)} \epsilon_{\text{far}} \|A\|_F. \end{aligned} \quad (54)$$

□

It is easy to see that the bound in (54) is quite pessimistic because we bound $\|A^{(l)}\|_F$ from above by $\|A\|_F$, where $A^{(l)}$ (defined in (42)) is a submatrix of A of size $|\mathbf{I}^{(l)}| \times |\mathbf{J}^{(l)}|$.

Corollary 4.1

Besides the assumptions in Theorem 4.1, if there also exists a constant $r \geq 2$ such that $r^{(l)} \leq r$ for each $l > 1$, then the approximation error E in SMASH satisfies the estimate:

$$\|E\|_F \leq (2r^2s^2)^L (16\epsilon_{\text{SVD}} + 8\epsilon_{\text{far}}) \|A\|_F.$$

Next we estimate the error in \mathcal{H}^2 approximation. Note that, by setting $\epsilon_{\text{SVD}} = 0$, the error estimate in Theorem 4.1 also holds for the \mathcal{H}^2 construction with a possibly different constant C_2 depending on dimension d . In fact, analogous to Lemma 4.4, we first have the following estimate for the \mathcal{H}^2 construction:

$$\|E^{(l)}\|_F \leq 4s^2r^{(l)}r^{(l+1)}\epsilon_{\text{far}}\|A^{(l)}\|_F,$$

where $E^{(l)}$ denotes the approximation error introduced at level l , as defined in (41). Assume the \mathcal{H}^2 matrix is associated with a perfect 2^d -tree \mathcal{T} ($d \in \{1, 2, 3\}$), i.e., each nonleaf node of \mathcal{T} has 2^d children and all leaves are at the same level. For a node i at level l , since $|\mathbf{i}| \leq 2^d r^{(l+1)} \leq 2^d r^{(l)}$, it follows from (36), (31) and Lemma 4.1 that

$$\|U^{(l)}\|_2 \leq s2^{d/2}r^{(l)}.$$

Notice that the column size of $(U^{(L)} \dots U^{(l+1)})$ is bounded by $r^{(l)}2^{dl}$. Then the counterpart of (53) can be obtained for the \mathcal{H}^2 construction as:

$$\|U^{(L)} \dots U^{(l+1)}\|_F \leq \sqrt{r^{(l)}2^{dl}} \|U^{(L)} \dots U^{(l+1)}\|_2 \leq 2^{dL/2} s^{L-l} \sqrt{r^{(l)}r^{(l+1)}} \dots r^{(L)},$$

and the same upper bound holds for $\|V^{(L)} \dots V^{(l+1)}\|_F$. The total approximation error can now be obtained by means of (52):

$$\begin{aligned} \|E\|_F &\leq \sum_{l=2}^{L-1} \|U^{(L)} \dots U^{(l+1)}\|_F \|E^{(l)}\|_F \|V^{(L)} \dots V^{(l+1)}\|_F + \|E^{(L)}\|_F \\ &\leq \sum_{l=2}^{L-1} 2^{dL} s^{2L-2l} r^{(l)} \left(r^{(l+1)} \dots r^{(L)} \right)^2 4s^2 r^{(l+1)} r^{(l)} \epsilon_{\text{far}} \|A^{(l)}\|_F + 4s^2 (r^{(L)})^2 \epsilon_{\text{far}} \|A\|_F \\ &\leq \sum_{l=2}^{L-1} 2^{dL+2} s^{2L-2l+2} \left(r^{(l)} r^{(l+1)} \dots r^{(L)} \right)^2 r^{(l+1)} \epsilon_{\text{far}} \|A\|_F + 4s^2 (r^{(L)})^2 \epsilon_{\text{far}} \|A\|_F \\ &\leq \sum_{l=2}^L 2^{dL+2} s^{2L-2l+2} \left(r^{(l)} r^{(l+1)} \dots r^{(L)} \right)^2 r^{(l+1)} \epsilon_{\text{far}} \|A\|_F. \end{aligned}$$

The above error analysis yields the following theorem.

Theorem 4.2

Suppose \mathcal{T} is a perfect 2^d -tree with L levels, associated with the \mathcal{H}^2 approximation of A . Under the assumptions in (43), SMASH produces the following factorization

$$\begin{aligned} A = &U^{(L)} \left(U^{(L-1)} \left(\dots \left(U^{(2)} B^{(1)} \left(V^{(2)} \right)^T + B^{(2)} \right) \dots \right) \left(V^{(L-1)} \right)^T \right. \\ &\left. + B^{(L-1)} \right) \left(V^{(L)} \right)^T + B^{(L)} + E, \end{aligned} \quad (55)$$

where the approximation error E satisfies the estimate

$$\|E\|_F \leq C\epsilon_{\text{far}} \|A\|_F,$$

with

$$C = \sum_{l=2}^L 2^{dL+2} s^{2L-2l+2} \left(r^{(l)} r^{(l+1)} \dots r^{(L)} \right)^2 r^{(l+1)}.$$

Corollary 4.2

Besides the assumptions in Theorem 4.2, if there also exists a constant $r \geq 2$ such that $r^{(l)} \leq r$ for each $l > 1$, then the approximation error E in SMASH satisfies the estimate:

$$\|E\|_F \leq (2^d r^2 s^2)^L 8\epsilon_{\text{far}} \|A\|_F.$$

Finally, we will show that the term $(2^d r^2 s^2)^L$ in Corollaries 4.1-4.2 can be bounded by functions of n which are independent of ϵ_{far} and ϵ_{svd} under certain assumptions.

Proposition 4.1

Given r , let \mathcal{T} be a perfect 2^d -tree of L levels such that $r^{2dL} \leq n$. Assume without loss of generality that $s \leq 2$. Then the constants in Corollary 4.1 and Corollary 4.2 satisfy the following estimation:

$$(2^d r^2 s^2)^L \leq n^{\frac{1}{2d} \log_2 n + \frac{2}{d} + 1}. \quad (56)$$

Proof

Since $r^{2dL} \leq n$, we have $L \leq \frac{1}{d} \log_2 \frac{n}{r}$. It follows that

$$2^{dL} s^{2L} \leq 2^{dL} 2^{2L} \leq \left(\frac{n}{r} \right)^{\frac{2}{d} + 1} \leq n^{\frac{2}{d} + 1}.$$

Next it remains to bound $r^{2L} \leq r^{\frac{2}{d} \log_2 \frac{n}{r}}$. To this end, we define $g(x) = x^{\log_2 \frac{n}{x}}$ for $x \geq 1$. Then $r^{2L} \leq (g(r))^{\frac{2}{d}}$. It can be verified that the global maximum of g is achieved at $x = \sqrt{n}$. Therefore,

$$r^{2L} \leq (g(r))^{\frac{2}{d}} \leq (g(\sqrt{n}))^{\frac{2}{d}} = n^{\frac{1}{2d} \log_2 n},$$

and we conclude that

$$(2^d r^2 s^2)^L \leq n^{\frac{1}{2d} \log_2 n + \frac{2}{d} + 1}. \quad \square$$

Proposition 4.1 also indicates that C_1 , C_2 in Theorem 4.1 and C in Theorem 4.2 all can be bounded by constants which are independent of ϵ_{far} and ϵ_{svd} under certain assumptions.

5. COMPLEXITY ANALYSIS

This section studies the complexity of SMASH for an $n \times n$ matrix. For simplicity, we only consider the case when $X = Y$ and the points are uniformly distributed. Under this assumption, a perfect tree \mathcal{T} will be used for both HSS and \mathcal{H}^2 structures.

