

A Superfast Direct Solver for 2D Type-II Inverse Nonuniform Discrete Fourier Transform Based on Hierarchically Semiseparable Matrix

Yingzhou Li*, Jingyu Liu†

July 2, 2026

Abstract

This paper proposes a direct inversion method for the 2D type-II nonuniform discrete Fourier transform (NUDFT). The NUDFT matrix \mathbf{A} is factored as $\mathbf{A} = \mathbf{GF}$, where \mathbf{G} can be expressed as a kernel matrix and \mathbf{F} is the 2D DFT matrix. We show that \mathbf{G} can be approximated by a hierarchically semiseparable (HSS) matrix and give an estimate of the HSS rank. Then, using the least-squares solver for HSS matrix and the two-dimensional inverse fast Fourier transform, the inverse NUDFT problem can be solved efficiently. Our algorithm has an offline complexity of $\mathcal{O}(M + N^{3/2} \log^3 N)$ where M and N are the size of rows and columns of the NUDFT matrix, respectively. Once the direct solver is built, it can be applied to a vector with an online complexity of $\mathcal{O}(M + N \log^3 N)$. The proposed method can be used as a preconditioner for iterative methods, especially when the sample points are distributed on a grid such that \mathbf{A} is ill-conditioned. Numerical results are provided to show the scaling performance of the inversion method and demonstrate the efficiency and robustness of it as a preconditioner.

Keywords nonuniform discrete Fourier transform, hierarchically semi-separable matrix

1 Introduction

This paper considers the two-dimensional type-II *nonuniform discrete Fourier transform* (NUDFT) of the following form:

$$f_j = \sum_{k^{[x]}=0}^{n^{[x]}-1} \sum_{k^{[y]}=0}^{n^{[y]}-1} c_{k^{[x]}, k^{[y]}} e^{-2\pi i(k^{[x]}x_j + k^{[y]}y_j)}, \quad 0 \leq j \leq M-1, \quad (1.1)$$

where the *sample points* $\{(x_j, y_j)\}$ are distributed arbitrarily in $[0, 1)^2$ and the *frequencies* $\{(k^{[x]}, k^{[y]})\}$ are distributed on a Cartesian grid of contiguous integers. Throughout the paper, the number of samples points M is assumed to be larger than or equal to the number of frequencies, i.e., $N = n^{[x]}n^{[y]}$. When the *coefficients* $\{c_{k^{[x]}, k^{[y]}}\}$ are given and one aims to compute the *target values* $\{f_j\}$, the problem is called the *forward*

*School of Mathematical Sciences, Fudan University; Shanghai Key Laboratory for Contemporary Applied Mathematics, Fudan University, yingzhouli@fudan.edu.cn

†School of Mathematical Sciences, Fudan University, jyliu22@m.fudan.edu.cn

NUDFT. When the target values $\{f_j\}$ are given and one aims to determine the coefficients $\{c_{k^{[x]},k^{[y]}}\}$, the problem is called the *inverse* NUDFT. Both the forward and inverse NUDFT have wide applications in scientific computing, such as signal processing [2], image reconstruction [7], and fast convolution [19] etc. Let the NUDFT matrix \mathbf{A} be defined as¹

$$\mathbf{A}(j, (k^{[x]}, k^{[y]})) = e^{-2\pi i(k^{[x]}x_j + k^{[y]}y_j)}, \quad 0 \leq j \leq M - 1, 0 \leq k^{[x]} \leq n^{[x]} - 1, 0 \leq k^{[y]} \leq n^{[y]} - 1, \quad (1.2)$$

where the column index $(k^{[x]}, k^{[y]})$ is ordered in a lexicographical manner. That is, we have an underlying 1-to-1 map $(k^{[x]}, k^{[y]}) \leftrightarrow k^{[y]} + k^{[x]}n^{[y]}$. The forward NUDFT can be expressed as a matrix-vector multiplication $\mathbf{f} = \mathbf{A}\mathbf{c}$ and the inverse NUDFT can be formulated as a linear least-squares problem $\min_{\mathbf{c}} \|\mathbf{A}\mathbf{c} - \mathbf{f}\|_2$, where $\mathbf{c} \in \mathbb{C}^N$ is the vectorized form of the coefficients $\{c_{k^{[x]},k^{[y]}}\}$ and $\mathbf{f} \in \mathbb{C}^M$ is the vector corresponding to the target values $\{f_j\}$. Typically, the forward and inverse NUDFT have a complexity of $\mathcal{O}(MN)$ and $\mathcal{O}(MN^2)$ respectively.

1.1 Related Work

If the sample points are distributed on a uniformly Cartesian grid, i.e., $(x_{j^{[x]}}, y_{j^{[y]}}) = (j^{[x]}/n^{[x]}, j^{[y]}/n^{[y]})$ for $0 \leq j^{[x]} \leq n^{[x]} - 1$ and $0 \leq j^{[y]} \leq n^{[y]} - 1$, the NUDFT reduces to the *discrete Fourier transform* (DFT). Consequently, the NUDFT matrix \mathbf{A} has a Kronecker product structure $\mathbf{A} = \mathbf{A}^{[x]} \otimes \mathbf{A}^{[y]}$, where $\mathbf{A}^{[x]}(j^{[x]}, k^{[x]}) = e^{-2\pi i k^{[x]} j^{[x]}/n^{[x]}}$ and $\mathbf{A}^{[y]}(j^{[y]}, k^{[y]}) = e^{-2\pi i k^{[y]} j^{[y]}/n^{[y]}}$ are the DFT matrices in the x and y directions, respectively. In this case, the forward NUDFT can be efficiently computed by the *fast Fourier transform* (FFT) algorithm [11] in $\mathcal{O}(N \log N)$ time. And the inverse NUDFT can also be efficiently solved by the *inverse fast Fourier transform* (iFFT) algorithm in $\mathcal{O}(N \log N)$ time. More generally, if $(x_{j^{[x]}}, y_{j^{[y]}}) = (j^{[x]}/m^{[x]}, j^{[y]}/m^{[y]})$ where $m^{[x]} \geq n^{[x]}$ and $m^{[y]} \geq n^{[y]}$, the matrix \mathbf{A} is a submatrix of the DFT matrix, and both the forward and inverse NUDFT can still be efficiently computed by the FFT algorithm in $\mathcal{O}(M \log M)$ time.

Unfortunately, when the sample points are not distributed on a uniformly Cartesian grid, FFT cannot be directly applied. Many algorithms are developed to compute the forward NUDFT efficiently, which are usually referred to as the *nonuniform fast Fourier transform* (NUFFT) algorithms [1, 3, 13, 14, 16, 22, 28, 29]. Commonly, the NUFFT algorithms have a complexity of $\mathcal{O}(M + N \log N)$.

In the nonuniform case, the pseudoinverse of the NUDFT matrix is no longer a scaled version of its adjoint. Consequently, specialized algorithms are required to solve the associated inverse problem efficiently. A widely adopted approach is the iterative method. Thanks to the efficient NUFFT algorithms for the forward computation, the inverse NUDFT is often solved by the *conjugate gradient* (CG) method. For example, using the MATLAB command “lsqr” with the NUFFT algorithm as the matrix-vector multiplication subroutine [27]. When the sample points are well-distributed, iterative methods usually converge in a small number of iterations. However, when the sample points are distributed in a highly nonuniform manner, which is often the case in practical applications, CG may require a large number of iterations to converge or even fail to converge. Preconditioners are needed to accelerate the convergence of iterative methods in this case.

Direct inversion methods are also developed. For example, Kircheis and Potts [20] proposed a direct inversion method for the type-II NUDFT in 1D case and they extended it to the 2D case [21]. Recently, in [30], the authors proposed a direct method to solve the type-II NUDFT in 1D based on the *hierarchically semiseparable* (HSS) matrix, which has a complexity of $\mathcal{O}((M + N) \log^2 N)$. Further, the authors in [23] extended this method to the case of type-III NUDFT in 1D, by factoring the type-III NUDFT matrix into a product of a type-II NUDFT matrix and an HSS matrix. For the 2D problem, if the sample points are distributed on a tensor grid, the NUDFT matrix has a Kronecker product structure, and one can apply the 1D solver in each direction to solve the problem. However, the extension of the method [30] to the general 2D case is nontrivial.

¹In this paper, the starting index of a matrix can be 0 or 1, depending on the context.

1.2 Contributions

In this paper, we propose a direct inversion method for the 2D type-II NUFT. Let $\mathbf{F} \in \mathbb{C}^{N \times N}$ be the 2D DFT matrix. We first derive an explicit formula for the entries of the matrix $\mathbf{G} = \mathbf{A}\mathbf{F}^{-1}$, showing that it can be expressed as a kernel matrix. Based on this expression, we analyze the low-rank structure of \mathbf{G} and show that it can be efficiently compressed into an HSS matrix.

The matrix \mathbf{G} can also be expressed as a face-splitting product of two matrices $\mathbf{G}^{[x]}$ and $\mathbf{G}^{[y]}$, where $\mathbf{G}^{[x]}$ and $\mathbf{G}^{[y]}$ are the transformed NUFT matrices in the x - and y - directions, respectively. Since $\mathbf{G}^{[x]}$ and $\mathbf{G}^{[y]}$ can be approximated by HSS matrices, this representation suggests that \mathbf{G} can also be approximated by an HSS matrix. To make this argument precise, we define the face-splitting product of two HSS trees and prove that this product is again an HSS tree. We further prove that the face-splitting product of two HSS matrices remains an HSS matrix, and provide an estimate for the resulting HSS rank. Finally, we establish an approximation error bound for the HSS approximation of \mathbf{G} .

Using the kernel matrix expression of \mathbf{G} , the HSS construction method based on the interpolative decomposition [10] and the proxy surface technique [25] can be applied to approximate \mathbf{G} by an HSS matrix $\tilde{\mathbf{G}}$ efficiently. After the construction, the URV factorization [31] is used as a direct solver for the least-squares problem associated with $\tilde{\mathbf{G}}$. The construction and factorization of the HSS matrix constitute the *offline* stage, whose complexity is $\mathcal{O}(M + N^{3/2} \log^3 N)$. Once the solver is built, the inverse NUFT problem can be solved by applying the HSS solver followed by the iFFT. This is referred to as the *online* stage, and its complexity is $\mathcal{O}(M + N \log^3 N)$.

The proposed method can also be used as a preconditioner for iterative methods, which is particularly useful when the sample points are distributed in a highly nonuniform manner. Numerical experiments agree with our theoretical analysis and demonstrate the efficiency of the proposed method.

1.3 Organization

The rest of the paper is organized as follows. In Section 2, we introduce some notations and preliminaries that will be used in the later paper. We derive the low-rank structure of the NUFT matrix in Section 3 and propose our direct solver in Section 4. In Section 5, we present numerical results to demonstrate the efficiency of the proposed method. Finally, we conclude the paper in Section 6 and discuss possible future directions.

2 Preliminaries

2.1 Notations and Preliminaries

We use $\mathbb{S} := \{z \in \mathbb{C} : |z| = 1\}$ to denote the unit circle. For a set \mathcal{J} , we use $|\mathcal{J}|$ to denote the cardinality of \mathcal{J} . We use MATLAB notation to denote the submatrix of \mathbf{A} with row and column indices in \mathcal{R} and \mathcal{C} , i.e., $\mathbf{A}(\mathcal{R}, \mathcal{C})$. A colon in the index means all the indices in that dimension, e.g., $\mathbf{A}(:, \mathcal{C})$ means the submatrix of \mathbf{A} with all rows and column indices in \mathcal{C} .

Let $\|\cdot\|$ be a unitarily invariant matrix norm. For a given tolerance $0 < \varepsilon < 1$, the numerical rank of a matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$, denoted by $\text{rank}_\varepsilon(\mathbf{A})$, is defined as the smallest integer r such that there exists a matrix $\tilde{\mathbf{A}} \in \mathbb{C}^{m \times n}$ with $\|\mathbf{A} - \tilde{\mathbf{A}}\| < \varepsilon \|\mathbf{A}\|$ and $\text{rank}(\tilde{\mathbf{A}}) = r$.

For two matrices $\mathbf{A}_1 \in \mathbb{C}^{m \times n_1}$ and $\mathbf{A}_2 \in \mathbb{C}^{m \times n_2}$ of the same size of rows, the *face-splitting product* of \mathbf{A}_1 and \mathbf{A}_2 is defined as the Kronecker product of the corresponding rows of \mathbf{A}_1 and \mathbf{A}_2 , i.e.,

$$\mathbf{A}_1 \bullet \mathbf{A}_2 := \begin{bmatrix} \mathbf{A}_1(1, :) \otimes \mathbf{A}_2(1, :) \\ \mathbf{A}_1(2, :) \otimes \mathbf{A}_2(2, :) \\ \vdots \\ \mathbf{A}_1(m, :) \otimes \mathbf{A}_2(m, :) \end{bmatrix} \in \mathbb{C}^{m \times (n_1 n_2)}.$$

The Frobenius norm of the face-splitting product can be estimated by the following lemma, which can be proved through direct calculation.

Lemma 2.1. *Let $\mathbf{A}_1 \in \mathbb{C}^{m \times n_1}$ and $\mathbf{A}_2 \in \mathbb{C}^{m \times n_2}$, then*

$$\|\mathbf{A}_1 \bullet \mathbf{A}_2\|_{\mathbb{F}} \leq \min \left\{ \max_{1 \leq j \leq m} \{\|\mathbf{A}_1(j, :)\|_2\} \|\mathbf{A}_2\|_{\mathbb{F}}, \max_{1 \leq j \leq m} \{\|\mathbf{A}_2(j, :)\|_2\} \|\mathbf{A}_1\|_{\mathbb{F}} \right\}.$$

2.2 HSS Matrices

HSS matrices are a class of hierarchical matrices [6, 18, 32]. Since the storage and algebraic operations for HSS matrices typically have linear complexity, they have been widely used in scientific computing [12, 26]. We first introduce the definition of HSS tree, which is the underlying data structure for HSS matrices.

Definition 2.2 (HSS tree, [31]). *A tree \mathbb{T} is called an HSS tree with row and column indices \mathcal{J} and \mathcal{K} if each node τ of \mathbb{T} is associated with two index sets \mathcal{J}_τ and \mathcal{K}_τ , satisfying the following conditions*

- (1) $\mathcal{J}_\rho = \mathcal{J}, \mathcal{K}_\rho = \mathcal{K}$ for the root node ρ .
- (2) For a nonleaf node τ , we have $\mathcal{J}_\tau = \sqcup_{\alpha \in \text{ch}(\tau)} \mathcal{J}_\alpha$ and $\mathcal{K}_\tau = \sqcup_{\alpha \in \text{ch}(\tau)} \mathcal{K}_\alpha$, where the notation \sqcup means the disjoint union.

For a node τ , we use $\text{ch}(\tau)$ and $\text{sib}(\tau)$ to denote the set of children of τ and the set of siblings of τ , respectively. For an HSS tree \mathbb{T} , we define the level of each node as follows: For the root node ρ , $\text{level}(\rho) = 0$. If τ is a nonleaf node and $\alpha \in \text{ch}(\tau)$, then $\text{level}(\alpha) = \text{level}(\tau) + 1$. The maximum level of the tree is denoted by L .

