

NFFT.jl: GENERIC AND FAST JULIA IMPLEMENTATION OF THE NONEQUIDISTANT FAST FOURIER TRANSFORM

TOBIAS KNOPP*, MARIJA BOBERG* , AND MIRCO GROSSER*

Abstract. The non-equidistant fast Fourier transform (NFFT) is an extension of the famous fast Fourier transform (FFT) that can be applied to non-equidistantly sampled data in time/space or frequency domain. It is an approximative algorithm that allows to control the approximation error in such a way that machine precision is reached while keeping the algorithmic complexity in the same order as a regular FFT. The NFFT plays a major role in many signal processing applications and has been intensively studied from a theoretical and computational perspective. The fastest CPU implementations of the NFFT are implemented in the low-level programming languages C and C++ and require a compromise between code generalizability, code readability, and code efficiency. The programming language Julia promises new opportunities in optimizing these three conflicting goals. In this work we show that Julia indeed allows to develop an NFFT implementation, which is completely generic, dimension-agnostic and requires about 2–3 times less code than the other famous libraries NFFT3 and FINUFFT while still being one of the fastest NFFT implementations developed to date.

Key words. Nonequidistant Fast Fourier Transform, Fast Implementation, Multi-Threading, Julia

AMS subject classifications. 65T50, 65T40, 65Y05, 68N01

1. Introduction. The Fourier transform plays an important role in many mathematical applications in particular those involving signal processing. Convolutions in time/spatial domain can be expressed as multiplications in frequency domain, which is often faster and allows one to invert convolutions in analytical form. In practice, signals are usually discrete and of finite length such that the discrete Fourier transform (DFT) needs to be applied. It maps a number of N samples in time/spatial domain to N samples in frequency domain and requires $\mathcal{O}(N^2)$ operations, which is often too expensive to apply the DFT in its ordinary form. The fast Fourier transform (FFT) [5] allows one to carry out the DFT in only $\mathcal{O}(N \log N)$ operations and enables many applications for which the DFT would be too expensive. Its high impact in various applications makes the FFT one of the most important numerical algorithms developed in the 20th century. One of the fastest FFT software libraries is the FFTW [9], which is used in most scientific applications.

An important limitation of the FFT is that it requires the input and output signals to be equidistantly sampled. This makes it unusable for some important applications, e.g., in optical coherence tomography [10] and magnetic resonance imaging [7, 17]. In this work we focus on the case, where the signal in one domain is sampled in a non-equidistant manner. In such cases, simple interpolation, to map from non-equidistant to equidistant points, leads to large numerical errors, which intuitively can be explained by the fact that local (interpolation) errors lead to global errors in the reciprocal domain. This motivated the development of the non-equidistant fast Fourier transform (NFFT), which allows to keep the approximation error below the floating point precision and thus can be used like an exact algorithm in most circumstances. It was developed in the 1990s and 2000s by various researchers [6, 3, 1, 26, 28, 20] with a special focus on developing error bounds that allow to predict the approximation error in dependence of NFFT hyperparameters. One important ingredient of the NFFT is

*University Medical Center Hamburg-Eppendorf and Hamburg University of Technology, Hamburg, Germany (t.knopp@uke.de, <http://www.tuhh.de/ibi>).

the use of a window function, which is convolved with the irregular data and sampled at equidistant points. In the past, several window functions have been proposed of which the Kaiser-Bessel function provides the highest accuracy, since it is well localized in time/spatial and frequency domain. Quite recently, new error bounds have been derived for different window functions [22] including the Kaiser-Bessel window [21].

Beside this theoretical work, there is also a strong demand for high-quality fast software packages implementing the NFFT. While the algorithm itself has a rather simple mathematical notation, a straight forward textbook implementation would suffer from suboptimal performance and be orders of magnitude slower than a tuned software library. For this reason, several high-performance NFFT libraries like NFFT3 [13] and FINUFFT [2] have been developed in the past. NFFT3 was developed in the late 2000s, is written in C and supports multi-threading. In addition to the NFFT it offers algorithms for the non-equidistant fast cosine transform (NDCT) and the non-equidistant fast sine transform (NDST). Two important features of NFFT3 are the ability to apply the algorithm to arbitrary dimensional signals and the ability to cache window function evaluations for faster computation. FINUFFT was developed in the late 2010s, is written in C++ and provides two major innovations: First, it uses a new window function that can be evaluated in a fast manner. Second, it uses a new multi-threading approach for the adjoint NFFT using a block-partitioning strategy. In combination with highly tuned C++ code, these two innovations make FINUFFT one of the fastest NFFT implementations developed to date. The NFFT has also been implemented on graphical processing units (GPU) [25, 18, 14, 29, 24], which can give additional speedups over CPU implementations but is not in the scope of the current paper.

When developing a new NFFT software library, there is a certain design space to be explored. An implementation can

1. be generic, i.e., allowing for different floating point types (`Float32/Float64`),
2. be dimension-agnostic,
3. allow changing the window function,
4. allow changing the precomputation strategy,
5. be fast/multi-threaded,
6. be readable/maintainable,
7. be re-usable/binding friendly.