5.1. Complexity for the HSS construction

We start with the HSS construction case. Since the HSS matrix structure is only efficient for one dimensional problems, we will focus on the one-dimensional problems in this section. Suppose \mathcal{T} has L levels such that $n = O(r^{2L})$, where r is a positive integer such that the rank of HSS generators is bounded above by r and

$$|\bar{\mathbf{i}}^{\text{row}}| \leq 2r, \quad |\bar{\mathbf{i}}^{\text{col}}| \leq 2r, \quad \forall i \neq \text{root}. \quad (57)$$

Notice that in the context of integral equations in potential theory, the assumption (57) in general holds only for integral equations defined on a curve. Since the points are uniformly distributed, for each nonroot node i , the number of nodes in \mathcal{N}_i (See Appendix) is very small, which we assume to be bounded above by 3. Under these assumptions, we have the following complexity estimate.

Theorem 5.1

Let \mathcal{T} be a perfect binary tree with L levels and (57) hold. Then the complexity of SMASH for the HSS construction in Section 3.3.1 is $O(n)$.

Proof

Based on (57), it is easy to see that, for each nonroot node i , the compression cost for its nearfield blocks in (21) is $O(r^3)$. This is because the size of the nearfield block row in (21) is no larger than $2r$ -by- $6r$ under the above assumption for \mathcal{N}_i . Besides, the farfield basis matrix \hat{U}_i has column size at most r , so the cost of an SRRQR procedure in (23) is $O(r^3)$. Therefore, the compression cost associated with each nonroot node i is $O(r^3)$ and the complexity of the HSS construction is $O(2^L r^3) = r^2 O(n) = O(n)$. \square

5.2. Complexity for the \mathcal{H}^2 construction

For the \mathcal{H}^2 construction case, we assume that when $X \subset \mathbb{R}^d$, \mathcal{T} is a perfect 2^d -tree with L levels such that $n = O(r2^{dL})$ and r is a positive integer such that the rank of \mathcal{H}^2 generators is bounded by r and

$$|\bar{\mathbf{i}}^{row}| \leq r2^d, \quad |\bar{\mathbf{i}}^{col}| \leq r2^d, \quad \forall i \neq \text{root}. \quad (58)$$

The analysis here is simpler than that of the HSS construction in Section 5.1. Since each node i only involves the compression of farfield basis $\hat{U}_i|_{\bar{\mathbf{i}}^{row}}$ (as well as $\hat{V}_i|_{\bar{\mathbf{i}}^{col}}$), whose size is no larger than $r2^d$ -by- r under the assumption (58), we deduce that the compression cost associated with each node is $O(r^3)$. As a result, the complexity of the \mathcal{H}^2 construction is $O(2^{dL} r^3) = r^2 O(n)$. Thus we conclude:

Theorem 5.2

Let \mathcal{T} be a perfect 2^d -tree with L levels and (58) hold. Then the complexity of SMASH for the \mathcal{H}^2 construction in Section 3.3.2 is $O(n)$.

6. NUMERICAL EXAMPLES

In this section, we present numerical examples to illustrate the performance of SMASH. All of the numerical results were performed in MATLAB R2014b on a macbook air with a 1.6 GHz CPU and 8 GB of RAM. The following notation is used throughout the section:

- n : the size of A ;
- t_{constr} : wall clock time for constructing \hat{A} in seconds;
- t_{matvec} : wall clock time for multiplying \hat{A} with a vector in seconds;
- t_{sol} : wall clock time for solving $\hat{A}x = b$ in seconds;
- ϵ_{svd} : relative tolerance used in the truncated SVD for the nearfield compression;
- $\text{rand}([0, 1])$: a random number sampled from the uniform distribution in $[0, 1]$.

6.1. Choice of parameters

Since the quality of a degenerate approximation depends on the underlying kernel function, there is no rule of thumb in general on choosing the parameters to satisfy a prescribed tolerance. For completeness, here we present a heuristic approach that we use in all numerical experiments on the choice of parameters.

Given a matrix A and a tolerance ϵ , suppose one wants to construct a hierarchical matrix \hat{A} (\mathcal{H} , \mathcal{H}^2 , or HSS) such that $\|A - \hat{A}\|_{\max} \approx \epsilon$. Then the following approach is adopted to determine parameters τ, r .

Table I. 2D test in Section 6.2.1: construction and matrix-vector product of an \mathcal{H}^2 matrix.

$n = m^2$	$\ \hat{A}u - Au\ /\ Au\ $	t_{constr}	t_{matvec}
1600	6.69×10^{-13}	0.52	0.02
6400	2.00×10^{-12}	1.97	0.07
25600	3.65×10^{-12}	9.53	0.30
102400	4.87×10^{-12}	39.47	1.18

The choice of separation ratio $\tau \in (0, 1)$ only depends on the dimension of the problem, so it is chosen first. We choose τ such that $\tau \leq 0.7$ and, in general, a slightly larger τ is preferred for higher dimensional problems. For example, we choose $\tau = 0.6$ for essentially one-dimensional problems, such as those in Section 6.3 and Section 6.4; we choose $\tau = 0.65$ for problems in two or three dimensions in Section 6.2.

Having chosen a separation ratio τ , we use the following function to determine the farfield approximation rank r used in constructing \hat{U}_i, \hat{V}_i (before the SRRQR postprocessing):

$$r = \begin{cases} \lfloor \log \epsilon / \log \tau - 20 \rfloor, & \text{if } \epsilon < 10^{-8}, \\ \lfloor \log \epsilon / \log \tau - 15 \rfloor, & \text{if } 10^{-8} \leq \epsilon < 10^{-6}, \\ \max\{\lfloor \log \epsilon / \log \tau - 10 \rfloor, 5\} & \text{otherwise,} \end{cases}$$

where $\lfloor x \rfloor$ yields the largest integer less than or equal to x . For example, in Section 6.2.1, $\epsilon = 10^{-7}$, $\tau = 0.65$, $r = 22$; in Section 6.3, $\epsilon = 10^{-8}$, $\tau = 0.6$, $r = 21$; in Section 6.4, $\epsilon = 10^{-10}$, $\tau = 0.6$, $r = 25$.

6.2. Construction and matrix-vector multiplication of \mathcal{H}^2 matrices

In this section, we perform numerical experiments to test the construction and matrix-vector multiplication of an \mathcal{H}^2 approximation associated with kernels in both two and three dimensions. A complicated three-dimensional geometry (see Figure 5) is presented to illustrate the robustness of the algorithm.

6.2.1. Two dimensions We first consider the kernel in (1) with $d_x = 1$. We chose X as a uniform $m \times m$ grid in $[0, 1]^2$ and $A = [\kappa(x, y)]_{x, y \in X}$. The computational domain $[0, 1]^2$ was recursively divided into 4 subdomains until the number of the points inside each domain was less than or equal to 50. We embedded $[0, 1]^2$ in the complex plane and used the truncated Taylor expansion (4) with $r = 22$ terms and the separation ratio $\tau = 0.65$ to compress farfield blocks. The approximation error of this construction is measured by the relative error $\|\hat{A}u - Au\|/\|Au\|$, where u is a random vector of length $n = m^2$ with entries generated by $\text{rand}([0, 1])$. The numerical results are reported in Table I.

As can be seen from Table I, SMASH for the \mathcal{H}^2 construction and the matrix-vector multiplication described in Section 3.4 scale linearly, which is consistent with the complexity analysis in Section 5.