Definition 2.3 (HSS matrix, [31]). *Let \mathbb{T} be an HSS tree of row and column indices \mathcal{J} and \mathcal{K} . A matrix $\mathbf{H} \in \mathbb{C}^{|\mathcal{J}| \times |\mathcal{K}|}$ is called an HSS matrix about \mathbb{T} if there are matrices $\mathbf{D}_\tau, \mathbf{U}_\tau, \mathbf{V}_\tau, \mathbf{R}_\tau, \mathbf{W}_\tau$ and $\mathbf{B}_{\tau, \sigma}$ (called HSS generators) associated with each node τ , which satisfy the following conditions:*

- (1) For a leaf node τ , $\mathbf{D}_\tau = \mathbf{H}(\mathcal{J}_\tau, \mathcal{K}_\tau)$ is a dense matrix.
- (2) For a nonleaf node τ with children $\{\alpha_p\}_{p=1}^b$, $\mathbf{D}_\tau = \mathbf{H}(\mathcal{J}_\tau, \mathcal{K}_\tau)$ is a $b \times b$ block matrix, with the (p, q) -th block $\mathbf{D}_{\tau; p, q}$ being \mathbf{D}_{α_p} for $p = q$ and $\mathbf{U}_{\alpha_p} \mathbf{B}_{\alpha_p, \alpha_q} \mathbf{V}_{\alpha_q}^*$ for $p \neq q$. The matrices \mathbf{U}_τ and \mathbf{V}_τ is a $b \times 1$ block matrix, with the p -th block being $\mathbf{U}_{\alpha_p} \mathbf{R}_{\alpha_p}$ and $\mathbf{V}_{\alpha_p} \mathbf{W}_{\alpha_p}$, respectively.

We call \mathbf{U}_τ and \mathbf{V}_τ^* the basis matrices, \mathbf{R}_τ and \mathbf{W}_τ the transfer matrices, and $\mathbf{B}_{\tau, \sigma}$ the interaction matrices.

One important property of HSS matrices, called the *shared basis property*, is that the basis matrices \mathbf{U}_τ and \mathbf{V}_τ^* are bases for the column and row spaces of the corresponding off-diagonal blocks (called the *HSS blocks*) $\mathbf{H}(\mathcal{J}_\tau, \mathcal{K}_\tau^c)$ and $\mathbf{H}(\mathcal{J}_\tau^c, \mathcal{K}_\tau)$ of \mathbf{H} , respectively. Here $\mathcal{J}_\tau^c = \mathcal{J} \setminus \mathcal{J}_\tau$ and $\mathcal{K}_\tau^c = \mathcal{K} \setminus \mathcal{K}_\tau$. the *HSS rank* r of \mathbf{H} is defined as the maximum rank among all the HSS blocks.

Throughout this paper, we assume that the HSS matrix is associated with a given HSS tree \mathbb{T} with maximum level L . For simplicity, we also assume that \mathbb{T} is a full tree, i.e., all the leaf nodes are on the same level L , and that the HSS row and column blocks of a node τ have a same rank r_τ .

2.3 The URV Factorization for HSS Matrices

Consider the least-squares problem $\min_{\mathbf{c}} \|\mathbf{H}\mathbf{c} - \mathbf{f}\|_2$, where \mathbf{H} is an HSS matrix and \mathbf{f} is a given vector. In this section, we review the URV factorization for HSS matrices [30, 31], which serves as a direct solver for the least-squares problem.

The URV factorization for HSS matrices is performed in a bottom-up manner on the HSS tree. Suppose τ is a leaf node and $\mathbf{D}_\tau \in \mathbb{C}^{m_\tau \times n_\tau}$. If $m_\tau \gg n_\tau$, a *size reduction* step is performed by QR factorization:

$$[\mathbf{U}_\tau \quad \mathbf{D}_\tau] = \Omega_\tau \begin{bmatrix} \mathbf{R}_{1,1} & \mathbf{R}_{1,2} \\ \mathbf{0} & \mathbf{R}_{2,2} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} = \Omega_\tau \begin{bmatrix} \check{\mathbf{U}}_\tau & \check{\mathbf{D}}_\tau \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad (2.1)$$

where $\mathbf{\Omega}_\tau \in \mathbb{C}^{m_\tau \times m_\tau}$ is unitary. The reduced row size, i.e., the number of rows of $\check{\mathbf{U}}_\tau$ and $\check{\mathbf{D}}_\tau$, is $\check{m}_\tau = n_\tau + r_\tau$. If no size reduction is adapted, we set $\mathbf{\Omega}_\tau$ to be the empty and let $\check{\mathbf{U}}_\tau = \mathbf{U}_\tau$, $\check{\mathbf{D}}_\tau = \mathbf{D}_\tau$ and $\check{m}_\tau = n_\tau + r_\tau$.

The next is a *basis elimination* step. For every leaf node τ , zeros are introduced into \mathbf{V}_τ by

$$\mathbf{V}_\tau = \mathbf{Q}_\tau \begin{bmatrix} \mathbf{0} \\ \widehat{\mathbf{V}}_{\tau;2} \end{bmatrix}, \quad (2.2)$$

where $\mathbf{Q}_\tau \in \mathbb{C}^{n_\tau \times n_\tau}$ is unitary and $\widehat{\mathbf{V}}_{\tau;2} \in \mathbb{C}^{r_\tau \times r_\tau}$. This can be done by a reverse QR decomposition. The diagonal block is then modified by $\widetilde{\mathbf{D}}_\tau = \widehat{\mathbf{D}}_\tau \mathbf{Q}_\tau$

Then a *partial decomposition* on the matrix $\widetilde{\mathbf{D}}_\tau$ is performed to decouple the unknowns corresponding to the zero and dense parts of (2.2). First, a QR decomposition is computed on the first block column of $\widetilde{\mathbf{D}}_\tau$ as

$$\begin{bmatrix} \widetilde{\mathbf{D}}_{\tau;1,1} \\ \widetilde{\mathbf{D}}_{\tau;2,1} \end{bmatrix} = \mathbf{P}_\tau \begin{bmatrix} \widehat{\mathbf{D}}_{\tau;1,1} \\ \mathbf{0} \end{bmatrix}. \quad (2.3)$$

Then it follows that

$$\widetilde{\mathbf{D}}_\tau = \mathbf{P}_\tau \begin{bmatrix} \widehat{\mathbf{D}}_{\tau;1,1} & \widehat{\mathbf{D}}_{\tau;1,2} \\ \mathbf{0} & \widehat{\mathbf{D}}_{\tau;2,2} \end{bmatrix}, \quad (2.4)$$

where

$$\begin{bmatrix} \widehat{\mathbf{D}}_{\tau;1,2} \\ \widehat{\mathbf{D}}_{\tau;2,2} \end{bmatrix} = \mathbf{P}_\tau^* \begin{bmatrix} \widetilde{\mathbf{D}}_{\tau;1,2} \\ \widetilde{\mathbf{D}}_{\tau;2,2} \end{bmatrix}. \quad (2.5)$$

The basis matrix \mathbf{U}_τ is then modified by

$$\mathbf{P}_\tau^* \mathbf{U}_\tau = \begin{bmatrix} \widehat{\mathbf{U}}_{\tau;1} \\ \widehat{\mathbf{U}}_{\tau;2} \end{bmatrix}. \quad (2.6)$$

After that, the variables corresponding to the first block column of $\widetilde{\mathbf{D}}_\tau$ can be decoupled from the rest variables, and the least-squares problem can be solved by a recursive process, which we describe in the following. For a nonleaf node τ with children $\{\alpha_p\}$, appropriate blocks of the children are merged to form the new generators. Let \mathbf{D}_τ^+ , \mathbf{U}_τ^+ and \mathbf{V}_τ^+ be the block matrices given by

$$\begin{aligned} \mathbf{D}_{\tau;p,q}^+ &= \begin{cases} \widehat{\mathbf{D}}_{\alpha_p;2,2}, & p = q, \\ \widehat{\mathbf{U}}_{\alpha_p;2} \mathbf{B}_{\alpha_p, \alpha_q} \widehat{\mathbf{V}}_{\alpha_q;2}^*, & p \neq q \end{cases} \\ \mathbf{U}_{\tau;p}^+ &= \widehat{\mathbf{U}}_{\alpha_p;2} \mathbf{R}_{\alpha_p}, \quad \mathbf{V}_{\tau;p}^+ = \widehat{\mathbf{V}}_{\alpha_p;2} \mathbf{W}_{\alpha_p}. \end{aligned} \quad (2.7)$$

Then the children of τ can be “removed” and the node τ can be treated as a leaf node. The generators $\{\mathbf{D}_\tau^+\}$, $\{\mathbf{U}_\tau^+\}$, $\{\mathbf{V}_\tau^+\}$, $\{\mathbf{R}_\tau\}$, $\{\mathbf{W}_\tau\}$ and $\{\mathbf{B}_{\tau,\sigma}\}$ define a new HSS matrix of maximum level $L - 1$. This HSS can be further factorized by the same discussion above (with \mathbf{D}_τ , \mathbf{U}_τ and \mathbf{V}_τ replaced by \mathbf{D}_τ^+ , \mathbf{U}_τ^+ and \mathbf{V}_τ^+) until the root node is reached. At the root node ρ , the QR factorization is performed, i.e., $\mathbf{D}_\rho^+ = \mathbf{P}_\rho \widehat{\mathbf{D}}_\rho^+$.

Once the URV factors are computed, the solution is obtained by a bottom-up pass and a top-down pass. The bottom-up pass is to apply the unitary transformations $\mathbf{\Omega}_\tau$ and \mathbf{P}_τ to the right-hand side \mathbf{f} . Starting from every leaf node τ , if size reduction is performed, the vector \mathbf{f}_τ is updated by

$$\mathbf{\Omega}_\tau^* \mathbf{f}_\tau = \begin{bmatrix} \check{\mathbf{f}}_\tau \\ \times \end{bmatrix}, \quad \widehat{\mathbf{f}}_\tau = \mathbf{P}_\tau^* \check{\mathbf{f}}_\tau = \begin{bmatrix} \widehat{\mathbf{f}}_{\tau;1} \\ \widehat{\mathbf{f}}_{\tau;2} \end{bmatrix},$$

where the \times means the corresponding entries are not used in the following steps and can be ignored. If no size reduction is performed, the first equation is omitted and $\check{\mathbf{f}}_\tau$ is replaced by \mathbf{f}_τ in the second equation. Then for a nonleaf node τ with children $\{\alpha_p\}$, \mathbf{f}_τ^+ is obtained by merging its children blocks, i.e.,

$\mathbf{f}_{\tau;p}^+ = \widehat{\mathbf{f}}_{\alpha p;2}$. This process is repeated (with \mathbf{f}_τ replaced by \mathbf{f}_τ^+ in the above equations) until the root node is reached.

The top-down pass is to perform back substitution to obtain the solution \mathbf{c} . At the root node ρ , after solving $\widehat{\mathbf{D}}_\rho^+ \mathbf{c}_\rho^+ = \widehat{\mathbf{f}}_\rho$, for each child $\alpha \in \text{ch}(\rho)$, the vectors $\widehat{\mathbf{c}}_{\alpha;2}$ is assigned as the corresponding block in $\widehat{\mathbf{c}}_\rho^+$. Further, the vectors $\widehat{\mathbf{h}}_\alpha$, $\widehat{\mathbf{g}}_\alpha$ and $\widehat{\mathbf{b}}_\alpha$ are initialized by $\mathbf{0}$, $\sum_{\beta \in \text{sib}(\alpha)} \mathbf{B}_{\alpha,\beta} \widehat{\mathbf{V}}_{\beta;2} \widehat{\mathbf{c}}_{\beta;2}$ and $\widehat{\mathbf{h}}_\alpha + \widehat{\mathbf{g}}_\alpha$, respectively. For a nonroot node τ , the vector \mathbf{c}_τ^+ is updated by $\mathbf{c}_\tau^+ = \mathbf{Q}_\tau \begin{bmatrix} \widehat{\mathbf{c}}_{\tau;1} \\ \widehat{\mathbf{c}}_{\tau;2} \end{bmatrix}$, where $\widehat{\mathbf{c}}_{\tau;1} = \widehat{\mathbf{D}}_{\tau;1,1}^{-1} (\widehat{\mathbf{f}}_{\tau;1} - \widehat{\mathbf{D}}_{\tau;1,2} \widehat{\mathbf{c}}_{\tau;2} - \widehat{\mathbf{U}}_{\tau;1} \widehat{\mathbf{b}}_\tau)$. If τ is nonleaf, then for every child $\alpha \in \text{ch}(\tau)$, $\widehat{\mathbf{c}}_{\alpha;2}$ is assigned as the corresponding block in \mathbf{c}_τ^+ and $\widehat{\mathbf{h}}_\alpha = \mathbf{R}_\alpha \widehat{\mathbf{h}}_\tau$. The vectors $\widehat{\mathbf{g}}_\alpha$ and $\widehat{\mathbf{b}}_\alpha$ are given by $\widehat{\mathbf{g}}_\alpha = \sum_{\beta \in \text{sib}(\alpha)} \mathbf{B}_{\alpha,\beta} \widehat{\mathbf{V}}_{\beta;2} \widehat{\mathbf{c}}_\beta$ and $\widehat{\mathbf{b}}_\alpha = \widehat{\mathbf{h}}_\tau + \widehat{\mathbf{g}}_\tau$. This process is repeated until the leaf nodes are reached, where $\mathbf{c}_\tau = \mathbf{c}_\tau^+$ for each leaf node τ .

The complexity of these two algorithms can be analyzed similarly as that in [31]. Assume that the HSS tree \mathbf{T} is a full quad tree with maximum level L , and that the number of rows and columns of the HSS matrix \mathbf{H} are $|\mathcal{J}| = M = 4^L M_L$ and $|\mathcal{K}| = N = 4^L N_L$ according to the HSS tree, where M_L and N_L are the row and column sizes of the leaf nodes, respectively. We also assume that $N = n^2$ where $n = 2^L n_L$. For each level ℓ , there are 4^ℓ nodes, and assume that the rank of the HSS blocks of a node on level ℓ is bounded by $r_\ell = \mathcal{O}(n_\ell \log n_\ell)$ where $n_\ell = 2^{L-\ell} n_L$. Then the complexity of the URV factorization and solution algorithms are $\mathcal{O}(M + N^{3/2} \log^3 N)$ and $\mathcal{O}(M + N \log^3 N)$, respectively. We point out here that the URV factorization can also be used compute $\mathbf{H}^{*,\dagger} \mathbf{f}$ for a given vector \mathbf{f} . The solution process is similar and we omit the details here for brevity.

2.4 A Review of the Direct Inverse Method in 1D Case

Consider the 1D type-II NUFFT matrix given by $\mathbf{A}(j, k) = e^{-2\pi i k x_j}$ where $x_j \in [0, 1)$ for $j = 0, 1, \dots, M-1$ and $k = 0, 1, \dots, N-1$. If we define $\gamma_j = e^{-2\pi i x_j} \in \mathbb{S}$, then $\mathbf{A}(j, k) = \gamma_j^k$ and it can be directly verified that \mathbf{A} satisfies the following Sylvester equation

$$\mathbf{\Gamma} \mathbf{A} - \mathbf{A} \mathbf{C} = \mathbf{u} \mathbf{e}_{N-1}^\top, \quad (2.8)$$

where $\mathbf{\Gamma} = \text{diag}(\gamma_0, \gamma_1, \dots, \gamma_{M-1})$ is a diagonal matrix, and

$$\mathbf{C} = \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \\ 1 & & & & \end{bmatrix}$$

is a circulant shift matrix and $u_j = \gamma_j^N - 1$ for $j = 0, 1, \dots, M-1$. The matrix \mathbf{D} can be diagonalized by the DFT matrix, i.e., $\mathbf{D} = \mathbf{F}^{-1} \mathbf{D} \mathbf{F}$ where \mathbf{F} is the DFT matrix given by $\mathbf{F}(k, \ell) = e^{-2\pi i k \ell / N}$ and $\mathbf{D} = \text{diag}(\xi_0, \xi_1, \dots, \xi_{N-1})$ is a diagonal matrix. Here $\xi_k = \zeta_N^k \in \mathbb{S}$ and $\zeta_N = e^{-2\pi i / N}$ is the N th root of unity. Substituting the diagonalization of \mathbf{C} into (2.8) and multiplying both sides by \mathbf{F}^{-1} on the right, it follows that

$$\mathbf{\Gamma} \mathbf{G} - \mathbf{G} \mathbf{D} = \mathbf{u} \mathbf{v}^\top, \quad (2.9)$$

where $\mathbf{G} = \mathbf{A} \mathbf{F}^{-1}$ and $\mathbf{v} = \mathbf{F}^{-1} \mathbf{e}_{N-1}$. The entries of \mathbf{v} are given by $v_k = \xi_k / N$ for $k = 0, 1, \dots, N-1$. Therefore, \mathbf{G} can be expressed as a Cauchy-like matrix with the entries given by

$$\mathbf{G}(j, k) = \frac{u_j v_k}{\gamma_j - \xi_k} = \frac{1}{N} \frac{(\gamma_j / \xi_k)^N - 1}{(\gamma_j / \xi_k) - 1} =: \Psi(\gamma_j, \xi_k). \quad (2.10)$$

In [30], the authors utilize the *displacement structure* [4] of \mathbf{G} in (2.9) to derive the *low-rank property* of \mathbf{G} . They showed that \mathbf{G} can be approximated by an HSS matrix $\widehat{\mathbf{G}}$ with numerical HSS rank $r =$

$\mathcal{O}(\log(1/\varepsilon)\log N)$, where ε is the accuracy parameter for the approximation. Specifically, we have the following Theorem.

Theorem 2.4 (Numerical rank of HSS blocks of \mathbf{G} , Theorem 3.2 in [30]). *Let \mathbf{G} be the matrix defined by (2.10) and $0 < \varepsilon < 1$ and $\mathcal{J} = \{0, 1, \dots, M-1\}$ and $\mathcal{K} = \{0, 1, \dots, N-1\}$ be the row and column index sets of \mathbf{G} respectively. Suppose that $[a, b] \subset [0, 1)$ is an interval and $\mathcal{R} = \{j : x_j \in [a, b]\} \subset \mathcal{J}$ and $\mathcal{C} = \{k : k/N \in [a, b]\} \subset \mathcal{K}$ are the row and column index sets corresponding to this interval. Then for the 2-norm or the Frobenius norm, the HSS row or column block of \mathbf{G} with respect to \mathcal{R} or \mathcal{C} has a numerical rank at most $r = \mathcal{O}(\log(1/\varepsilon)\log n_c)$ where $n_c = |\mathcal{C}|$.*

There are two main differences between Theorem 2.4 and the original result in [30]. First, the original result in [30] is for the 2-norm while we add the Frobenius norm in Theorem 2.4. This is because the singular values of the corresponding submatrix decay exponentially. The second is that the rank bound here is more refined by carefully tracking the inequality in the proof of in [30, Theorem 3.2]. From the Sylvester equation (2.9), the HSS approximation $\tilde{\mathbf{G}}$ can be constructed by combining the *factored alternating direction implicit* (FADI) method [5] and interpolative decomposition [10], at a complexity of $\mathcal{O}((M+N)\log^2 N)$.

Once the HSS approximation $\tilde{\mathbf{G}}$ is constructed, the URV factorization discussion in Section 2.3 is performed, these are the offline steps of the direct inverse method. In the online step, given a vector $\mathbf{f} \in \mathbb{C}^M$, the solution of the least-squares problem associated with \mathbf{A} can be obtained by $\mathbf{c} = \mathbf{F}^{-1}\tilde{\mathbf{G}}^\dagger\mathbf{f}$, which involves one URV solution of $\tilde{\mathbf{G}}$ and one iFFT. Therefore, the online complexity is $\mathcal{O}((M+N)\log N)$.

3 The Low-Rank Property of the NUDFT Matrix

This section discusses the low-rank property of the 2D type-II NUDFT matrix. Section 3.1 gives a kernel matrix expression of the transformed 2D type-II NUDFT matrix and derives the low-rank property of the matrix. In Section 3.2, we define the face-splitting product of two HSS matrices and show that the face-splitting product of two HSS matrices is still an HSS matrix.

3.1 Kernel Matrix Perspective

Consider the 2D type-II NUDFT matrix $\mathbf{A} \in \mathbb{C}^{M \times N}$ given by (1.2). Let $\mathbf{A}^{[x]} \in \mathbb{C}^{M \times n^{[x]}}$ and $\mathbf{A}^{[y]} \in \mathbb{C}^{M \times n^{[y]}}$ be the type-II NUDFT in x and y direction respectively, i.e., $\mathbf{A}^{[x]}(j, k^{[x]}) = e^{-2\pi i k^{[x]} x_j}$ and $\mathbf{A}^{[y]}(j, k^{[y]}) = e^{-2\pi i k^{[y]} y_j}$. Then \mathbf{A} can be written as the face-splitting product $\mathbf{A} = \mathbf{A}^{[x]} \bullet \mathbf{A}^{[y]}$.

Let $\mathbf{F}^{[x]} \in \mathbb{C}^{n^{[x]} \times n^{[x]}}$ and $\mathbf{F}^{[y]} \in \mathbb{C}^{n^{[y]} \times n^{[y]}}$ be the DFT matrices in x and y direction respectively and $\mathbf{F} = \mathbf{F}^{[x]} \otimes \mathbf{F}^{[y]} \in \mathbb{C}^{N \times N}$ be the DFT matrix in 2D. Following the discussions in Section 2.4, if we define $\mathbf{G}^{[x]} = \mathbf{A}^{[x]} \mathbf{F}^{[x], -1}$, $\mathbf{G}^{[y]} = \mathbf{A}^{[y]} \mathbf{F}^{[y], -1}$ and $\mathbf{G} = \mathbf{A} \mathbf{F}^{-1}$, then it can also be verified that $\mathbf{G} = \mathbf{G}^{[x]} \bullet \mathbf{G}^{[y]}$. Using (2.10), the entries of \mathbf{G} can be written as

$$\begin{aligned} \mathbf{G}(j, (k^{[x]}, k^{[y]})) &= \frac{1}{n^{[x]}} \frac{(\gamma_j^{[x]} / \xi_{k^{[x]}}^{[x]})^{n^{[x]}} - 1}{(\gamma_j^{[x]} / \xi_{k^{[x]}}^{[x]}) - 1} \frac{1}{n^{[y]}} \frac{(\gamma_j^{[y]} / \xi_{k^{[y]}}^{[y]})^{n^{[y]}} - 1}{(\gamma_j^{[y]} / \xi_{k^{[y]}}^{[y]}) - 1} \\ &= \Psi^{[x]}(\gamma_j^{[x]}, \xi_{k^{[x]}}^{[x]}) \Psi^{[y]}(\gamma_j^{[y]}, \xi_{k^{[y]}}^{[y]}) \\ &=: \Psi(\boldsymbol{\gamma}_j, \boldsymbol{\xi}_{(k^{[x]}, k^{[y]})}), \end{aligned} \quad (3.1)$$

where $\gamma_j^{[x]} = e^{-2\pi i x_j}$, $\gamma_j^{[y]} = e^{-2\pi i y_j}$, $\xi_{k^{[x]}}^{[x]} = e^{-2\pi i k^{[x]} / n^{[x]}}$ and $\xi_{k^{[y]}}^{[y]} = e^{-2\pi i k^{[y]} / n^{[y]}}$ belong to the unit circle \mathbb{S} , $\Psi^{[x]}$ and $\Psi^{[y]}$ are corresponding kernel functions in x and y direction respectively, and $\boldsymbol{\gamma}_j = (\gamma_j^{[x]}, \gamma_j^{[y]})$ and $\boldsymbol{\xi}_{(k^{[x]}, k^{[y]})} = (\xi_{k^{[x]}}^{[x]}, \xi_{k^{[y]}}^{[y]})$ are points on \mathbb{S}^2 . Consequently, \mathbf{G} can be viewed as a kernel matrix with the kernel function

$$\Psi(\mathbf{z}, \mathbf{w}) = \Psi^{[x]}(z_1, w_1) \Psi^{[y]}(z_2, w_2), \quad \mathbf{z} = (z_1, z_2) \in \mathbb{S}^2, \quad \mathbf{w} = (w_1, w_2) \in \mathbb{S}^2. \quad (3.2)$$




Figure 3.1: The partition of $[0, 1]^2$ according to a rectangular domain. The red box represents the rectangular domain $[a, b] \times [c, d]$. The blue boxes represent the points whose x -coordinate not in $[a, b]$ but y -coordinate in $[c, d]$. The green boxes represent the points whose x -coordinate in $[a, b]$ but y -coordinate not in $[c, d]$. The yellow boxes represent the points whose x -coordinate not in $[a, b]$ and y -coordinate not in $[c, d]$.

The equation (3.1) gives a kernel matrix representation of \mathbf{G} , implying that it can be approximated by an HSS matrix. In the remaining part of this section, we give some intuition and discussion on the HSS rank. We first give a preliminary lemma, which shows that the face-splitting product of two low-rank matrices is still low-rank.

Lemma 3.1 (Rank of the matrix face-splitting product). *Let $\mathbf{A}_1 \in \mathbb{C}^{m \times n_1}$ and $\mathbf{A}_2 \in \mathbb{C}^{m \times n_2}$ be two matrices and $\mathbf{A} = \mathbf{A}_1 \bullet \mathbf{A}_2$, then $\text{rank}(\mathbf{A}) \leq \text{rank}(\mathbf{A}_1) \text{rank}(\mathbf{A}_2)$.*

Proof. Suppose $\text{rank}(\mathbf{A}_1) = r_1$ and $\text{rank}(\mathbf{A}_2) = r_2$. We factorize them as $\mathbf{A}_t = \mathbf{U}_t \mathbf{S}_t \mathbf{V}_t^*$ for $t = 1, 2$, where $\mathbf{U}_t \in \mathbb{C}^{m \times r_t}$, $\mathbf{S}_t \in \mathbb{C}^{r_t \times r_t}$ and $\mathbf{V}_t \in \mathbb{C}^{n_t \times r_t}$. Then for $1 \leq j \leq m$, $1 \leq k_1 \leq n_1$ and $1 \leq k_2 \leq n_2$, the $(j, (k_1, k_2))$ entry of \mathbf{A} is given by

$$\begin{aligned} \mathbf{A}(j, (k_1, k_2)) &= \mathbf{A}_1(j, k_1) \mathbf{A}_2(j, k_2) \\ &= \left(\sum_{\alpha_1, \alpha_2=1}^{r_1} \mathbf{U}_1(j, \alpha_1) \mathbf{S}_1(\alpha_1, \alpha_2) \bar{\mathbf{V}}_1(k_1, \alpha_2) \right) \left(\sum_{\beta_1, \beta_2=1}^{r_2} \mathbf{U}_2(j, \beta_1) \mathbf{S}_2(\beta_1, \beta_2) \bar{\mathbf{V}}_2(k_2, \beta_2) \right) \\ &= \sum_{\alpha_1=1}^{r_1} \sum_{\beta_1=1}^{r_2} \left(\mathbf{U}_1(j, \alpha_1) \mathbf{U}_2(j, \beta_1) \right) \sum_{\alpha_2=1}^{r_1} \sum_{\beta_2=1}^{r_2} \left(\mathbf{S}_1(\alpha_1, \alpha_2) \mathbf{S}_2(\beta_1, \beta_2) \right) \left(\bar{\mathbf{V}}_1(k_1, \alpha_2) \bar{\mathbf{V}}_2(k_2, \beta_2) \right). \end{aligned}$$

Therefore, $\mathbf{A} = (\mathbf{U}_1 \bullet \mathbf{U}_2) (\mathbf{S}_1 \otimes \mathbf{S}_2) (\mathbf{V}_1 \otimes \mathbf{V}_2)^*$ and $\text{rank}(\mathbf{A})$ is at most $r_1 r_2$. \square

Since each $z \in \mathbb{S}$ can be represented by $z = e^{-2\pi i t}$ where $t \in [0, 1)$, every point in \mathbb{S}^2 can be identified with a point in a periodic domain $[0, 1)^2$. Accordingly, the kernel function in (3.2) can be viewed as a function defined on $[0, 1)^2$. Consider a rectangular domain $[a, b] \times [c, d]$ in $[0, 1)^2$, shown as the red rectangle in Figure 3.1. We partition $[0, 1)^2$ into nine rectangles by cutting along the boundaries of $[a, b] \times [c, d]$. These rectangles can be further classified into four classes. The first class is the rectangle $[a, b] \times [c, d]$ itself. The second class consists of points whose x -coordinate lies outside $[a, b]$ while whose y -coordinate lies in $[c, d]$. The third class consists of points whose x -coordinate lies in $[a, b]$ while whose y -coordinate lies outside $[c, d]$. These two classes are colored as shown in blue and green, respectively, in Figure 3.1. The fourth class consists of points whose x -coordinate and y -coordinate both lie outside $[a, b]$ and $[c, d]$, respectively, and is colored as yellow in Figure 3.1.

Let \mathcal{R} and $\mathcal{C} = \mathcal{C}^{[x]} \times \mathcal{C}^{[y]}$ be row and column index sets of \mathbf{G} respectively, such that the corresponding row and column points belong to the first class. Define $\mathcal{R}^c = \mathcal{J} \setminus \mathcal{R}$ and $\mathcal{C}^c = \mathcal{K} \setminus \mathcal{C}$. Suppose $|\mathcal{C}^{[x]}| = n_c^{[x]}$ and $|\mathcal{C}^{[y]}| = n_c^{[y]}$ and $n_c = \max\{n_c^{[x]}, n_c^{[y]}\}$.