6.2.2. Three dimensions In three dimensions, we consider the kernel

$$\kappa(x, y) = \begin{cases} \frac{\log|x-y|}{|x-y|}, & \text{if } x \neq y, \\ 1, & \text{if } x = y, \end{cases}$$

which has a stronger singularity than $1/|x-y|$ near $x=y$. The geometry here is a triceratops embedded in the cube $[-100, 100]^3$ as illustrated in Figure 5[†]. The points in X are drawn randomly

[†]The datasets are from the point cloud tools <http://www.geo.tuwien.ac.at/downloads/pg/pctools/pctools.html>

Table II. 3D test in Section 6.2.2: construction and matrix-vector product of an \mathcal{H}^2 matrix.

n	$\ \hat{A}u - Au\ /\ Au\ $	t_{constr}	t_{matvec}
10000	1.98×10^{-6}	9.12	0.14
20000	3.83×10^{-6}	23.20	0.38
40000	5.83×10^{-6}	50.39	0.74
80000	7.20×10^{-6}	115.03	1.57

from the triceratops and $A = [\kappa(x, y)]_{x, y \in X}$. The \mathcal{H}^2 approximation \hat{A} is constructed based on interpolation in three dimensions, where 5 Chebyshev points are used in each direction.

It is easily seen from Table II that the time for construction or matrix-vector multiplication scales linearly in terms of the matrix size. Moreover, despite the complicated geometry, the approximation accuracy remains relatively high as problem sizes increase.

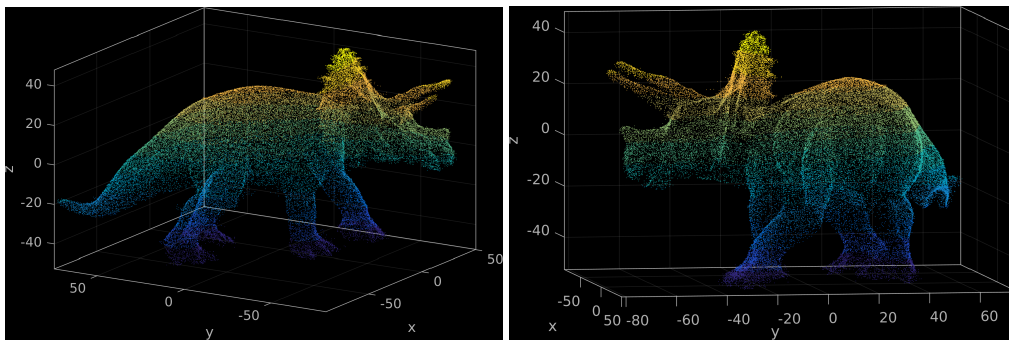


Figure 5. The 3D triceratops geometry used for the numerical experiments in Table II.

6.3. Cauchy-like matrices

We consider in this section the numerical solution of Cauchy-like matrices. It is known that Cauchy-like matrices are related to other types of structured matrices including Toeplitz matrices, Vandermonde matrices, Hankel matrices and their variants [55, 56, 57, 58, 59, 60]. Consider the kernel $\kappa(x, y) = 1/(x - y)$, $x \neq y \in \mathbb{C}$. Let $x_i, y_j (i, j = 1 : n)$ be $2n$ pairwise distinct points in \mathbb{C} . The Cauchy matrix is then given by $\mathcal{C} = [\kappa(x_i, y_j)]_{i, j=1:n}$, which is known to be invertible [61]. Given two matrices $w, v \in \mathbb{C}^{n \times p}$, the (i, j) -entry of a Cauchy-like matrix A associated with generators w, v is defined by [62]

$$a_{i,j} = \frac{1}{x_i - y_j} \sum_{l=1}^p w_{i,l} v_{j,l}. \quad (59)$$

For simplicity, we consider the case $p = 2$, i.e., w (as well as v) is composed of two column vectors. Denote by $\hat{w}_1, \hat{w}_2, \hat{v}_1, \hat{v}_2$ the column vectors in w, v , i.e., $w = [\hat{w}_1, \hat{w}_2], v = [\hat{v}_1, \hat{v}_2]$. It can be seen that A can be written as

$$A = \text{diag}(\hat{w}_1) \mathcal{C} \text{diag}(\hat{v}_1) + \text{diag}(\hat{w}_2) \mathcal{C} \text{diag}(\hat{v}_2). \quad (60)$$

Existing approaches for solving Cauchy or Cauchy-like linear systems associated with points in \mathbb{R} mainly rely on some variants of Gaussian elimination with pivoting techniques. For example, fast $O(n^2)$ algorithms for solving Cauchy linear systems can be found in [63, 64, 65, 66], etc.; a superfast $O(n \log^3 n)$ algorithm based on a sequential block Gaussian elimination process was proposed in [67]. The performance of most existing methods depends on the the distribution of

point sets x, y . As pointed out in [64], if two sets of points x, y can not be separated, for example, when they are interlaced, existing algorithms (for example, BP-type algorithm of [63]) suffer from backward stability issues. Moreover, due to the use of pivoting techniques, the accuracy of existing algorithms heavily depend on the ordering of points [63, 64] and the analysis is limited to the case when the points are in \mathbb{R} .

Therefore, in view of the issues mentioned above, we assume x_i, y_j are mixed together such that in adaptive partitioning (see Section 3.1), each box contains the same number of points from x_i and y_j . We also consider that x_i, y_j are distributed on a curve in \mathbb{R}^2 as illustrated in Fig. 7 to demonstrate that the algorithm is independent of the ordering of points and is applicable for points in \mathbb{C} .

We construct the HSS approximation \hat{A} to A using SMASH discussed in Section 3.3.1 and then solve the linear system associated with \hat{A} using a fast ULV factorization solver [5]. Due to the choice of stable expansion in (4), arbitrarily high approximation accuracy can be achieved without stability issues [49].

Note that the HSS approximation to \mathcal{C} can be readily obtained as in Section 3. Consequently, the HSS approximations to $\text{diag}(\hat{w}_1)\mathcal{C}\text{diag}(\hat{v}_1)$ and $\text{diag}(\hat{w}_2)\mathcal{C}\text{diag}(\hat{v}_2)$ can be derived by modifying U, V, D generators, respectively. The sum of these two HSS representations is also an HSS matrix whose generators can be easily obtained using the technique presented in [68] by merging the two sets of HSS generators. Hence the HSS approximation to A is derived.

In the first experiment, the point sets $\{x_k\}_{k=1}^n, \{y_k\}_{k=1}^n$ are chosen as follows:

$$x_k = k/(n+1), \quad y_k = x_k + 10^{-7} * \text{rand}([0, 1]), \quad k = 1, \dots, n. \quad (61)$$

In the second experiment, the point sets are distributed on the curve illustrated in Fig. 6 that is parametrized by

$$\gamma(t) = e^{-\pi i/6} * [(0.5 + \sin(4\pi t)) \cos(2\pi t) + i(0.5 + \sin(4\pi t)) \sin(2\pi t)], \quad t \in [0, 1],$$

and $\{x_k\}_{k=1}^n, \{y_k\}_{k=1}^n$ are given by

$$x_k = \gamma(k/(n+1)), \quad y_k = \gamma(x_k + 10^{-7} * \text{rand}([0, 1])), \quad k = 1, \dots, n. \quad (62)$$

In the third experiment, the point set $\{x_k\}_{k=1}^n$ is on the snail geometry in \mathbb{C} as illustrated in Fig. 7 and the point set $\{y_k\}_{k=1}^n$ is given by

$$y_k = x_k + 10^{-7} * \text{rand}([0, 1]), \quad k = 1, \dots, n. \quad (63)$$

For generators of this Cauchy-like matrix, i.e., $w = [\hat{w}_1, \hat{w}_2], v = [\hat{v}_1, \hat{v}_2]$, we chose $\hat{w}_l, \hat{v}_l (l = 1, 2)$ such that each entry in those vectors was given by $\text{rand}([0, 1])$. In order to solve the linear system $Au = b$, we constructed an HSS approximation \hat{A} to A in (60) by SMASH. 1D boxes (i.e., intervals) and 2D boxes (i.e., rectangles) are used in adaptive partitioning for point sets in (61) and (63), respectively. The right subfigures in Fig. 7 illustrate the adaptive partitioning using 2D boxes as described in Remark 1. We chose separation ratio as $\tau = 0.6$ and adaptive partitioning stopped when the number of points inside each box was less than or equal 50. The nearfield blocks were compressed through SVD with truncation tolerance 10^{-9} . The exact solution was set to be a column vector u of length n with entries generated by $\text{rand}([0, 1])$, and the right-hand side b was formed by $b = Au$.