Consider the HSS row block $\mathbf{G}(\mathcal{R}, \mathcal{C}^c)$. We can partition it as

$$\mathbf{G}(\mathcal{R}, \mathcal{C}^c) = [\mathbf{G}(\mathcal{R}, \mathcal{C}_b) \quad \mathbf{G}(\mathcal{R}, \mathcal{C}_g) \quad \mathbf{G}(\mathcal{R}, \mathcal{C}_y)], \quad (3.3)$$

where \mathcal{C}_b , \mathcal{C}_g and \mathcal{C}_y denote the column index sets corresponding to the second, third and fourth classes respectively. Note that these submatrices can be further factorized as

$$\begin{aligned} \mathbf{G}(\mathcal{R}, \mathcal{C}_b) &= \mathbf{G}^{[x]}(\mathcal{R}, \mathcal{C}^{[x],c}) \bullet \mathbf{G}^{[y]}(\mathcal{R}, \mathcal{C}^{[y]}), \\ \mathbf{G}(\mathcal{R}, \mathcal{C}_g) &= \mathbf{G}^{[x]}(\mathcal{R}, \mathcal{C}^{[x]}) \bullet \mathbf{G}^{[y]}(\mathcal{R}, \mathcal{C}^{[y],c}), \\ \mathbf{G}(\mathcal{R}, \mathcal{C}_y) &= \mathbf{G}^{[x]}(\mathcal{R}, \mathcal{C}^{[x],c}) \bullet \mathbf{G}^{[y]}(\mathcal{R}, \mathcal{C}^{[y],c}). \end{aligned} \quad (3.4)$$

From Theorem 2.4, for a given accuracy parameter $0 < \varepsilon < 1$, the numerical ranks of $\mathbf{G}^{[x]}(\mathcal{R}, \mathcal{C}^{[x],c})$ and $\mathbf{G}^{[y]}(\mathcal{R}, \mathcal{C}^{[y],c})$ are bounded by $r = \mathcal{O}(\log(1/\varepsilon) \log n_c)$. Suppose for simplicity that the ranks of these two matrices are exactly at most r . Then from Lemma 3.1, the ranks of $\mathbf{G}(\mathcal{R}, \mathcal{C}_b)$, $\mathbf{G}(\mathcal{R}, \mathcal{C}_g)$ and $\mathbf{G}(\mathcal{R}, \mathcal{C}_y)$ are at most rn_c , rn_c and r^2 respectively. Consequently, the rank of $\mathbf{G}(\mathcal{R}, \mathcal{C}^c)$ is at most $2rn_c + r^2 = \mathcal{O}(\log(1/\varepsilon)n_c \log n_c)$. The same procedure also applies to the HSS column block $\mathbf{G}(\mathcal{R}^c, \mathcal{C})$. Note that in the HSS tree, the largest size of n_c is $\mathcal{O}(\sqrt{N})$. Therefore, intuitively one may argue that the HSS rank of the HSS approximation of \mathbf{G} is at most $\mathcal{O}(\log(1/\varepsilon)\sqrt{N} \log N)$.

3.2 Face-Splitting Product of HSS Trees and HSS Matrices

Let $\mathbb{T}^{[x]}$ be an HSS tree with row and column index sets \mathcal{J} and $\mathcal{K}^{[x]}$ and $\mathbb{T}^{[y]}$ be an HSS tree with row and column index sets \mathcal{J} and $\mathcal{K}^{[y]}$. Suppose $\mathcal{K} = \mathcal{K}^{[x]} \times \mathcal{K}^{[y]}$. We construct a new HSS tree \mathbb{T} with row and column index sets \mathcal{J} and \mathcal{K} as follows.

Definition 3.2 (Face-splitting product of HSS trees). *Let $\mathbb{T}^{[x]}$ and $\mathbb{T}^{[y]}$ be two HSS trees of row and column index sets \mathcal{J} and $\mathcal{K}^{[x]}$ and \mathcal{J} and $\mathcal{K}^{[y]}$ respectively, then the face-splitting product of $\mathbb{T}^{[x]}$ and $\mathbb{T}^{[y]}$, denoted by $\mathbb{T} = \mathbb{T}^{[x]} \bullet \mathbb{T}^{[y]}$, is defined as follows.*

- (1) The root of \mathbb{T} is $\rho = (\rho^{[x]}, \rho^{[y]})$, where $\rho^{[x]}$ and $\rho^{[y]}$ are the roots of $\mathbb{T}^{[x]}$ and $\mathbb{T}^{[y]}$ respectively.
- (2) If $\tau^{[x]}$ is a leaf and $\tau^{[y]}$ is a nonleaf, then $\tau = (\tau^{[x]}, \tau^{[y]})$ is a nonleaf node of \mathbb{T} and its children are given by $\text{ch}(\tau) = \{(\tau^{[x]}, \alpha^{[y]}) : \alpha^{[y]} \in \text{ch}(\tau^{[y]})\}$.
- (3) If $\tau^{[x]}$ is a nonleaf and $\tau^{[y]}$ is a leaf, then $\tau = (\tau^{[x]}, \tau^{[y]})$ is a nonleaf node of \mathbb{T} and its children are given by $\text{ch}(\tau) = \{(\alpha^{[x]}, \tau^{[y]}) : \alpha^{[x]} \in \text{ch}(\tau^{[x]})\}$.
- (4) If $\tau^{[x]}$ and $\tau^{[y]}$ are nonleaf, then $\tau = (\tau^{[x]}, \tau^{[y]})$ is a nonleaf node of \mathbb{T} and its children are given by $\text{ch}(\tau) = \{(\alpha^{[x]}, \alpha^{[y]}) : \alpha^{[x]} \in \text{ch}(\tau^{[x]}), \alpha^{[y]} \in \text{ch}(\tau^{[y]})\}$.
- (5) If $\tau^{[x]}$ and $\tau^{[y]}$ are both leaves, then $\tau = (\tau^{[x]}, \tau^{[y]})$ is a leaf node of \mathbb{T} .
- (6) Each node $\tau = (\tau^{[x]}, \tau^{[y]})$ is associated with the row index set $\mathcal{J}_\tau = \mathcal{J}_{\tau^{[x]}} \cap \mathcal{J}_{\tau^{[y]}}$ and the column index set $\mathcal{K}_\tau = \mathcal{K}_{\tau^{[x]}} \times \mathcal{K}_{\tau^{[y]}}$.

Typically, from the discussions in Section 2.4, the HSS tree $\mathbb{T}^{[x]}$ and $\mathbb{T}^{[y]}$ are constructed by the bisection of the periodic interval $[0, 1)$. Therefore, the HSS tree \mathbb{T} can be viewed as a tree constructed by the quadrisection of $[0, 1)^2$. Note that each node of \mathbb{T} is a ‘‘tensor product of’’ a node of $\mathbb{T}^{[x]}$ and a node of $\mathbb{T}^{[y]}$, corresponding to a rectangular domain in $[0, 1)^2$. We first prove the face-splitting product of two HSS trees is still an HSS tree and then prove the face-splitting product of two HSS matrices is still an HSS matrix.

Theorem 3.3 (Face-splitting product of HSS trees is an HSS tree). *Let T be a face-splitting product of two HSS trees $\mathsf{T}^{[x]}$ and $\mathsf{T}^{[y]}$ of row and column index sets \mathcal{J} and $\mathcal{K}^{[x]}$ and \mathcal{J} and $\mathcal{K}^{[y]}$ respectively, then T is an HSS tree of row and column index sets \mathcal{J} and \mathcal{K} .*

Proof. It suffices to verify that the conditions in Definition 2.2 for T .

- (1) For the root node $\rho = (\rho^{[x]}, \rho^{[y]})$, we have $\mathcal{J}_\rho = \mathcal{J}_{\rho^{[x]}} \cap \mathcal{J}_{\rho^{[y]}} = \mathcal{J} \cap \mathcal{J} = \mathcal{J}$ and $\mathcal{K}_\rho = \mathcal{K}_{\rho^{[x]}} \times \mathcal{K}_{\rho^{[y]}} = \mathcal{K}^{[x]} \times \mathcal{K}^{[y]} = \mathcal{K}$.
- (2) If $\tau = (\tau^{[x]}, \tau^{[y]})$ is a nonleaf node, then there are three cases. We only give the proof for the first case and the proof for the other two cases can be done by similar arguments. If $\tau^{[x]}$ is a leaf and $\tau^{[y]}$ is a nonleaf, then the children of τ are given by $\text{ch}(\tau) = \{(\tau^{[x]}, \alpha^{[y]}) : \alpha^{[y]} \in \text{ch}(\tau^{[y]})\}$. Therefore,

$$\begin{aligned} \mathcal{J}_\tau &= \mathcal{J}_{\tau^{[x]}} \cap \mathcal{J}_{\tau^{[y]}} = \mathcal{J}_{\tau^{[x]}} \cap \left(\bigsqcup_{\alpha^{[y]} \in \text{ch}(\tau^{[y]})} \mathcal{J}_{\alpha^{[y]}} \right) = \bigsqcup_{\alpha^{[y]} \in \text{ch}(\tau^{[y]})} \left(\mathcal{J}_{\tau^{[x]}} \cap \mathcal{J}_{\alpha^{[y]}} \right) = \bigsqcup_{\alpha \in \text{ch}(\tau)} \mathcal{J}_\alpha, \\ \mathcal{K}_\tau &= \mathcal{K}_{\tau^{[x]}} \times \mathcal{K}_{\tau^{[y]}} = \mathcal{K}_{\tau^{[x]}} \times \left(\bigsqcup_{\alpha^{[y]} \in \text{ch}(\tau^{[y]})} \mathcal{K}_{\alpha^{[y]}} \right) = \bigsqcup_{\alpha^{[y]} \in \text{ch}(\tau^{[y]})} \left(\mathcal{K}_{\tau^{[x]}} \times \mathcal{K}_{\alpha^{[y]}} \right) = \bigsqcup_{\alpha \in \text{ch}(\tau)} \mathcal{K}_\alpha. \end{aligned}$$

This completes the proof. \square

Theorem 3.4 (Face-splitting product of HSS matrices is an HSS matrix). *Let T be a face-splitting product of two HSS trees $\mathsf{T}^{[x]}$ and $\mathsf{T}^{[y]}$ of row and column index sets \mathcal{J} and $\mathcal{K}^{[x]}$ and \mathcal{J} and $\mathcal{K}^{[y]}$ respectively. If $\mathbf{H}^{[x]}$ and $\mathbf{H}^{[y]}$ are two HSS matrices corresponding to $\mathsf{T}^{[x]}$ and $\mathsf{T}^{[y]}$, then $\mathbf{H} = \mathbf{H}^{[x]} \bullet \mathbf{H}^{[y]}$ is an HSS matrix for the HSS tree T .*

Furthermore, suppose that the rank of HSS blocks in $\mathsf{T}^{[x]}$ and $\mathsf{T}^{[y]}$ only depend on the level, that is, for any node $\tau^{[x]}$ in $\mathsf{T}^{[x]}$ with $\text{level}(\tau^{[x]}) = \ell$ and any node $\tau^{[y]}$ in $\mathsf{T}^{[y]}$ with $\text{level}(\tau^{[y]}) = \ell$, the rank of the HSS blocks corresponding to $\tau^{[x]}$ and $\tau^{[y]}$ is at most $r_\ell^{[x]}$ and $r_\ell^{[y]}$ respectively, then for any node $\tau = (\tau^{[x]}, \tau^{[y]})$ in T with $\text{level}(\tau) = \ell$, the rank of the HSS blocks corresponding to τ is at most $r_\ell^{[x]} d_\ell^{[y]} + r_\ell^{[y]} d_\ell^{[x]} + r_\ell^{[x]} r_\ell^{[y]}$ where $d_\ell^{[x]}$ and $d_\ell^{[y]}$ are the maximum size of all diagonal blocks of $\mathbf{H}^{[x]}$ and $\mathbf{H}^{[y]}$ at level ℓ respectively.

Proof. Without loss of generality, assume that $\mathsf{T}^{[x]}$ and $\mathsf{T}^{[y]}$ are full binary HSS trees of maximum level L . This means that the cases (2) and (3) in Definition 3.2 will not happen. The maximum level of T is L , and every nonleaf node of T has exactly four children. General cases can be proved by similar arguments.

Let $\{\mathbf{D}^{[x]}, \mathbf{U}^{[x]}, \mathbf{V}^{[x]}, \mathbf{R}^{[x]}, \mathbf{W}^{[x]}, \mathbf{B}^{[x]}\}$ and $\{\mathbf{D}^{[y]}, \mathbf{U}^{[y]}, \mathbf{V}^{[y]}, \mathbf{R}^{[y]}, \mathbf{W}^{[y]}, \mathbf{B}^{[y]}\}$ the generator of $\mathbf{H}^{[x]}$ and $\mathbf{H}^{[y]}$ respectively. For a node $\lambda = (\tau, \sigma) \in \mathsf{T}$, we define

$$\mathbf{D}_\lambda := \mathbf{H}(\mathcal{J}_\lambda, \mathcal{K}_\lambda) = \mathbf{H}(\mathcal{J}_\lambda, \mathcal{K}_\tau^{[x]} \times \mathcal{K}_\sigma^{[y]}) = \mathbf{H}^{[x]}(\mathcal{J}_\lambda, \mathcal{K}_\tau^{[x]}) \bullet \mathbf{H}^{[y]}(\mathcal{J}_\lambda, \mathcal{K}_\sigma^{[y]}).$$

By $\mathbf{D}_\tau^{[x]} = \mathbf{H}^{[x]}(\mathcal{J}_\tau, \mathcal{K}_\tau^{[x]})$ and $\mathcal{J}_\lambda = \mathcal{J}_\tau \cap \mathcal{J}_\sigma$, $\mathbf{H}^{[x]}(\mathcal{J}_\lambda, \mathcal{K}_\tau^{[x]})$ is a submatrix corresponding to some rows of $\mathbf{D}_\tau^{[x]}$. We denote $\mathbf{H}^{[x]}(\mathcal{J}_\lambda, \mathcal{K}_\tau^{[x]}) = \mathbf{D}_\lambda^{[x]}$. Similarly, we denote $\mathbf{H}^{[y]}(\mathcal{J}_\lambda, \mathcal{K}_\sigma^{[y]}) = \mathbf{D}_\lambda^{[y]}$. Therefore,

$$\mathbf{D}_\lambda = \mathbf{D}_\lambda^{[x]} \bullet \mathbf{D}_\lambda^{[y]}.$$

If $\lambda = (\tau, \sigma)$ is a leaf, then both τ and σ are leaf, meaning that both $\mathbf{D}_\tau^{[x]}$ and $\mathbf{D}_\sigma^{[y]}$ are dense matrices, thus \mathbf{D}_λ defined above is a dense matrix as well.

If $\lambda = (\tau, \sigma)$ is a nonleaf, then by our assumption both τ and σ are nonleaf and $\mathbf{D}_\tau^{[x]}$ and $\mathbf{D}_\sigma^{[y]}$ are 2×2 block matrices. Let $\text{ch}(\tau) = \{\alpha_1, \alpha_2\}$ and $\text{ch}(\sigma) = \{\beta_1, \beta_2\}$, then

$$\text{ch}(\lambda) = \{(\alpha_1, \beta_1), (\alpha_1, \beta_2), (\alpha_2, \beta_1), (\alpha_2, \beta_2)\}.$$

The columns of \mathbf{D}_λ can be partitioned into 4 column blocks corresponding to $\text{ch}(\lambda)$, with the (r, s) -th column block of is

$$\mathbf{D}_{\lambda;::,(r,s)} = \mathbf{D}_{\lambda;::,r}^{[x]} \bullet \mathbf{D}_{\lambda;::,s}^{[y]} = \mathbf{H}^{[x]}(\mathcal{J}_\lambda, \mathcal{K}_{\alpha_r}^{[x]}) \bullet \mathbf{H}^{[y]}(\mathcal{J}_\lambda, \mathcal{K}_{\beta_s}^{[y]}).$$

We partition the rows of D_λ into 4 blocks according to its children, where the (p, q) -th row block of D_λ is

$$D_{\lambda; (p, q), :} = D_{\lambda; p, :}^{[x]} \bullet D_{\lambda; q, :}^{[y]} = H^{[x]}(\mathcal{J}_{(\alpha_p, \beta_q)}, \mathcal{K}_\tau^{[x]}) \bullet H^{[y]}(\mathcal{J}_{(\alpha_p, \beta_q)}, \mathcal{K}_\sigma^{[y]}).$$

Consequently, the $((p, q), (r, s))$ -th block of D_λ is

$$D_{\lambda; (p, q), (r, s)} = D_{\lambda; p, r}^{[x]} \bullet D_{\lambda; q, s}^{[y]} = H^{[x]}(\mathcal{J}_{(\alpha_p, \beta_q)}, \mathcal{K}_{\alpha_r}^{[x]}) \bullet H^{[y]}(\mathcal{J}_{(\alpha_p, \beta_q)}, \mathcal{K}_{\beta_s}^{[y]}).$$

In the following, we use the one-to-one mapping of the indices

$$(1, 1) \leftrightarrow (\alpha_1, \beta_1) \leftrightarrow 1, \quad (1, 2) \leftrightarrow (\alpha_1, \beta_2) \leftrightarrow 2, \quad (2, 1) \leftrightarrow (\alpha_2, \beta_1) \leftrightarrow 3, \quad (2, 2) \leftrightarrow (\alpha_2, \beta_2) \leftrightarrow 4.$$

We fix the row child node to be $\xi_1 = (\alpha_1, \beta_1)$ and discuss the block structure of D_λ . Note that for the off-diagonal blocks, we have

$$D_{\tau; 1, 2}^{[x]} = U_{\alpha_1}^{[x]} B_{\alpha_1, \alpha_2}^{[x]} V_{\alpha_2}^{[x], *}, \quad D_{\sigma; 1, 2}^{[y]} = U_{\beta_1}^{[y]} B_{\beta_1, \beta_2}^{[y]} V_{\beta_2}^{[y], *}.$$

- Case: The column child node is $\xi_1 = (\alpha_1, \beta_1)$, corresponding to block $(1, 1) = ((1, 1), (1, 1))$.

In this case,

$$D_{\lambda; 1, 1} = H^{[x]}(\mathcal{J}_{\xi_1}, \mathcal{K}_{\alpha_1}^{[x]}) \bullet H^{[y]}(\mathcal{J}_{\xi_1}, \mathcal{K}_{\beta_1}^{[y]}) = D_{\xi_1}.$$

- Case: The column child node is $\xi_2 = (\alpha_1, \beta_2)$, corresponding to block $(1, 2) = ((1, 1), (1, 2))$.

Using the proof of Lemma 3.1, we have

$$\begin{aligned} D_{\lambda; 1, 2} &= D_{\lambda; 1, 1}^{[x]} \bullet D_{\lambda; 1, 2}^{[y]} = D_{\alpha_1}^{[x]} \bullet (U_{\xi_1}^{[y]} B_{\beta_1, \beta_2}^{[y]} V_{\beta_2}^{[y], *}) \\ &= (D_{\xi_1}^{[x]} \bullet U_{\xi_1}^{[y]}) (I \otimes B_{\beta_1, \beta_2}^{[y]}) (I \otimes V_{\beta_2}^{[y], *}), \end{aligned} \quad (3.5)$$

where $U_{\xi_1}^{[y]}$ is the submatrix of $U_{\beta_1}^{[y]}$ corresponding to the row index set \mathcal{J}_{ξ_1} .

- Case: The column child node is $\xi_3 = (\alpha_2, \beta_1)$, corresponding to block $(1, 3) = ((1, 1), (2, 1))$.

In this case, by a similar argument as in (3.5), we have

$$D_{\lambda; 1, 3} = (U_{\xi_1}^{[x]} B_{\alpha_1, \alpha_2}^{[x]} V_{\alpha_2}^{[x], *}) \bullet D_{\xi_1}^{[y]} = (U_{\xi_1}^{[x]} \bullet D_{\xi_1}^{[y]}) (B_{\alpha_1, \alpha_2}^{[x]} \otimes I) (V_{\alpha_2}^{[x]} \otimes I)^*. \quad (3.6)$$

- Case: The column child node is $\xi_4 = (\alpha_2, \beta_2)$, corresponding to block $(1, 4) = ((1, 1), (2, 2))$.

In this case, we have

$$\begin{aligned} D_{\lambda; 1, 4} &= (U_{\xi_1}^{[x]} B_{\alpha_1, \alpha_2}^{[x]} V_{\alpha_2}^{[x], *}) \bullet (U_{\xi_1}^{[y]} B_{\beta_1, \beta_2}^{[y]} V_{\beta_2}^{[y], *}) \\ &= (U_{\xi_1}^{[x]} \bullet U_{\xi_1}^{[y]}) (B_{\alpha_1, \alpha_2}^{[x]} \otimes B_{\beta_1, \beta_2}^{[y]}) (V_{\alpha_2}^{[x]} \otimes V_{\beta_2}^{[y], *}). \end{aligned} \quad (3.7)$$

Therefore, for a nonroot node $\xi = (\alpha, \beta)$, if we define

$$U_\xi = \begin{bmatrix} D_\xi^{[x]} \bullet U_\xi^{[y]} & U_\xi^{[x]} \bullet D_\xi^{[y]} & U_\xi^{[x]} \bullet U_\xi^{[y]} \end{bmatrix}, \quad V_\xi = \begin{bmatrix} I \otimes V_\beta^{[y]} & V_\alpha^{[x]} \otimes I & V_\alpha^{[x]} \otimes V_\beta^{[y]} \end{bmatrix}, \quad (3.8)$$

then the off-diagonal blocks (3.5), (3.6) and (3.7) can be written as

$$\begin{aligned} D_{\lambda; 1, 2} &= U_{\xi_1; \xi_2} B_{\xi_1, \xi_2} V_{\xi_2}^*, & B_{\xi_1, \xi_2} &= \text{diag}(I \otimes B_{\beta_1, \beta_2}^{[y]}, \mathbf{0}, \mathbf{0}) \\ D_{\lambda; 1, 3} &= U_{\xi_1; \xi_3} B_{\xi_1, \xi_3} V_{\xi_3}^*, & B_{\xi_1, \xi_3} &= \text{diag}(\mathbf{0}, I \otimes B_{\alpha_1, \alpha_2}^{[x]}, \mathbf{0}) \\ D_{\lambda; 1, 4} &= U_{\xi_1; \xi_4} B_{\xi_1, \xi_4} V_{\xi_4}^*, & B_{\xi_1, \xi_4} &= \text{diag}(\mathbf{0}, \mathbf{0}, B_{\alpha_1, \alpha_2}^{[x]} \otimes B_{\beta_1, \beta_2}^{[y]}). \end{aligned}$$

The remaining is to verify the shared basis property. Suppose $\lambda = (\tau, \sigma)$ is a nonroot and nonleaf node. We only check it for the first row block corresponding to $\xi_1 = (\alpha_1, \beta_1)$. Consider \mathbf{V}_λ in (3.8) first, following a similar argument as above, we have

$$\begin{aligned}\mathbf{V}_{\lambda;1} &= \begin{bmatrix} \mathbf{I} \otimes \mathbf{V}_{\sigma;1}^{[y]} & \mathbf{V}_{\tau;1}^{[x]} \otimes \mathbf{I} & \mathbf{V}_{\tau;1}^{[x]} \otimes \mathbf{V}_{\sigma;1}^{[y]} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{I} \otimes (\mathbf{V}_{\beta_1}^{[y]} \mathbf{W}_{\beta_1}^{[y]}) & (\mathbf{V}_{\alpha_1}^{[x]} \mathbf{W}_{\alpha_1}^{[x]}) \otimes \mathbf{I} & (\mathbf{V}_{\alpha_1}^{[x]} \mathbf{W}_{\alpha_1}^{[x]}) \otimes (\mathbf{V}_{\beta_1}^{[y]} \mathbf{W}_{\beta_1}^{[y]}) \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{I} \otimes \mathbf{V}_{\beta_1}^{[y]} & \mathbf{V}_{\alpha_1}^{[x]} \otimes \mathbf{I} & \mathbf{V}_{\alpha_1}^{[x]} \otimes \mathbf{V}_{\beta_1}^{[y]} \end{bmatrix} \text{diag}(\mathbf{I} \otimes \mathbf{W}_{\beta_1}^{[y]}, \mathbf{W}_{\alpha_1}^{[x]} \otimes \mathbf{I}, \mathbf{W}_{\alpha_1}^{[x]} \otimes \mathbf{W}_{\beta_1}^{[y]}) \\ &= \mathbf{V}_{\xi_1}; \mathbf{W}_{\xi_1},\end{aligned}$$

where \mathbf{W}_{ξ_1} is defined as the block diagonal matrix in the above equation. Similarly, for \mathbf{U}_λ in (3.8), we have

$$\mathbf{U}_{\lambda;1} = \begin{bmatrix} \mathbf{D}_{\lambda;1,:}^{[x]} \bullet \mathbf{U}_{\lambda;1}^{[y]} & \mathbf{U}_{\lambda;1}^{[x]} \bullet \mathbf{D}_{\lambda;1,:}^{[y]} & \mathbf{U}_{\lambda;1}^{[x]} \bullet \mathbf{U}_{\lambda;1}^{[y]} \end{bmatrix}.$$

The first column block can be further factorized as

$$\begin{aligned}\mathbf{U}_{\lambda;1}^{[1]} &= \begin{bmatrix} \mathbf{D}_{\xi_1}^{[x]} & \mathbf{U}_{\xi_1}^{[x]} \mathbf{B}_{\alpha_1, \alpha_2}^{[x]} \mathbf{V}_{\alpha_2}^{[x],*} \end{bmatrix} \bullet (\mathbf{U}_{\xi_1}^{[y]} \mathbf{R}_{\beta_1}^{[y]}) \\ &= \begin{bmatrix} \mathbf{D}_{\xi_1}^{[x]} \bullet (\mathbf{U}_{\xi_1}^{[y]} \mathbf{R}_{\beta_1}^{[y]}) & (\mathbf{U}_{\xi_1}^{[x]} \mathbf{B}_{\alpha_1, \alpha_2}^{[x]} \mathbf{V}_{\alpha_2}^{[x],*}) \bullet (\mathbf{U}_{\xi_1}^{[y]} \mathbf{R}_{\beta_1}^{[y]}) \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{D}_{\xi_1}^{[x]} \bullet \mathbf{U}_{\xi_1}^{[y]} & \mathbf{U}_{\xi_1}^{[x]} \bullet \mathbf{D}_{\xi_1}^{[y]} & \mathbf{U}_{\xi_1}^{[x]} \bullet \mathbf{U}_{\xi_1}^{[y]} \end{bmatrix} \text{diag}\left(\mathbf{I} \otimes \mathbf{R}_{\beta_1}^{[y]}, \mathbf{0}, (\mathbf{B}_{\alpha_1, \alpha_2}^{[x]} \mathbf{V}_{\alpha_2}^{[x],*}) \otimes \mathbf{R}_{\beta_1}^{[y]}\right) \\ &= \mathbf{U}_{\xi_1} \mathbf{R}_{\xi_1}^{[1]},\end{aligned}$$

where $\mathbf{R}_{\xi_1}^{[1]}$ is defined as the block diagonal matrix in the above equation. Similarly, the second and third column blocks of $\mathbf{U}_{\lambda;1}$ can be factorized as $\mathbf{U}_{\lambda;1}^{[2]} = \mathbf{U}_{\xi_1} \mathbf{R}_{\xi_1}^{[2]}$ and $\mathbf{U}_{\lambda;1}^{[3]} = \mathbf{U}_{\xi_1} \mathbf{R}_{\xi_1}^{[3]}$. Therefore,

$$\mathbf{U}_{\lambda;1}^{[1]} = \begin{bmatrix} \mathbf{U}_{\lambda;1}^{[1]} & \mathbf{U}_{\lambda;1}^{[2]} & \mathbf{U}_{\lambda;1}^{[3]} \end{bmatrix} = \mathbf{U}_{\xi_1} \begin{bmatrix} \mathbf{R}_{\xi_1}^{[1]} & \mathbf{R}_{\xi_1}^{[2]} & \mathbf{R}_{\xi_1}^{[3]} \end{bmatrix} = \mathbf{U}_{\xi_1} \mathbf{R}_{\xi_1}.$$

Combining all the above results, we conclude that \mathbf{H} is an HSS matrix about T. The results on the rank of HSS blocks can be easily verified by the construction process. \square

For our NUDFT matrix, assume that $n^{[x]} = n^{[y]} = n$ and ε is a given accuracy parameter. Then for both $\mathbf{H}^{[x]}$ and $\mathbf{H}^{[y]}$, the rank of HSS blocks at level ℓ is $\mathcal{O}(\log(1/\varepsilon) \log(n/2^\ell))$ and the maximum size of the diagonal blocks at level ℓ is $\mathcal{O}(n/2^\ell)$. Therefore the rank of \mathbf{H} at level ℓ is $\mathcal{O}(\log(1/\varepsilon)(n/2^\ell) \log(n/2^\ell))$, implying that the HSS rank of \mathbf{H} is $\mathcal{O}(\log(1/\varepsilon)\sqrt{N} \log N)$.

Next we give a result on the approximation error of the face-splitting product of HSS approximations of $\mathbf{G}^{[x]}$ and $\mathbf{G}^{[y]}$ to \mathbf{G} . Before that, we point out that since every row of $\mathbf{A}^{[x]}$ has the same norm and $\mathbf{F}^{[x]}$ is unitary up to a scaling, every row of $\mathbf{G}^{[x]} = \mathbf{A}^{[x]} \mathbf{F}^{[x],-1}$ has the same norm. The similar result also holds for $\mathbf{G}^{[y]}$ and consequently for \mathbf{G} . These facts serve as a key tool in our proofs and we write it as a lemma here, whose proof is straightforward and omitted.