The numerical results for three Cauchy-like matrix problems are reported in Table III. From Table III, we see that the computational time for both construction and the solution scale linearly, and SMASH in Section 3.3.1 is quite robust with respect to complex geometries. Moreover, it can be seen that SMASH is independent of the ordering of points.

6.4. Integral equations

In this section, we solve Laplace boundary value problems via the integral equation method. Assume Ω is a smooth simply-connected domain in \mathbb{R}^2 and let $\Gamma = \partial\Omega$ be the boundary of Ω of class C^2 .




Figure 6. Honeybee geometry used for the numerical experiments in Table III. Left: Original curve; Right: Adaptive partitioning of the curve for the case when $n = 12800$ in Table III.




Figure 7. Snail geometry used for the numerical experiments in Table III. Left: Original curve; Right: Adaptive partitioning of the curve for the case when $n = 12800$ in Table III.

Table III. Numerical results for solving the Cauchy-like matrix when $\{x_k\}_{k=1}^n$ are distributed on three different curves.

curve	n	$\ u - \hat{u}\ /\ u\ $	$\ Au - A\hat{u}\ /\ Au\ $	t_{constr}	t_{sol}
[0, 1]	1600	7.69×10^{-12}	5.56×10^{-15}	0.33	0.11
	3200	1.01×10^{-09}	3.87×10^{-14}	0.63	0.19
	6400	5.58×10^{-11}	5.58×10^{-14}	1.29	0.35
	12800	1.47×10^{-08}	5.87×10^{-14}	2.58	0.69
honeybee	1600	9.37×10^{-11}	8.09×10^{-14}	1.14	0.28
	3200	9.78×10^{-10}	5.60×10^{-13}	2.22	0.51
	6400	1.55×10^{-09}	9.16×10^{-13}	4.42	0.96
	12800	2.76×10^{-09}	1.54×10^{-12}	8.49	1.87
snail	1600	1.34×10^{-11}	1.02×10^{-15}	1.75	0.39
	3200	2.98×10^{-11}	6.61×10^{-16}	2.93	0.69
	6400	3.65×10^{-10}	5.88×10^{-15}	5.81	1.30
	12800	2.78×10^{-10}	6.27×10^{-15}	10.57	2.39

Consider the interior Dirichlet problem: find $u \in C^2(\Omega) \cap C(\bar{\Omega})$ such that

$$\begin{aligned} \Delta u &= 0 \quad \text{in } \Omega, \\ u &= u_D \quad \text{on } \Gamma, \end{aligned} \tag{64}$$

where $u_D \in C(\Gamma)$ is given. With smooth boundary curves and Dirichlet data, the wellposedness of this problem is well studied in potential theory [69, 70].

The fundamental solution and its gradient (in terms of \mathbf{y}) for the Laplace equation in \mathbb{R}^2 are given by:

$$\Phi(\mathbf{x}, \mathbf{y}) = -\frac{1}{2\pi} \log |\mathbf{x} - \mathbf{y}|, \quad \text{and} \quad \nabla_{\mathbf{y}} \Phi(\mathbf{x}, \mathbf{y}) = -\frac{1}{2\pi} \frac{\mathbf{y} - \mathbf{x}}{|\mathbf{x} - \mathbf{y}|^2}.$$

Let ν_y denote the unit outer normal at point $\mathbf{y} \in \Gamma$. The double layer potential with continuous density σ is given by

$$K\sigma(\mathbf{x}) := \int_{\Gamma} \frac{\partial\Phi(\mathbf{x}, \mathbf{y})}{\partial\nu_y} \sigma(\mathbf{y}) ds_y = \int_0^1 \frac{\partial\Phi(\mathbf{x}, \mathbf{r}(t))}{\partial\nu_y} |\mathbf{r}'(t)| \sigma(\mathbf{r}(t)) dt, \quad \mathbf{x} \in \Omega, \quad (65)$$

where we assume Γ is parametrized by $\mathbf{r}(t) : [0, 1] \rightarrow \mathbb{R}^2$.

Given Dirichlet data $u_D \in C(\Gamma)$ in (64), we solve the following integral equation for $\sigma \in C(\Gamma)$:

$$\left(K - \frac{1}{2}I\right)\sigma = u_D, \quad \text{on } \Gamma. \quad (66)$$

It is well-known ([70]) that the problem above for $\sigma \in C(\Gamma)$ is well-posed, and the corresponding double layer potential $u := K\sigma$ solves the interior Dirichlet problem (64).

Denote the kernel in the second integral in (65) by

$$\kappa(s, t) := \frac{\partial\Phi(\mathbf{r}(s), \mathbf{r}(t))}{\partial\nu_y} |\mathbf{r}'(t)|. \quad (67)$$

It is easily seen that the smoothness of the above kernel function depends on the underlying boundary curve and it is known that ([71]) if $\mathbf{r} \in C^\infty([0, 1])$, then $\kappa \in C^\infty([0, 1] \times [0, 1])$.

Several Laplace problems (64) with the same exact solution but different domains are considered here. The first domain Ω is a ram head whose boundary curve Γ is parametrized by $\mathbf{r}(t) = (r_1(t), r_2(t))$ for $t \in [0, 1]$:

$$\begin{aligned} r_1(t) &= 2 \cos(2\pi t), \\ r_2(t) &= 1 + \sin(2\pi t) - 1.4 \cos^4(4\pi t). \end{aligned} \quad (68)$$

The second domain is a sunflower whose boundary curve Γ is parametrized by:

$$\begin{aligned} r_1(t) &= (1.3 + 1.25 \cos(40\pi t)) \cos(2\pi t), \\ r_2(t) &= (1.3 + 1.25 \cos(40\pi t)) \sin(2\pi t). \end{aligned} \quad (69)$$

We chose the Dirichlet data u_D such that the exact solution of (64) is

$$u(\mathbf{x}) = \log |\mathbf{x} - \mathbf{x}_0|,$$

where the source point $\mathbf{x}_0 = (2, 1.5)$ is in the exterior of Ω . Illustrations for the curves parametrized in (68) and (69) are shown in Fig. 8 and Fig. 9, respectively. We used Nyström method with trapezoidal rule to discretize (66). Since the curve Γ and the kernel are both smooth, Nyström discretization converges with a convergence rate proportional to that of the quadrature rule.