Lemma 3.5 (Equal row norm property). *Consider the 2D type-II NUDFT matrix $\mathbf{A} = \mathbf{A}^{[x]} \bullet \mathbf{A}^{[y]}$ defined by (1.2), where $\mathbf{A}^{[x]}$ and $\mathbf{A}^{[y]}$ be the type-II NUDFT matrices in x and y direction respectively. Let $\mathbf{F} = \mathbf{F}^{[x]} \otimes \mathbf{F}^{[y]}$ be the DFT matrix in 2D and $\mathbf{G} = \mathbf{A} \mathbf{F}^{-1} = \mathbf{G}^{[x]} \bullet \mathbf{G}^{[y]}$ where $\mathbf{G}^{[x]} = \mathbf{A}^{[x]} \mathbf{F}^{[x],-1}$ and $\mathbf{G}^{[y]} = \mathbf{A}^{[y]} \mathbf{F}^{[y],-1}$. Then every row of $\mathbf{G}^{[x]} = \mathbf{A}^{[x]} \mathbf{F}^{[x],-1}$ has the same norm, denoted by $Q^{[x]}$ and every row of $\mathbf{G}^{[y]} = \mathbf{A}^{[y]} \mathbf{F}^{[y],-1}$ has the same norm, denoted by $Q^{[y]}$. Consequently, every row of $\mathbf{G} = \mathbf{A} \mathbf{F}^{-1} = \mathbf{G}^{[x]} \bullet \mathbf{G}^{[y]}$ has the same norm $Q^{[x]} Q^{[y]}$. In particular, $\|\mathbf{G}^{[x]}\|_{\text{F}} = \sqrt{M} Q^{[x]}$, $\|\mathbf{G}^{[y]}\|_{\text{F}} = \sqrt{M} Q^{[y]}$ and $\|\mathbf{G}\|_{\text{F}} = \sqrt{M} Q^{[x]} Q^{[y]}$.*

Theorem 3.6. *Using the same notations as in Lemma 3.5, for a given $0 < \varepsilon < 1$, suppose $\tilde{\mathbf{G}}^{[x]}$ and $\tilde{\mathbf{G}}^{[y]}$ are two HSS approximations of $\mathbf{G}^{[x]}$ and $\mathbf{G}^{[y]}$ such that $\|\mathbf{G}^{[x]} - \tilde{\mathbf{G}}^{[x]}\|_{\text{F}} \leq \varepsilon\|\mathbf{G}^{[x]}\|_{\text{F}}$ and $\|\mathbf{G}^{[y]} - \tilde{\mathbf{G}}^{[y]}\|_{\text{F}} \leq \varepsilon\|\mathbf{G}^{[y]}\|_{\text{F}}$. Assume that exists a constant C such that $\max_j \{\|\tilde{\mathbf{G}}^{[x]}(j, :)\|_2\} \leq C \max_j \{\|\mathbf{G}^{[x]}(j, :)\|_2\}$. Then $\tilde{\mathbf{G}} = \tilde{\mathbf{G}}^{[x]} \bullet \tilde{\mathbf{G}}^{[y]}$ is an HSS approximation of $\mathbf{G} = \mathbf{G}^{[x]} \bullet \mathbf{G}^{[y]}$ and $\|\tilde{\mathbf{G}} - \mathbf{G}\|_{\text{F}} \leq (1 + C)\varepsilon\|\mathbf{G}\|_{\text{F}}$.*

Proof. From Lemma 2.1 and Lemma 3.5, we have

$$\begin{aligned} \|\mathbf{G} - \tilde{\mathbf{G}}\|_{\text{F}} &= \|(\mathbf{G}^{[x]} - \tilde{\mathbf{G}}^{[x]}) \bullet \mathbf{G}^{[y]} + \tilde{\mathbf{G}}^{[x]} \bullet (\mathbf{G}^{[y]} - \tilde{\mathbf{G}}^{[y]})\|_{\text{F}} \\ &\leq \|(\mathbf{G}^{[x]} - \tilde{\mathbf{G}}^{[x]}) \bullet \mathbf{G}^{[y]}\|_{\text{F}} + \|\tilde{\mathbf{G}}^{[x]} \bullet (\mathbf{G}^{[y]} - \tilde{\mathbf{G}}^{[y]})\|_{\text{F}} \\ &\leq \varepsilon\|\mathbf{G}^{[x]}\|_{\text{F}}Q^{[y]} + CQ^{[x]}\varepsilon\|\mathbf{G}^{[y]}\|_{\text{F}} \\ &= \varepsilon\sqrt{M}Q^{[x]}Q^{[y]} + C\varepsilon\sqrt{M}Q^{[x]}Q^{[y]} = (1 + C)\varepsilon\|\mathbf{G}\|_{\text{F}}. \end{aligned}$$

This completes the proof. \square

From $\|\mathbf{G}^{[x]} - \tilde{\mathbf{G}}^{[x]}\|_{\text{F}} \leq \varepsilon\|\mathbf{G}^{[x]}\|_{\text{F}}$, it is not hard to prove that C can be chosen as $1 + \sqrt{M}\varepsilon$. Therefore, one can expect the approximation error in Theorem 3.6 is not large when $\sqrt{M}\varepsilon$ is small. However, this estimate is quite pessimistic. If every row of $\mathbf{G}^{[x]} - \tilde{\mathbf{G}}^{[x]}$ has approximately the same norm (and this could be the case in practice), then one could expect that C is a small constant independent of M and N , meaning that the approximation error in Theorem 3.6 is small.

4 A Direct Solver for the Inverse NUDFT

In this section, we present a superfast solver for the inverse NUDFT problem (1.1) based on the HSS approximation of \mathbf{G} . Section 4.1 describes the construction of the HSS approximation of \mathbf{G} and Section 4.2 presents the algorithm for the approximate direct inversion.

4.1 The Construction of the HSS Approximation

Theorem 3.4 shows that $\mathbf{G} = \mathbf{G}^{[x]} \bullet \mathbf{G}^{[y]}$ could be approximated by an HSS matrix $\tilde{\mathbf{G}}$, and the proof of it also provides a way to construct $\tilde{\mathbf{G}}$ directly from the HSS approximations of $\mathbf{G}^{[x]}$ and $\mathbf{G}^{[y]}$. But, instead, our construction of $\tilde{\mathbf{G}}$ works on the kernel matrix expression (3.1) of \mathbf{G} directly, which is a common way in many literatures such as [9, 25, 26]. The main tool in the construction is the *interpolative decomposition* (ID) [10], which is defined as follows.

Lemma 4.1 (Interpolative decomposition [10]). *Let \mathbf{A} be an $m \times n$ matrix with rank $r \leq \min\{m, n\}$. Then there exist a subset $\hat{\mathcal{K}} \subset \{1, 2, \dots, n\}$ of size r and a matrix $\mathbf{Y} \in \mathbb{C}^{n \times r}$ such that $\mathbf{A} = \mathbf{A}(:, \hat{\mathcal{K}})\mathbf{Y}^*$ and $\mathbf{Y}(\hat{\mathcal{K}}, :) = \mathbf{I}$. This is called the column ID of \mathbf{A} . Similarly, there exist a subset $\hat{\mathcal{J}} \subset \{1, 2, \dots, m\}$ of size r and a matrix $\mathbf{L} \in \mathbb{C}^{m \times r}$ such that $\mathbf{A} = \mathbf{L}\mathbf{A}(\hat{\mathcal{J}}, :)$ and $\mathbf{L}(\hat{\mathcal{J}}, :) = \mathbf{I}$. This is called the row ID of \mathbf{A} . The matrices \mathbf{Y} and \mathbf{L} are called the interpolation matrices.*

Numerically, ID can be computed by rank-revealing QR factorization (RRQR) [17] or be approximated using the QR factorization with column pivoting (QRCP). Since \mathbf{G} has a kernel matrix expression, the *proxy surface* technique [25, 26] is used to accelerate the construction process. By (3.1), for any row and column index sets \mathcal{R} and \mathcal{C} , the submatrix can be expressed as $\mathbf{G}(\mathcal{R}, \mathcal{C}) = \Psi(\mathcal{Z}_{\mathcal{R}}, \mathcal{W}_{\mathcal{C}})$ where Ψ is the kernel function, $\mathcal{Z}_{\mathcal{R}} = \{\gamma_j\}_{j \in \mathcal{R}}$ and $\mathcal{W}_{\mathcal{C}} = \{\xi_k\}_{k \in \mathcal{C}}$.

We give a brief introduction of the construction process. Starting from the leaf nodes, we set $\mathbf{D}_{\tau} = \mathbf{G}(\mathcal{J}_{\tau}, \mathcal{K}_{\tau})$ directly. For each node τ , let \mathcal{P}_{τ} be a set of proxy points such that

$$\mathbf{G}(\mathcal{J}_{\tau}, \mathcal{K}_{\tau}^c) = \Psi(\mathcal{Z}_{\mathcal{J}_{\tau}}, \mathcal{W}_{\mathcal{K}_{\tau}^c}) \approx \Psi(\mathcal{Z}_{\mathcal{J}_{\tau}}, \mathcal{P}_{\tau})\mathbf{Q}_{\tau}^*, \quad \mathbf{G}(\mathcal{J}_{\tau}^c, \mathcal{K}_{\tau}) = \Psi(\mathcal{Z}_{\mathcal{J}_{\tau}^c}, \mathcal{W}_{\mathcal{K}_{\tau}}) \approx \mathbf{P}_{\tau}\Psi(\mathcal{P}_{\tau}, \mathcal{W}_{\mathcal{K}_{\tau}})$$

for some matrix \mathbf{P}_τ and \mathbf{Q}_τ . If we have the row ID $\Psi(\mathcal{L}_{\mathcal{J}_\tau}, \mathcal{P}_\tau) \approx \mathbf{L}_\tau \Psi(\mathcal{L}_{\mathcal{J}_\tau}, \mathcal{P}_\tau)$, then

$$\mathbf{G}(\mathcal{J}_\tau, \mathcal{K}_\tau^c) \approx \Psi(\mathcal{L}_{\mathcal{J}_\tau}, \mathcal{W}_{\mathcal{K}_\tau^c}) \approx \mathbf{L}_\tau \Psi(\mathcal{L}_{\mathcal{J}_\tau}, \mathcal{P}_\tau) \mathbf{Q}_\tau^*,$$

which gives the row ID of $\mathbf{G}(\mathcal{J}_\tau, \mathcal{K}_\tau^c)$. Similarly, the column ID of $\mathbf{G}(\mathcal{J}_\tau^c, \mathcal{K}_\tau)$ can be obtained from the column ID of $\Psi(\mathcal{P}_\tau, \mathcal{W}_{\mathcal{K}_\tau})$. Therefore, after computing the row and column IDs of $\Psi(\mathcal{L}_{\mathcal{K}_\tau}, \mathcal{P}_\tau)$ and $\Psi(\mathcal{P}_\tau, \mathcal{W}_{\mathcal{K}_\tau})$, the basis matrices can be obtained as $\mathbf{U}_\tau = \mathbf{L}_\tau$ and $\mathbf{V}_\tau = \mathbf{Y}_\tau$.

After computing the generators for all leaf nodes on level L , we move to level $L-1$. For a node τ on level $L-1$, we define the row index sets $\mathcal{J}_\tau^+ = \cup_{\alpha \in \text{ch}(\tau)} \hat{\mathcal{J}}_\alpha$, and set $\mathcal{J}^+ = \cup_{\text{level}(\tau)=L-1} \mathcal{J}_\tau^+$ and $\mathcal{J}_\tau^{c,+} = \mathcal{J}^+ \setminus \mathcal{J}_\tau^+$. The column index sets \mathcal{K}_τ^+ , \mathcal{K}^+ and $\mathcal{K}_\tau^{c,+}$ are defined similarly. For a node τ on level $L-1$, the interaction matrices of its children are given by $\mathbf{B}_{\alpha_p, \alpha_q} = \mathbf{G}(\hat{\mathcal{J}}_{\alpha_p}, \hat{\mathcal{K}}_{\alpha_q})$ for $p \neq q$. We then compute the interpolation matrices \mathbf{L}_τ and \mathbf{Y}_τ by the row and column IDs of $\Psi(\mathcal{L}_{\mathcal{J}_\tau^+}, \mathcal{P}_\tau)$ and $\Psi(\mathcal{P}_\tau, \mathcal{W}_{\mathcal{K}_\tau^+})$ respectively. Partitioning \mathbf{L}_τ and \mathbf{Y}_τ according to the children of τ , the transfer matrices are given as $\mathbf{R}_{\alpha_p} = \mathbf{L}_{\tau;p}$ and $\mathbf{W}_{\alpha_p} = \mathbf{Y}_{\tau;p}$. This process is repeated until we reach the root node ρ .

Typically, the size of the proxy point set is $\mathcal{O}(r_\tau)$ where r_τ is the rank of the HSS blocks corresponding to τ . In our case, \mathcal{P}_τ is chosen as the random points surrounding the bounding box of $\mathcal{L}_{\mathcal{J}_\tau}$ and $\mathcal{W}_{\mathcal{K}_\tau}$, which is shown in Figure 4.1. More specifically, suppose the box contains $n_c \times n_c$ column points. Let the 1D rank estimate in Theorem 2.4 be given as s_τ (hence $r_\tau \approx s_\tau n_c$). Then we sample $s_\tau n_c$ points uniformly on the horizontal and vertical sides of the box, and sample s_τ^2 points uniformly on the diagonal boundary of the box. Following the discussions in [25], the complexity of the construction of the HSS approximation of \mathbf{G} is $\mathcal{O}(M + N^{3/2} \log^3 N)$.

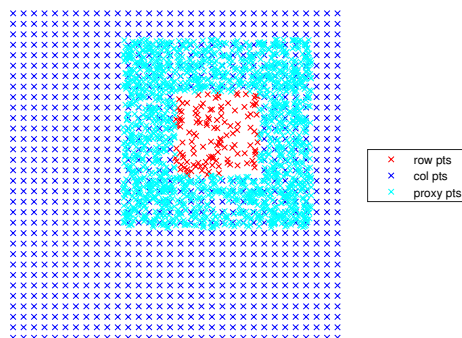


Figure 4.1: Proxy surfaces for the row ID in the construction of the HSS approximation of \mathbf{G} . The red points are the row points $\mathcal{L}_{\mathcal{J}_\tau}$, the blue points are the the column points $\mathcal{W}_{\mathcal{K}_\tau}$ and the cyan points are the proxy points \mathcal{P}_τ .

4.2 An Algorithm for the Direct Inversion of Type-II 2D NUDFT

We design an algorithm for solving the 2D type-II NUDFT based on the HSS matrix approximation of \mathbf{G} , based on the discussions from the previous sections. The algorithm consists of two stages. The *offline* stage first constructs the HSS approximation of \mathbf{G} using the construction method described in Section 4.1, then computes the URV factorization of the HSS matrix as described in Section 2.3. This is independent of the target values \mathbf{f} and can be reused. The *online* stage computes the solution of the inverse NUDFT problem by the URV solution of HSS matrix and two-dimensional iFFT. The pseudocode of the algorithm is summarized in Algorithm 1. Algebraically, we have a factorization

$$\mathbf{A} = \mathbf{G}\mathbf{F} \approx \mathbf{A}_{\text{fast}} := \tilde{\mathbf{G}}\mathbf{F},$$

for the type-II 2D NUFFT matrix \mathbf{A} , where $\tilde{\mathbf{G}}$ is the HSS approximation of \mathbf{G} and \mathbf{F} is the 2D DFT matrix. The pseudo-inverse of \mathbf{A} can be approximated by

$$\mathbf{A}^\dagger \approx \mathbf{A}_{\text{fast}}^\dagger = \mathbf{F}^{-1} \tilde{\mathbf{G}}^\dagger,$$

where $\tilde{\mathbf{G}}^\dagger$ is represented by its URV factorization and \mathbf{F}^{-1} can be computed by two-dimensional iFFT. We also remark that our fast direct solver for the 2D type-II NUFFT can be used as a preconditioner for iterative solvers such as lsqr, as long as we can efficiently apply $\mathbf{A}_{\text{fast}}^{\dagger,*} = \tilde{\mathbf{G}}^{\dagger,*} \mathbf{F}/N$, which involves the URV solution of HSS matrix mentioned in the end of Section 2.3 and two-dimensional FFT.

Algorithm 1 A direct solver for type-II 2D NUFFT problem

Input: Sample points $\{(x_j, y_j)\}_{j=0}^{M-1}$ and number of frequencies $n^{[x]}$ and $n^{[y]}$, target values $\{f_j\}_{j=0}^{M-1}$.

Output: Coefficients \mathbf{c} .

- 1: Construct the HSS approximation $\tilde{\mathbf{G}}$ of $\mathbf{G} = \mathbf{A}\mathbf{F}^{-1}$ using the kernel matrix expression (3.1).
 - 2: Compute the URV factorization of $\tilde{\mathbf{G}}$.
 - 3: Compute $\mathbf{c} = \mathbf{F}^{-1} \tilde{\mathbf{G}}^\dagger \mathbf{f}$ by the URV solution and two-dimensional iFFT.
-

We analyze the complexity of Algorithm 1 under a simplified setting. Suppose $\mathbb{T}^{[x]}$ and $\mathbb{T}^{[y]}$ are both full binary HSS trees with maximum level L constructed by bisecting the column index sets, and the maximum level of both trees is L , and all leaf nodes have exact column size n_L , where n_L is a small constant independent of $n^{[x]}$ and $n^{[y]}$. That is, $n^{[x]} = n^{[y]} = n = 2^L n_L$ and $N = 4^L N_L$ where $N_L = n_L^2$. Let $\mathbb{T} = \mathbb{T}^{[x]} \bullet \mathbb{T}^{[y]}$ and assume further that $M = 4^L M_L$ correspondingly where M_L is the number of sample points in the leaf nodes. For level ℓ , there are 4^ℓ nodes and the HSS blocks on level ℓ be denoted by r_ℓ . By the discussions in Section 3.1 and Section 3.2, we assume that $r_\ell = \mathcal{O}(n_\ell \log n_\ell)$. Following the results in Section 4.1 and Section 2.3, the complexity of the offline stage is $\mathcal{O}(M + N^{3/2} \log^3 N)$, while the complexity of the URV solution is $\mathcal{O}(M + N \log^3 N)$.

5 Numerical Results

In this section, we present numerical results of our direct solver for the 2D type-II NUFFT. Two types of nonuniform grids are considered: the random grid and the (modified) polar grid [15]. The random grid is generated by sampling M points distributed uniformly randomly in $[0, 1]^2$. The polar grid is defined as the zero point $(0, 0)$ together with points of the form

$$\begin{aligned} x_{(p,q)} &= \frac{1}{2} + r_p \cos(2\pi t_q), & y_{(p,q)} &= \frac{1}{2} + r_p \sin(2\pi t_q), \\ r_p &= \frac{\sqrt{2}}{2} \frac{p}{n_r}, & t_q &= \frac{q}{n_t}, & p &= 1, \dots, n_r - 1, & q &= 0, 1, \dots, n_t - 1, \end{aligned}$$

where we exclude the points outside $[0, 1]^2$. Figure 5.1 shows the two grid examples.

In our experiments, we first let n ranges in $\{32, 64, 128, 256, 512\}$ and set $n^{[x]} = n^{[y]} = n$ and $N = n^{[x]} n^{[y]} = n^2$. For random grid, the number of sample points is set as $M = \alpha N$, where α is selected to be 1.5. For polar grid, we set $n_r = n$ and $n_t \approx \beta n \log_2 n$ for $\beta = 0.6$, which gives $M \approx 0.8\beta N \log_4 N$, where the factor 0.8 comes from the sample area that is inside $[0, 1]^2$. The well-known Shepp-Logan phantom, which is widely used in the field of magnetic resonance imaging (MRI) [7], is used as the exact solution. That is, the coefficients \mathbf{c} are sampled on a $n \times n$ grid of the Shepp-Logan phantom, and the target values \mathbf{f} are generated by applying the type-II 2D NUFFT to \mathbf{c} . The tolerance parameter ε in the construction of the HSS approximation is set as 10^{-2} or 10^{-4} for all experiments. We use the MATLAB command “lsqr” [27]




Figure 5.1: Nonuniform grids. $N = 32^2$. Left: Random grid with $M = 2N$ points. Right: Polar grid with $M \approx 0.54N \log_4(N)$ points.

as the iterative solver, with the maximum number of iterations set as 500 and the tolerance parameter set as 10^{-12} . All algorithms are implemented in MATLAB R2023b and the FINUFFT library [3] are utilized for forward NUFFT, and are carried out on a server with two Intel Gold 6226R CPUs at 2.90 GHz and 1000.6 GB of RAM.

The following quantities are used to evaluate the performance of the algorithms. We use r_h to denote the HSS rank of the HSS approximation of \mathbf{G} , t_c to denote the walltime of constructing the HSS approximation, t_f to denote the walltime of computing the URV factorization of the HSS approximation, t_s to denote the walltime of computing the solution by the URV solution of HSS matrix and iFFT, r_s to denote the relative residual of the solution, i.e., $\|\mathbf{Ac} - \mathbf{f}\|_2 / \|\mathbf{f}\|_2$, and e_s to denote the relative error of the solution, i.e., $\|\mathbf{c} - \mathbf{c}_{\text{ex}}\|_2 / \|\mathbf{c}_{\text{ex}}\|_2$ where \mathbf{c}_{ex} is the exact solution. For iterative solvers, we use t_{pre} to denote the walltime of constructing the preconditioner, i.e., the offline computation, which equals to $t_c + t_f$, t_{iter} to denote the walltime of all iterations when solving the problem by MATLAB’s “lsqr” command with or without a preconditioner, and n_{iter} to denote the number of iterations.

5.1 Random Grid

Figure 5.2 shows the scaling of the HSS rank, construction time, factorization time, and solution time of Algorithm 1 on random grids with $M = 1.5N$. The corresponding numerical results are summarized in Table 5.1. It can be observed that the HSS rank grows as $\mathcal{O}(\sqrt{N} \log N)$, which is consistent with our analysis. Moreover, both the construction and factorization times follow the predicted $\mathcal{O}(M + N^{3/2} \log^3 N)$ complexity. The solving time is significantly smaller than the construction and factorization times, indicating that the dominant computational cost lies in the offline stage, while the online stage is negligible in comparison. For example, when $N = 512^2$, $M = 1.5N$ and $\varepsilon = 10^{-2}$, the construction takes approximately 1.4 hours, whereas the factorization and solution times are only about 28 minutes and 50 seconds, respectively. Furthermore, the tolerance parameter ε has a significant impact on the computational cost. For the same example, when ε is decreased from 10^{-2} to 10^{-4} , the construction, factorization, and solution times increase by factors of approximately 4.7, 3.2, and 1.3, respectively. In most cases, the relative residuals and relative errors of the direct solver are on the order of 10^{-2} for $\varepsilon = 10^{-2}$ and 10^{-4} for $\varepsilon = 10^{-4}$, suggesting that the algorithm remains stable as the problem size increases.

Table 5.2 demonstrates the behavior of the iterative solver using our direct solver as a preconditioner, where lsqr denotes the standard lsqr method and plsqr denotes the preconditioned lsqr method. First, the results show that, for random grids with $M = 1.5N$, standard lsqr fails to converge within 500 iterations for all tested problem sizes. In contrast, plsqr converges in 8 to 40 iterations when $\varepsilon = 10^{-2}$ and 4 to 14 iterations when $\varepsilon = 10^{-4}$, and the relative residual and relative error achieve at a level of 10^{-12} in both

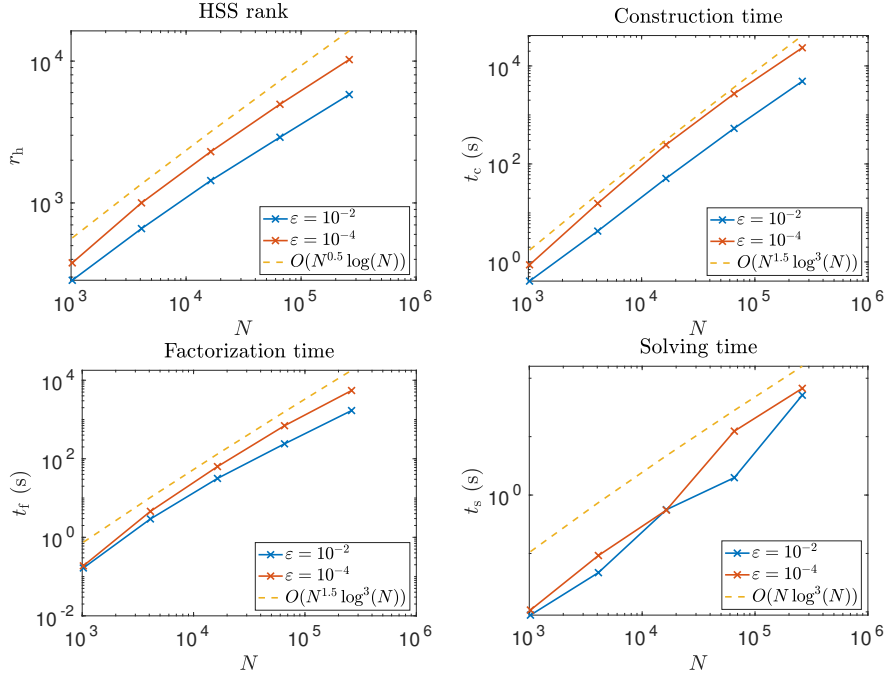


Figure 5.2: HSS rank, construction time, factorization time and solution time of the direct solver on a random grid. $M = 1.5N$.

N	M	ϵ	t_c	t_f	t_s	r_s	e_s
32^2	1536	10^{-2}	4.1e-01	1.7e-01	8.8e-03	2.1e-03	2.3e-03
		10^{-4}	8.7e-01	1.9e-01	1.1e-02	1.0e-05	1.2e-05
64^2	6144	10^{-2}	4.3e+00	2.9e+00	4.7e-02	3.9e-03	5.9e-03
		10^{-4}	1.6e+01	4.6e+00	9.2e-02	3.0e-05	4.3e-05
128^2	24576	10^{-2}	5.1e+01	3.2e+01	5.6e-01	5.0e-03	7.1e-03
		10^{-4}	2.5e+02	6.3e+01	5.5e-01	1.4e-04	2.1e-04
256^2	98304	10^{-2}	5.3e+02	2.4e+02	2.0e+00	7.6e-03	1.1e-02
		10^{-4}	2.7e+03	6.9e+02	1.2e+01	4.2e-04	6.0e-04
512^2	393216	10^{-2}	4.9e+03	1.7e+03	5.1e+01	1.1e-02	1.7e-02
		10^{-4}	2.3e+04	5.5e+03	6.7e+01	7.3e-04	1.1e-03

Table 5.1: Construction time, factorization time, solve time, relative residual and relative error of the direct solver on a random grid. $M = 1.5N$.

N	M	method	t_{pre}	t_{iter}	n_{iter}	r_s	e_s
32^2	1536	lsqr	-	1.2e+00	500	5.2e-10	5.7e-10
		plsqr ($\varepsilon = 10^{-2}$)	5.7e-01	1.4e-01	8	2.6e-12	3.0e-12
		plsqr ($\varepsilon = 10^{-4}$)	1.1e+00	1.1e-01	4	4.4e-15	5.2e-15
64^2	6144	lsqr	-	2.7e+00	500	2.8e-08	5.4e-08
		plsqr ($\varepsilon = 10^{-2}$)	7.2e+00	1.0e+00	12	5.6e-13	9.6e-13
		plsqr ($\varepsilon = 10^{-4}$)	2.0e+01	8.9e-01	4	1.3e-13	2.0e-13
128^2	24576	lsqr	-	4.8e+00	500	6.5e-05	9.4e-05
		plsqr ($\varepsilon = 10^{-2}$)	8.2e+01	8.5e+00	14	4.3e-12	7.0e-12
		plsqr ($\varepsilon = 10^{-4}$)	3.1e+02	7.6e+00	6	2.5e-14	3.5e-14
256^2	98304	lsqr	-	1.2e+01	500	1.2e-04	2.1e-04
		plsqr ($\varepsilon = 10^{-2}$)	7.7e+02	6.6e+01	23	2.0e-12	3.1e-12
		plsqr ($\varepsilon = 10^{-4}$)	3.4e+03	5.6e+01	8	8.8e-13	1.3e-12
512^2	393216	lsqr	-	3.0e+01	500	1.9e-04	3.2e-04
		plsqr ($\varepsilon = 10^{-2}$)	6.5e+03	5.1e+02	39	3.6e-12	5.5e-12
		plsqr ($\varepsilon = 10^{-4}$)	2.9e+04	4.6e+02	14	2.7e-12	4.0e-12

Table 5.2: Time cost, iteration number, relative residual and relative error of lsqr and plsqr on a random grid. $M = 1.5N$.

cases. It means that the preconditioner is effective and efficient.

5.2 Polar Grid

Following the setting introduced at the beginning of Section 5, for the polar grid we set $n_r = n$ and $n_t \approx \beta n \log_2 n$ with $\beta = 0.6$. This gives $M \approx 0.47N \log_4 N$. Figure 5.3 reports the scaling of the HSS rank, construction time, factorization time, and solution time of Algorithm 1 on the polar grid with $\beta = 0.6$. The corresponding numerical results are summarized in Table 5.3. Similar to the random-grid case, the HSS rank grows approximately as $\mathcal{O}(\sqrt{N} \log N)$, in agreement with the estimation in Section 4.2. The construction and factorization times also follow the predicted $\mathcal{O}(M + N^{3/2} \log^3 N)$ scaling. In particular, since $M \approx 0.47N \log_4 N$ for the polar grid, the dominant cost in both stages comes from the $\mathcal{O}(N^{3/2} \log^3 N)$ term rather than the linear $\mathcal{O}(M)$ term. The solution time is much smaller than the construction and factorization times. For instance, when $N = 512^2$ and $\varepsilon = 10^{-2}$, the construction takes about 1.7 hours, whereas the factorization and solution times are approximately 33 minutes and 34 seconds, respectively. When the tolerance is tightened to $\varepsilon = 10^{-4}$, the construction, factorization, and solution times increase to 7.2 hours, 1.9 hours, and 92 seconds, respectively. Moreover, the relative residuals and relative errors of the direct solver are generally on the order of 10^{-2} for $\varepsilon = 10^{-2}$ and 10^{-4} for $\varepsilon = 10^{-4}$, indicating that the proposed direct solver remains stable for the polar grid as the problem size increases.