As in Section 6.3, we applied the HSS matrix techniques to approximate and solve the resulting matrix from the Nyström discretization of the integral equation in (66). The adaptive partitioning based on bisection (cf. Remark .1) was applied to a box covering the domain Ω in \mathbb{R}^2 and each box in leaf level contained no more than 50 quadrature points on Γ . Empty boxes were discarded during the partitioning. Illustrations of adaptive partitioning are shown in the right subfigures of Fig. 8 and Fig. 9, when 10240 quadrature points are in use. A binary tree \mathcal{T} was then generated corresponding to each adaptive partitioning. In the construction of the HSS matrix, the bases for the farfield blocks were approximated by polynomial interpolation (with 25 interpolating points) with respect to the separation ratio 0.6 and the bases for the nearfield blocks are computed by the truncated SVD with the tolerance $\epsilon_{\text{svd}} = 10^{-11}$. In order to test the convergence of the discretization, for curves in Fig. 8, we compared the numerical solution \hat{u} with the exact solution u by evaluating them at point $\mathbf{x}^* = (0.1, 0.1)$ inside Ω . For the curve in Fig. 9, the evaluation point is chosen as $\mathbf{x}^* = (1.5, 0)$ inside Ω . The numerical results for the ram head and sunflower problems are shown in Table IV and Table V, respectively.

From Table IV and Table V, it can be seen that the HSS matrix methods achieve linear complexity at both the construction and solution stages. In addition, it is worth noting that the

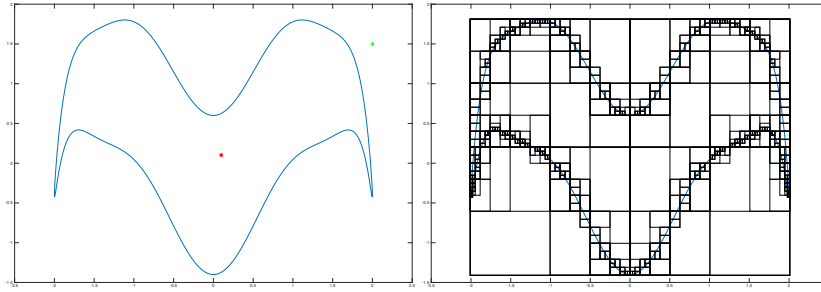


Figure 8. Ram head domain for the Dirichlet problem (64) with source point and evaluation point marked as green '+' and red '*', respectively. Left: Original curve; Right: Adaptive partitioning of the ram head boundary curve for the case when 10240 quadrature points are used in Table IV.




Figure 9. Sunflower domain for the Dirichlet problem (64) with source point and evaluation point marked as green '+' and red '*', respectively. Left: Original curve; Right: Adaptive partitioning of the sunflower boundary curve for the case when 10240 quadrature points are used in Table V.

Table IV. Numerical results for solving a 2D Laplace Dirichlet problem in a ram head domain as shown in Fig. 8.

n	$ u(\mathbf{x}^*) - \hat{u}(\mathbf{x}^*) $	$\ A - \hat{A}\ _{\max}$	$\text{cond}(A)$	t_{constr}	t_{sol}
160	5.03×10^{-08}	1.06×10^{-10}	6.15×10^{01}	0.42	0.107
320	9.54×10^{-11}	6.81×10^{-10}	6.92×10^1	1.54	0.042
640	1.91×10^{-12}	7.98×10^{-10}	6.00×10^1	2.93	0.127
1280	8.22×10^{-13}	2.26×10^{-09}	6.02×10^1	5.41	0.103
2560	7.78×10^{-13}	3.90×10^{-09}	6.02×10^1	9.76	0.177
5120	1.50×10^{-13}	9.63×10^{-09}	6.02×10^1	18.74	0.282
10240	1.96×10^{-12}	1.01×10^{-08}	6.02×10^{01}	34.78	0.591

Table V. Numerical results for solving a 2D Laplace Dirichlet problem in a sunflower domain as shown in Fig. 9.

n	$ u(\mathbf{x}^*) - \hat{u}(\mathbf{x}^*) $	$\ A - \hat{A}\ _{\max}$	$\text{cond}(A)$	t_{constr}	t_{sol}
640	2.96×10^{-02}	1.20×10^{-08}	5.65×10^3	5.42	0.105
1280	1.25×10^{-03}	2.65×10^{-08}	2.23×10^3	16.80	0.256
2560	1.88×10^{-06}	3.56×10^{-08}	1.85×10^3	41.76	0.624
5120	1.02×10^{-10}	4.27×10^{-07}	1.72×10^4	102.65	1.429
10240	1.66×10^{-11}	3.80×10^{-07}	1.12×10^4	192.78	2.165
20480	8.03×10^{-10}	9.55×10^{-07}	6.86×10^3	316.68	3.219

Table VI. Comparison of exact ϵ -rank and approximation rank of $M = A_{i^{row} \times j^{col}}$ for the ram head geometry in Fig.8, the sunflower geometry in Fig.9 with $\epsilon = 10^{-3}, 10^{-6}, 10^{-10}$

geometry	n	size(M)	$\epsilon = 10^{-3}$		$\epsilon = 10^{-6}$		$\epsilon = 10^{-10}$	
			$r_\epsilon(M)$	size(B_i)	$r_\epsilon(M)$	size(B_i)	$r_\epsilon(M)$	size(B_i)
ram head	1280	640	13	19	25	45	43	70
	2560	1280	13	18	25	45	43	71
	5120	2560	13	18	25	45	43	72
	10240	5120	13	19	25	45	43	72
sun-flower	1280	640	67	84	111	141	151	185
	2560	1280	83	117	159	187	213	251
	5120	2560	83	125	182	226	298	347
	10240	5120	83	123	187	237	328	380
	20480	10240	83	120	187	238	327	382

numerical solutions for these Laplace Dirichlet problems converge exponentially fast regardless of the complicated geometries and the solver is quite robust. For example, in view of Table V which corresponds to the seemingly complicated geometry in Fig. 9, 10 digits of accuracy can be achieved using only 5120 quadrature points.

6.5. Nearly optimal compression

In this section, we compare the exact numerical rank of the largest off-diagonal block of A with the approximation rank obtained from SMASH. The numerical results show that the approximation rank is nearly optimal in the sense that it differs from the exact numerical rank by a small constant that is roughly independent of the kernel and the matrix size.

Definition 6.1 (ϵ -rank)

Let $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r$ be singular values of a nonzero matrix A . Given a tolerance $\epsilon \in (0, 1)$, the relative ϵ -rank of a matrix A , denoted by $r_\epsilon(A)$, is the largest number i such that $\sigma_i \geq \epsilon \sigma_1$.

Consider the numerical examples in Section 6.4, where different curves give rise to different kernels according to (67). Let i, j be two children of the root node. We focus on the (largest) off-diagonal block $A_{i^{row} \times j^{col}}$.

We consider three tolerances: $\epsilon = 10^{-3}, 10^{-6}, 10^{-10}$. We list the size of $A_{i^{row} \times j^{col}}$, the exact ϵ -rank, and the approximation rank characterized by the size of B_i generator with $\text{size}(B_i) := \max(\text{row size}, \text{column size})$. The results are reported in Table VI.

Note that no a priori information is needed to determine the approximation rank as it is solely derived from the prescribed tolerance and the construction algorithm. Thus the numerical results also imply that the proposed method in Section 6.1 for choosing parameters is satisfactory.

6.6. Storage cost

In this section, we demonstrate the benefit of the special structure in the generators produced by SMASH. We store the HSS generators using the strategy mentioned at the end of Section 2.3. For comparison, we also compute the cost by storing the generators as dense matrices, denoted by HSS_0 , as well as the storage cost for the original dense matrix. The test matrix A of order 10240 is derived from kernels in Section 6.4 and all entries are stored in double precision. The results are collected in Table VII for different geometries and different approximation accuracy. The reduction in storage justifies the use of strong rank-revealing QR algorithm in the construction and we see that SMASH is quite cheap even when a high approximation accuracy is in use.

Table VII. A comparison of storage costs (MB) of HSS generators the storage strategy in Section 2.3 (SMASH) and the standard approach by storing dense generators (HSS₀) for a square matrix A of order 10240.

ϵ_{far}	ϵ_{SVD}	geometry	storage(A)	HSS ₀	SMASH
10^{-4}	10^{-5}	ram head	800	5.9	4.4
10^{-4}	10^{-5}	sunflower	800	35.2	11.0
10^{-10}	10^{-11}	ram head	800	23.2	8.8
10^{-10}	10^{-11}	sunflower	800	145.7	31.3

7. CONCLUSION

We presented a unified framework, called SMASH, to construct either an $n \times n$ HSS or \mathcal{H}^2 matrix with an $O(n)$ cost. One appealing feature of this scheme is its simple implementation which only requires a routine to compress far field blocks. In addition, SMASH can greatly reduce the memory cost relative to existing analytic construction schemes. The numerical experiments illustrated the efficiency and robustness of SMASH through a few examples with various point distributions and kernel matrices.

We plan to extend this scheme to highly oscillatory kernels and to develop approximate inverse-type preconditioners for solving the resulting linear systems with \mathcal{H}^2 matrix representations.

ACKNOWLEDGMENTS

We would like to thank anonymous referees for their useful suggestions which led to substantial improvements of the original version of this paper. YX would like to thank Ming Gu for fruitful discussions about the strong rank revealing QR algorithm.

REFERENCES


1. Rokhlin V. Rapid solution of integral equations of classical potential theory. *Journal of Computational Physics* 1985; **60**(2):187–207, doi:[http://dx.doi.org/10.1016/0021-9991\(85\)90002-6](http://dx.doi.org/10.1016/0021-9991(85)90002-6). URL <http://www.sciencedirect.com/science/article/pii/0021999185900026>.
2. Greengard L, Rokhlin V. A fast algorithm for particle simulations. *J. Comput. Phys.* 1987; **73**:325–348.
3. Börm S. *Efficient numerical methods for non-local operators : \mathcal{H}^2 -matrix compression, algorithms and analysis*. EMS Tracts in Mathematics, European Mathematical Society: Zürich, 2010. URL <http://opac.inria.fr/record=b1133579>.
4. Börm S, Grasedyck L, Hackbusch W. Introduction to hierarchical matrices with applications. *ENG. ANAL. BOUND. ELEM.* 2003; **27**(5):405–422, doi:10.1016/S0955-7997(02)00152-2. URL <http://www.mis.mpg.de/de/publications/preprints/2002/prepr2002-18.html>.
5. Chandrasekaran S, Gu M, Pals T. A fast ULV decomposition solver for hierarchically semiseparable representations. *SIAM J. Matrix Anal. Appl.* 2006; **28**(3):603–622, doi:10.1137/S0895479803436652. URL <http://dx.doi.org/10.1137/S0895479803436652>.
6. Hackbusch W. *Hierarchical matrices : algorithms and analysis*. Springer, 2015.
7. Hackbusch W, Khoromskij B, Sauter S. On \mathcal{H}^2 -matrices. *Lectures on applied mathematics*, Bungartz HJ, Hoppe RHW, Zenger C (eds.). Springer: Berlin, 2000; 9–29. URL <http://www.mis.mpg.de/de/publications/preprints/1999/prepr1999-50.html>.
8. Xi Y, Xia J, Cauley S, Balakrishnan V. Superfast and stable structured solvers for toeplitz least squares via randomized sampling. *SIAM J. Matrix Anal. Appl.* 2014; **35**(1):44–72, doi:10.1137/120895755. URL <http://dx.doi.org/10.1137/120895755>.
9. Xia J. Multi-layer hierarchical structures and factorizations. *submitted to SIAM J. Matrix Anal. Appl.* ; .
10. Xia J, Xi Y, Gu M. A superfast structured solver for Toeplitz linear systems via randomized sampling. *SIAM J. Matrix Anal. Appl.* 2012; **33**(3):837–858, doi:10.1137/110831982. URL <http://dx.doi.org/10.1137/110831982>.
11. Bebendorf M. Approximation of boundary element matrices. *Numer. Math.* 2000; **86**(4):565–589, doi:10.1007/PL00005410. URL <http://dx.doi.org/10.1007/PL00005410>.