Table 5.4 shows the performance of lsqr and plsqr on the polar grid. For $\beta = 0.6$, lsqr fails to converge within 500 iterations for all but the smallest problem size. In contrast, plsqr converges in 9 to 38 iterations when $\varepsilon = 10^{-2}$, and in 4 to 10 iterations when $\varepsilon = 10^{-4}$. In both cases, the relative residuals and relative errors reach the level of 10^{-11} to 10^{-12} , demonstrating the effectiveness of the proposed preconditioner. Overall, these results confirm that the proposed preconditioner significantly improves the convergence of lsqr on polar grids, especially for larger and more ill-conditioned problems.

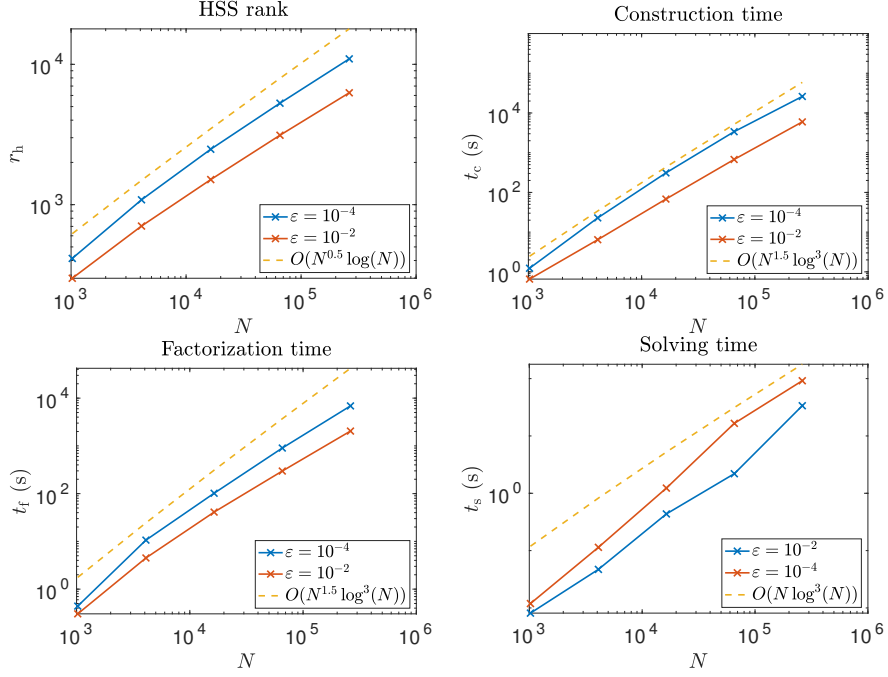


Figure 5.3: HSS rank, construction time, factorization time and solution time of the direct solver on a polar grid. $\beta = 0.6$. $M \approx 0.47N \log_4(N)$.

N	M	ϵ	t_c	t_f	t_s	r_s	e_s
32^2	2397	10^{-2}	6.6e-01	3.0e-01	8.1e-03	6.6e-04	7.0e-03
		10^{-4}	1.2e+00	4.4e-01	1.2e-02	3.3e-06	1.3e-05
64^2	11620	10^{-2}	6.4e+00	4.5e+00	4.7e-02	2.9e-02	3.1e-02
		10^{-4}	2.3e+01	1.1e+01	1.1e-01	2.1e-04	2.2e-04
128^2	54373	10^{-2}	6.8e+01	4.1e+01	4.4e-01	2.3e-03	3.2e-02
		10^{-4}	3.1e+02	1.0e+02	1.2e+00	8.0e-05	2.0e-03
256^2	249074	10^{-2}	6.8e+02	3.0e+02	2.2e+00	5.5e-02	5.4e-02
		10^{-4}	3.4e+03	9.0e+02	1.7e+01	2.9e-03	2.8e-03
512^2	1122019	10^{-2}	6.0e+03	2.0e+03	3.4e+01	6.1e-03	5.7e-03
		10^{-4}	2.6e+04	6.9e+03	9.2e+01	4.6e-04	4.2e-04

Table 5.3: Construction time, factorization time, solve time, relative residual and relative error of the direct solver on a polar grid. $M \approx 0.47N \log_4(N)$.

N	M	method	t_{pre}	t_{iter}	n_{iter}	r_s	e_s
32 ²	2397	lsqr	-	1.3e+00	470	8.9e-13	6.8e-12
		plsqr ($\varepsilon = 10^{-2}$)	9.6e-01	1.9e-01	9	2.9e-13	9.9e-12
		plsqr ($\varepsilon = 10^{-4}$)	1.7e+00	1.5e-01	4	7.7e-16	6.5e-15
64 ²	11620	lsqr	-	2.9e+00	500	3.8e-03	4.0e-03
		plsqr ($\varepsilon = 10^{-2}$)	1.1e+01	2.2e+00	21	1.5e-11	1.6e-11
		plsqr ($\varepsilon = 10^{-4}$)	3.3e+01	1.4e+00	5	8.9e-11	9.3e-11
128 ²	54373	lsqr	-	5.6e+00	500	2.2e-05	1.6e-03
		plsqr ($\varepsilon = 10^{-2}$)	1.1e+02	2.8e+01	38	2.6e-13	1.4e-11
		plsqr ($\varepsilon = 10^{-4}$)	4.1e+02	1.5e+01	10	2.2e-13	6.7e-11
256 ²	249074	lsqr	-	1.7e+01	500	1.9e-03	1.8e-03
		plsqr ($\varepsilon = 10^{-2}$)	9.7e+02	1.3e+02	37	1.3e-11	1.2e-11
		plsqr ($\varepsilon = 10^{-4}$)	4.3e+03	1.0e+02	12	5.7e-11	5.5e-11
512 ²	1122019	lsqr	-	6.5e+01	500	4.8e-04	4.6e-04
		plsqr ($\varepsilon = 10^{-2}$)	8.0e+03	4.2e+02	24	2.2e-12	2.1e-12
		plsqr ($\varepsilon = 10^{-4}$)	3.3e+04	4.2e+02	10	2.5e-12	2.4e-12

Table 5.4: Time cost, iteration number, relative residual and relative error of lsqr and plsqr on a polar grid. $M \approx 0.47N \log_4(N)$.

6 Conclusions and Future Directions

This paper presents a direct solver for the 2D type-II NUDFT based on the HSS matrix approximation $\tilde{\mathbf{G}}$ of the transformed NUDFT matrix $\mathbf{G} = \mathbf{A}\mathbf{F}^{-1}$. The construction of the HSS approximation is based on the kernel matrix expression of \mathbf{G} and the proxy surface technique, resulting in the complexity of $\mathcal{O}(M + N^{3/2} \log^3 N)$. The URV factorization is then applied to $\tilde{\mathbf{G}}$, which also takes $\mathcal{O}(M + N^{3/2} \log^3 N)$ time. Finally, the solution can be obtained using the URV factorization of $\tilde{\mathbf{G}}$ and two-dimensional FFT, which takes $\mathcal{O}(M + N \log^3 N)$ time. The proposed method can also be used as a preconditioner for iterative solvers. Numerical results show the efficiency and accuracy of our direct solver for the 2D type-II NUDFT with both random and polar grids. It also demonstrates the effectiveness of the proposed solver as a preconditioner for iterative solvers, especially when the grid is ill-conditioned.

There are several future directions. The first is to extend the proposed method to the 2D type-III NUDFT. This could be done by using the same idea proposed in [23]. The second is to give a rigorous discussion to the estimation of the numerical rank of the HSS blocks of \mathbf{G} . The discussion in Section 3.1 provides an intuition of the numerical rank of the HSS blocks. To achieve this we need a numerical rank version of Lemma 3.1, which is not straightforward. The last is about the extension to higher dimensions. Our method factorizes \mathbf{A} approximately as a product of an HSS matrix $\tilde{\mathbf{G}}$ and a fast transform \mathbf{F} . As we have shown intuitively and numerically, the HSS rank is $r = \mathcal{O}(\sqrt{N} \log N)$, which is quite large compared to the 1D case, where the HSS rank in 1D is $r = \mathcal{O}(\log N)$. Moreover, by the same discussion one can also prove that in 3D case the HSS rank will be $r = \mathcal{O}(N^{2/3} \log N)$, which is even larger. Consequently, the complexity of our algorithm in 3D case will be $\mathcal{O}(N^2 \log^\alpha N)$ for some $\alpha > 0$, which is less efficient, especially when N is large. Therefore, it is interesting to investigate whether there are other data-sparse approximations of \mathbf{G} with smaller rank, which can further reduce the complexity of the algorithm. For example, the *butterfly factorization* [8, 24] for Fourier integral operators may be a good candidate, which achieves a quasi-linear complexity $\mathcal{O}(N \log N)$ for the representation of the matrix when $M = N$.

References

- [1] C. ANDERSON AND M. D. DAHLEH, *Rapid computation of the discrete Fourier transform*, SIAM Journal on Scientific Computing, 17 (1996), pp. 913–919.
- [2] S. BAGCHI AND S. K. MITRA, *The nonuniform discrete Fourier transform and its applications in signal processing*, Springer US, 2012.
- [3] A. H. BARNETT, J. MAGLAND, AND L. AF KLINTEBERG, *A parallel nonuniform fast fourier transform library based on an “exponential of semicircle” kernel*, SIAM Journal on Scientific Computing, 41 (2019), pp. C479–C504.
- [4] B. BECKERMANN AND A. TOWNSEND, *On the singular values of matrices with displacement structure*, SIAM Journal on Matrix Analysis and Applications, 38 (2017), pp. 1227–1248.
- [5] P. BENNER, R.-C. LI, AND N. TRUHAR, *On the ADI method for Sylvester equations*, J. Comput. Appl. Math., 233 (2009), pp. 1035–1045.
- [6] S. BÖRM, L. GRASEDYCK, AND W. HACKBUSCH, *Introduction to hierarchical matrices with applications*, Engineering Analysis with Boundary Elements, 27 (2003), pp. 405–422.
- [7] M. BOURGEOIS, F. T. A. W. WAJER, D. VAN ORMONDT, AND D. GRAVERON-DEMILLY, *Reconstruction of MRI images from non-uniform sampling and its application to intrascan motion correction in functional MRI*, Birkhäuser Boston, Boston, MA, 2001, pp. 343–363.
- [8] E. CANDÈS, L. DEMANET, AND L. YING, *A fast butterfly algorithm for the computation of Fourier integral operators*, Multiscale Model. Simul., 7 (2009), pp. 1727–1750.
- [9] S. CHANDRASEKARAN, M. GU, AND T. PALS, *A fast ULV decomposition solver for hierarchically semiseparable representations*, SIAM Journal on Matrix Analysis and Applications, 28 (2006), pp. 603–622.
- [10] H. CHENG, Z. GIMBUTAS, P.-G. MARTINSSON, AND V. ROKHLIN, *On the compression of low rank matrices*, SIAM Journal on Scientific Computing, 26 (2005), pp. 1389–1404.
- [11] J. W. COOLEY AND J. W. TUKEY, *An algorithm for the machine calculation of complex Fourier series*, Mathematics of Computation, 19 (1965), pp. 297–301.
- [12] E. CORONA, P.-G. MARTINSSON, AND D. ZORIN, *An $O(N)$ direct solver for integral equations on the plane*, Applied and Computational Harmonic Analysis, 38 (2015), pp. 284–317.
- [13] A. DUTT AND V. ROKHLIN, *Fast Fourier transforms for nonequispaced data*, SIAM Journal on Scientific Computing, 14 (1993), pp. 1368–1393.
- [14] ———, *Fast Fourier transforms for nonequispaced data, II*, Applied and Computational Harmonic Analysis, 2 (1995), pp. 85–100.
- [15] M. FENN, S. KUNIS, AND D. POTTS, *On the computation of the polar FFT*, Appl. Comput. Harmon. Anal., 22 (2007), pp. 257–263.
- [16] L. GREENGARD AND J.-Y. LEE, *Accelerating the nonuniform fast Fourier transform*, SIAM Review, 46 (2004), pp. 443–454.
- [17] M. GU AND S. C. EISENSTAT, *Efficient algorithms for computing a strong rank-revealing QR factorization*, SIAM Journal on Scientific Computing, 17 (1996), pp. 848–869.
- [18] W. HACKBUSCH, *A sparse matrix arithmetic based on \mathcal{H} -matrices. Part I: Introduction to \mathcal{H} -matrices*, Computing, 62 (1999), pp. 89–108.
- [19] S. JIANG AND L. GREENGARD, *A dual-space multilevel kernel-splitting framework for discrete and continuous convolution*, Commun. Pure Appl. Math., 78 (2025), pp. 1086–1143.
- [20] M. KIRCHEIS AND D. POTTS, *Direct inversion of the nonequispaced fast Fourier transform*, Linear Algebra and its Applications, 575 (2019), pp. 106–140.

- [21] ———, *Fast and direct inversion methods for the multivariate nonequispaced fast Fourier transform*, Front. Appl. Math. Stat., Volume 9 (2023).
- [22] J.-Y. LEE AND L. GREENGARD, *The type 3 nonuniform FFT and its applications*, Journal of Computational Physics, 206 (2005), pp. 1–5.
- [23] Y. LI AND J. LIU, *A superfast direct solver for type-III inverse nonuniform discrete Fourier transform*, arxiv, (2025).
- [24] Y. LI, H. YANG, E. R. MARTIN, K. L. HO, AND L. YING, *Butterfly factorization*, Multiscale Model. Simul., 13 (2015), pp. 714–732.
- [25] P.-G. MARTINSSON, *Fast direct solvers for elliptic PDEs*, Fast direct solvers for elliptic partial differential equations, Society for Industrial Applied Mathematics, 2019.
- [26] P.-G. MARTINSSON AND V. ROKHLIN, *A fast direct solver for boundary integral equations in two dimensions*, Journal of Computational Physics, 205 (2005), pp. 1–23.
- [27] C. C. PAIGE AND M. A. SAUNDERS, *Lsqr: An algorithm for sparse linear equations and sparse least squares*, ACM Trans. Math. Softw., 8 (1982), pp. 43–71.
- [28] D. POTTS, G. STEIDL, AND M. TASCHE, *Fast Fourier transforms for nonequispaced data: A tutorial*, Birkhäuser Boston, Boston, MA, 2001, pp. 247–270.
- [29] D. RUIZ-ANTOLÍN AND A. TOWNSEND, *A nonuniform fast Fourier transform based on low rank approximation*, SIAM Journal on Scientific Computing, 40 (2018), pp. A529–A547.
- [30] H. WILBER, E. N. EPPERLY, AND A. H. BARNETT, *Superfast direct inversion of the nonuniform discrete fourier transform via hierarchically semiseparable least squares*, SIAM Journal on Scientific Computing, (2025), pp. A1702–A1732.
- [31] Y. XI, J. XIA, S. CAULEY, AND V. BALAKRISHNAN, *Superfast and stable structured solvers for Toeplitz least squares via randomized sampling*, SIAM Journal on Matrix Analysis and Applications, 35 (2014), pp. 44–72.
- [32] J. XIA, S. CHANDRASEKARAN, M. GU, AND X. S. LI, *Fast algorithms for hierarchically semiseparable matrices*, Numerical Linear Algebra with Applications, 17 (2010), pp. 953–976.