12. Gillman A, Young PM, Martinsson PG. A direct solver with $\mathcal{O}(n)$ complexity for integral equations on one-dimensional domains. *Frontiers of Mathematics in China* 2012; **7**(2):217–247, doi:10.1007/s11464-012-0188-3. URL <http://dx.doi.org/10.1007/s11464-012-0188-3>.
13. Ho KL, Greengard L. A fast direct solver for structured linear systems by recursive skeletonization. *SIAM Journal on Scientific Computing* 2012; **34**(5):A2507–A2532, doi:10.1137/120866683. URL <http://dx.doi.org/10.1137/120866683>.
14. L Ho K, Ying L. Hierarchical interpolative factorization for elliptic operators: Integral equations. *Communications on Pure and Applied Mathematics* 2016; **69**(7):1314–1353, doi:10.1002/cpa.21577. URL <http://dx.doi.org/10.1002/cpa.21577>.
15. Liu X, Xia J, de Hoop MV. Parallel randomized and matrix-free direct solvers for large structured dense linear systems. *SIAM J. Sci. Comput.* 2016; **38**(5):S508–S538, doi:10.1137/15M1023774. URL <http://dx.doi.org/10.1137/15M1023774>.
16. Martinsson P, Rokhlin V. A fast direct solver for boundary integral equations in two dimensions. *J. Comput. Phys.* 2005; **205**(1):1–23, doi:http://dx.doi.org/10.1016/j.jcp.2004.10.033.
17. Borne SL, Grasedyck L. H-matrix preconditioners in convection-dominated problems. *SIAM J. Matrix Anal. Appl.* 2006; **27**(4):1172–1183, doi:10.1137/040615845. URL <http://dx.doi.org/10.1137/040615845>.
18. Xia J. Efficient structured multifrontal factorization for general large sparse matrices. *SIAM J. Sci. Comput.* 2013; **35**(2):A832–A860, doi:10.1137/120867032. URL <http://dx.doi.org/10.1137/120867032>.
19. Grasedyck L. Existence of a low rank or \mathcal{H} -matrix approximant to the solution of a Sylvester equation. *Numer. Linear Algebra Appl.* 2004; **11**(4):371–389, doi:10.1002/nla.366. URL <http://dx.doi.org/10.1002/nla.366>.
20. Grasedyck L, Hackbusch W, Khoromskij NB. Solution of large scale algebraic matrix riccati equations by use of hierarchical matrices. *Computing* 2003; **70**(2):121–165, doi:10.1007/s00607-002-1470-0. URL <http://dx.doi.org/10.1007/s00607-002-1470-0>.
21. Benner P, Mach T. Computing all or some eigenvalues of symmetric \mathcal{H}_l -matrices. *SIAM J. Sci. Comput.* 2012; **34**(1):A485–A496, doi:10.1137/100815323. URL <http://dx.doi.org/10.1137/100815323>.
22. Xi Y, Xia J, Chan R. A fast randomized eigensolver with structured ldl factorization update. *SIAM J. Matrix Anal. Appl.* 2014; **35**(3):974–996, doi:10.1137/130914966. URL <http://dx.doi.org/10.1137/130914966>.
23. Hackbusch W, Börm S. Data-sparse approximation by adaptive \mathcal{H}^2 -matrices. *Computing* 2002; **69**(1):1–35, doi:10.1007/s00607-002-1450-4. URL <http://www.mis.mpg.de/de/publications/preprints/2001/prepr2001-86.html>.
24. Xi Y, Xia J. On the stability of some hierarchical rank structured matrix algorithms. *SIAM J. Matrix Anal. Appl.* 2016; **37**(3):1279–1303.
25. Xia J, Chandrasekaran S, Gu M, Li XS. Fast algorithms for hierarchically semiseparable matrices. *Numer. Linear Algebra Appl.* 2010; **17**(6):953–976, doi:10.1002/nla.691. URL <http://dx.doi.org/10.1002/nla.691>.
26. Gu M, Eisenstat SC. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM J. Sci. Comput.* 1996; **17**(4):848–869, doi:10.1137/0917055. URL <http://dx.doi.org/10.1137/0917055>.
27. Börm S. Data-sparse approximation of non-local operators by \mathcal{H}^2 -matrices. *Linear Algebra Appl.* 2007; **422**(2):380–403, doi:http://dx.doi.org/10.1016/j.laa.2006.10.021. URL <http://www.sciencedirect.com/science/article/pii/S0024379506004745>.
28. Bebendorf M, Rjasanow S. Adaptive low-rank approximation of collocation matrices. *Computing* 2003; **70**(1):1–24, doi:10.1007/s00607-002-1469-6. URL <http://dx.doi.org/10.1007/s00607-002-1469-6>.
29. Bebendorf M, Venn R. Constructing nested bases approximations from the entries of non-local operators. *Numer. Math.* 2012; **121**(4):609–635, doi:10.1007/s00211-012-0449-9. URL <http://dx.doi.org/10.1007/s00211-012-0449-9>.
30. Börm S, Grasedyck L. Hybrid cross approximation of integral operators. *Numer. Math.* 2005; **101**(2):221–249, doi:10.1007/s00211-005-0618-1. URL <http://dx.doi.org/10.1007/s00211-005-0618-1>.
31. Tyrtshnikov E. Mosaic-skeleton approximations. *CALCOLO* Jun 1996; **33**(1):47–57.
32. Goreinov S, Tyrtshnikov EE, Yereimin AY. Matrix-free iterative solution strategies for large dense linear systems. *Numer. Linear Algebra Appl.* 1997; **4**(4):273–294, doi:10.1002/(SICI)1099-1506(199707/08)4:4<273::AID-NLA97>3.0.CO;2-T. URL [http://dx.doi.org/10.1002/\(SICI\)1099-1506\(199707/08\)4:4<273::AID-NLA97>3.0.CO;2-T](http://dx.doi.org/10.1002/(SICI)1099-1506(199707/08)4:4<273::AID-NLA97>3.0.CO;2-T).
33. Goreinov S, Tyrtshnikov E, Zamarashkin N. A theory of pseudoskeleton approximations. *Linear Algebra Appl.* 1997; **261**(1):1–21, doi:https://doi.org/10.1016/S0024-3795(96)00301-1. URL <http://www.sciencedirect.com/science/article/pii/S0024379596003011>.
34. Tyrtshnikov E. Incomplete cross approximation in the mosaic-skeleton method. *Computing* Jun 2000; **64**(4):367–380.
35. Pan VY, Zhao L. Low-rank approximation of a matrix: Novel insights, new progress, and extensions. *arXiv* 2015; doi:http://arxiv.org/abs/1510.06142.
36. Pan VY, Luan Q, Svadlenka J, Zhao L. Primitive and cynical low rank approximation, preprocessing and extensions. *arXiv* 2016; doi:http://arxiv.org/abs/1611.01391.
37. Hackbusch W, Börm S. \mathcal{H}^2 -matrix approximation of integral operators by interpolation. *Appl. Numer. Math.* 2002; **43**(1):129–143.
38. Hackbusch W, Khoromskij BN, Kriemann R. Hierarchical matrices based on a weak admissibility criterion. *Computing* 2004; **73**(3):207–243, doi:10.1007/s00607-004-0080-4. URL <http://dx.doi.org/10.1007/s00607-004-0080-4>.
39. Anderson CR. An implementation of the fast multipole method without multipoles. *SIAM J. Sci. Statist. Comput.* 1992; **13**(4):923–947, doi:10.1137/0913055. URL <http://dx.doi.org/10.1137/0913055>.
40. Ying L, Biros G, Zorin D. A kernel-independent adaptive fast multipole algorithm in two and three dimensions. *J. Comput. Phys.* 2004; **196**(2):591–626, doi:http://dx.doi.org/10.1016/j.jcp.2003.11.021. URL <http://www.sciencedirect.com/science/article/pii/S0021999103006090>.

41. Lin L, Lu J, Ying L. Fast construction of hierarchical matrix representation from matrixvector multiplication. *J. Comput. Phys.* 2011; **230**(10):4071–4087, doi:<http://dx.doi.org/10.1016/j.jcp.2011.02.033>. URL <http://www.sciencedirect.com/science/article/pii/S0021999111001227>.
42. Martinsson PG. A fast randomized algorithm for computing a hierarchically semiseparable representation of a matrix. *SIAM J. Matrix Anal. Appl.* 2011; **32**(4):1251–1274, doi:10.1137/100786617. URL <http://dx.doi.org/10.1137/100786617>.
43. Rouet FH, Li XS, Ghysels P, Napov A. A distributed-memory package for dense hierarchically semi-separable matrix computations using randomization. *submitted to ACM Trans. Math. Softw.* ; .
44. Barnes J, Hut P. A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature* Dec 1986; **324**:446–449, doi:10.1038/324446a0.
45. Carrier J, Greengard L, Rokhlin V. A fast adaptive multipole algorithm for particle simulations. *SIAM Journal on Scientific and Statistical Computing* 1988; **9**(4):669–686, doi:10.1137/0909044.
46. Börm S. Construction of data-sparse \mathcal{H}^2 -matrices by hierarchical compression. *SIAM J. Sci. Comput.* 2009; **31**(3):1820–1839, doi:10.1137/080720693. URL <http://dx.doi.org/10.1137/080720693>.
47. Sun X, Pitsianis N. A matrix version of the fast multipole method. *SIAM Rev.* Feb 2001; **43**(2):289–300, doi:10.1137/S0036144500370835. URL <http://dx.doi.org/10.1137/S0036144500370835>.
48. Bebendorf M. *Hierarchical Matrices: A Means to Efficiently Solve Elliptic Boundary Value Problems*. Lecture Notes in Computational Science and Engineering, Springer Berlin Heidelberg, 2008. URL <https://books.google.com/books?id=hnIJOwyz9Z4C>.
49. Cai D, Xia J. A stable and efficient matrix version of the fast multipole method. *to be submitted.* ; .
50. Halko N, Martinsson PG, Tropp J. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review* 2011; **53**.
51. Gu M. Subspace iteration randomization and singular value problems. *preprint* ; URL <https://arxiv.org/pdf/1408.2208v1.pdf>.
52. Börm S. Approximation of integral operators by \mathcal{H}^2 -matrices with adaptive bases. *Computing* May 2005; **74**(3):249–271, doi:10.1007/s00607-004-0106-y. URL <https://doi.org/10.1007/s00607-004-0106-y>.
53. Börm S, Christophersen S. Approximation of integral operators by green quadrature and nested cross approximation. *Numerische Mathematik* Jul 2016; **133**(3):409–442, doi:10.1007/s00211-015-0757-y. URL <https://doi.org/10.1007/s00211-015-0757-y>.
54. Chandrasekaran S, Gu M, Lyons W. A fast adaptive solver for hierarchically semiseparable representations. *CALCOLO* 2005; **42**(3):171–185, doi:10.1007/s10092-005-0103-3. URL <http://dx.doi.org/10.1007/s10092-005-0103-3>.
55. Pan VY. *Structured Matrices and Polynomials: Unified Superfast Algorithms*. Springer-Verlag New York, Inc.: New York, NY, USA, 2001.
56. Pan VY. Transformations of matrix structures work again. *Linear Algebra and its Applications* 2015; **465**:107–138, doi:<http://dx.doi.org/10.1016/j.laa.2014.09.004>. URL <http://www.sciencedirect.com/science/article/pii/S0024379514005898>.
57. Pan VY. *Fast Approximate Computations with Cauchy Matrices, Polynomials and Rational Functions*. Springer International Publishing: Cham, 2014; 287–299, doi:10.1007/978-3-319-06686-8_22. URL http://dx.doi.org/10.1007/978-3-319-06686-8_22.
58. Pan VY, Wang X. Inversion of displacement operators. *SIAM Journal on Matrix Analysis and Applications* 2003; **24**(3):660–677, doi:10.1137/S089547980238627X. URL <http://dx.doi.org/10.1137/S089547980238627X>.
59. Pan VY. How Bad Are Vandermonde Matrices? *SIAM J. Matrix Anal. Appl.* 2016; **37**(2):676–694, doi:10.1137/15M1030170.
60. Pan VY. Fast approximate computations with cauchy matrices and polynomials. *Math. Comput.* 2017; **86**(308):2799–2826, doi:10.1090/mcom/3204. URL <https://doi.org/10.1090/mcom/3204>.
61. Calvetti D, Reichel L. Factorizations of Cauchy matrices. *J. Comput. Appl. Math.* 1997; **86**(1):103–123, doi:[http://dx.doi.org/10.1016/S0377-0427\(97\)00150-7](http://dx.doi.org/10.1016/S0377-0427(97)00150-7). URL <http://www.sciencedirect.com/science/article/pii/S0377042797001507>.
62. Bini DA, Meini B, Poloni F. Fast solution of a certain riccati equation through cauchy-like matrices. *ETNA. Electronic Transactions on Numerical Analysis [electronic only]* 2008; **33**:84–104. URL <http://eudml.org/doc/223449>.
63. Boros T, Kailath T, Olshevsky V. A fast parallel BjörckPereyra-type algorithm for solving Cauchy linear equations. *Linear Algebra Appl.* 1999; **302**:265–293, doi:[http://dx.doi.org/10.1016/S0024-3795\(99\)00115-9](http://dx.doi.org/10.1016/S0024-3795(99)00115-9). URL <http://www.sciencedirect.com/science/article/pii/S0024379599001159>.
64. Boros T, Kailath T, Olshevsky V. Special issue on structured and infinite systems of linear equations pivoting and backward stability of fast algorithms for solving Cauchy linear equations. *Linear Algebra Appl.* 2002; **343**:63–99, doi:[http://dx.doi.org/10.1016/S0024-3795\(01\)00519-5](http://dx.doi.org/10.1016/S0024-3795(01)00519-5). URL <http://www.sciencedirect.com/science/article/pii/S0024379501005195>.
65. Gohbert I, Kailath T, Olshevsky V. Fast Gaussian elimination with partial pivoting for matrices with displacement structure. *Math. Comput.* Oct 1995; **64**(212):1557–1576, doi:10.2307/2153371. URL <http://dx.doi.org/10.2307/2153371>.
66. Gu M. Stable and efficient algorithms for structured systems of linear equations. *SIAM J. Matrix Anal. Appl.* 1998; **19**(2):279–306, doi:10.1137/S0895479895291273. URL <http://dx.doi.org/10.1137/S0895479895291273>.
67. Pan VY, Zheng A. Superfast algorithms for Cauchy-like matrix computations and extensions. *Linear Algebra Appl.* 2000; **310**(1):83–108, doi:[http://dx.doi.org/10.1016/S0024-3795\(00\)00041-0](http://dx.doi.org/10.1016/S0024-3795(00)00041-0). URL <http://www.sciencedirect.com/science/article/pii/S0024379500000410>.

68. Xia J. On the complexity of some hierarchical structured matrix algorithms. *SIAM Journal on Matrix Analysis and Applications* 2012; **33**(2):388–410, doi:10.1137/110827788. URL <http://dx.doi.org/10.1137/110827788>.
69. Hsiao GC, Wendland WL. *Boundary integral equations*. Applied Mathematical Sciences, Springer: Berlin, Heidelberg, 2008.
70. Kress R. *Linear Integral Equations*. Applied Mathematical Sciences, Springer New York, 2013. URL <https://books.google.com/books?id=-L-8BAAAQBAJ>.
71. Atkinson K. *The Numerical Solution of Integral Equations of the Second Kind*. Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, 1997.

Figure 10. Illustration of the sets \mathcal{N}_i used in HSS constructions.



APPENDIX: DETERMINATION OF NEARFIELD BLOCKS

In building HSS structures, we need to collect nodes corresponding to blocks that are nearfield with respect to node i . The set of such nodes is denoted by \mathcal{N}_i for node i . We set $\mathcal{N}_i = \emptyset$ when $i = root$. For the other cases, \mathcal{N}_i is defined below where p_i denotes the parent of i :

$$\mathcal{N}_i = \{k \in \mathcal{T} \text{ such that : } \begin{array}{l} \text{either } k \text{ is a sibling of } i \\ \text{or } p_k \in \mathcal{N}_{p_i} \text{ and } (i, k) \notin \text{Sep} \\ \text{or } k \text{ is a leaf such that } k \in \mathcal{N}_{p_i} \text{ and } (i, k) \notin \text{Sep} \end{array} \} \quad (70)$$

Note that when i is a child of root then \mathcal{N}_{p_i} is empty and so only the first case can take place (k is a sibling of i). We also remark that the third case (k is a leaf such that $k \in \mathcal{N}_{p_i}$ and $(i, k) \notin \text{Sep}$) is required for non-uniform distributions and that it is empty if the distribution is uniform. It is easy to see that if $\Omega = [0, 1]$ and $X = Y$ is uniformly distributed, \mathcal{T} is a perfect binary tree. In addition, if the separation ratio is set to $\tau = 0.5$, then for any nonroot node i , \mathcal{N}_i contains at most two nodes.

For each node i , let \mathbf{i} denote the index set of the points in $X \cap \Omega_i$. Similarly, \mathbf{j} represents the index set of the points in $Y \cap \Omega_j$. Namely, $X_{\mathbf{i}} = X \cap \Omega_i$ and $Y_{\mathbf{j}} = Y \cap \Omega_j$.

Remark .1

Since the HSS structure [54, 5] is associated with a binary tree regardless of the dimension of the problem (see Section 3.2), to construct HSS matrices, bisection is used throughout the adaptive partitioning procedure. For example, given a domain or a curve enclosed in a square in \mathbb{R}^2 , we use bisection in the horizontal direction and the vertical direction alternatively in consecutive stages of the adaptive partitioning, i.e., if horizontal bisection is used at partitioning level l , then vertical bisection will be employed at level $l + 1$. The numerical experiments in Section 6.3 and Section 6.4 provide illustrations. This partitioning strategy corresponds to the *geometrically regular* clustering (cf.[3]), and can be generalized into the *geometrically balanced* clustering (cf.[3]).