Since not all of these goals can be achieved at the same time, certain compromises must be made, many of which also depend on the programming language being used. Both NFFT3 and FINUFFT are implemented in C/C++ and made the following design decisions:

- Both libraries are partly generic, i.e., they allow for `Float32` and `Float64` floating point types. To avoid manual code duplication, the number type is defined using a C macro allowing for automatic code copies and textual replacements. In fact, most of the NFFT3 code is written in macro form to avoid code duplication.
- It is difficult to achieve a fast and dimension-independent implementation in C at the same time. This is because a static dimension, known at compile time, allows the compiler to generate much more efficient machine code. In particular nested for loops slow down code execution significantly if they are emulated in a dimension-agnostic fashion during runtime. Both libraries solve this problem by independent and redundant implementation of specific dimensions (NFFT3: 1D–5D, FINUFFT: 1D–3D). NFFT3 has an additional dynamic and slow fallback for higher dimensions while FINUFFT is restricted

to 1D–3D. Thus, speed and dimension-agnosticity can be combined in C but this decreases the code readability/maintainability.

- Hardcoding the window function can lead to additional speedups since it avoids a dynamic dispatch in a hot code path. For this reason the window function is a compile time option in NFFT3 while it is hardcoded in FINUFFT.

The purpose of this work is to develop a new NFFT package that makes less trade-offs than its C/C++ counterparts. Our package, named NFFT.jl, is implemented in the programming language Julia [4], which was first published in 2012 and has since then been developed into a popular programming language for scientific computing. Julia code can be both high-level like Matlab/Python and low-level like C/C++. The resulting machine code is usually as fast as comparable C/C++ machine code. This is achieved by a proper language design and the use of a just-in-time (JIT) compiler that allows to dynamically generate efficient machine code during runtime. Nowadays, Julia also supports multi-threading, which originates back to the prototype published in [15].

This paper outlines the key implementation aspects of NFFT.jl and makes an in-depth performance analysis compared to NFFT3 and FINUFFT. Since the architecture of NFFT.jl is rather flexible, we were able to implement several features that are only available in either NFFT3 or FINUFFT. This allows us to compare and benchmark certain implementation strategies within the same code base. We also make algorithmical improvements to the strategies proposed in [2], making NFFT.jl one of the fastest NFFT implementation developed to date.

2. Mathematical Description. We start with a mathematical description of the NFFT. Throughout this work, we stick to a multi-dimensional formulation and keep the notation close to the one used in [13]. Before discussing the NFFT, we first introduce the underlying mathematical transform to be computed: the non-equidistant discrete Fourier transform (NDFT).

2.1. NDFT. We let $\mathbf{N} \in \mathbb{N}^D$ be the size of the NDFT in the equidistant sampling domain and $D \in \mathbb{N}$ be its dimensionality. We further use the multi-dimensional index set

$$(2.1) \quad I_{\mathbf{N}} := \mathbb{Z}^D \cap \prod_{d=1}^D \left[-\frac{N_d}{2}, \frac{N_d}{2} \right)$$

representing all equidistant sampling points. The subindex n_d covers $-\frac{N_d}{2}, \dots, \frac{N_d}{2} - 1$ for even N_d and $-\frac{N_d-1}{2}, \dots, \frac{N_d-1}{2}$ for odd N_d . The corresponding signal is denoted by $f_{\mathbf{n}} \in \mathbb{C}$, $\mathbf{n} \in I_{\mathbf{N}}$. In the frequency domain, the signal is sampled at non-equidistant sampling points $\mathbf{k}_j = (k_{d,j})_{d=1}^D \in \mathbb{T}^D$, $j = 1, \dots, J$ with $J \in \mathbb{N}$ and $\mathbb{T} := [-1/2, 1/2)$. The resulting Fourier coefficients are denoted by $\hat{f}_j \in \mathbb{C}$, $j = 1, \dots, J$. The (direct) NDFT is then defined as

$$(2.2) \quad \hat{f}_j := \sum_{\mathbf{n} \in I_{\mathbf{N}}} f_{\mathbf{n}} e^{-2\pi i \mathbf{n} \cdot \mathbf{k}_j}, \quad j = 1, \dots, J \quad (\text{equidistant to non-equidistant}),$$

where $\mathbf{n} \cdot \mathbf{k}_j$ is the standard inner product between the two sampling points. This transformation is also known as the type-2 NDFT [2].

Remark 2.1. There are different conflicting definitions of the NDFT that consider the input signal to be sampled either in frequency [13] or in spatial domain [8].

Although we largely follow the notation of [13] we decided to switch frequency and spatial domain to make the NDFT consistent with the DFT, which is commonly defined to map from spatial to frequency domain. We note that this change only affects notation and that the computations remain the same.

The NDFT has an associated adjoint that can be expressed as

$$(2.3) \quad y_{\mathbf{n}} := \sum_{j=1}^J \hat{f}_j e^{2\pi i \mathbf{n} \cdot \mathbf{k}_j}, \quad \mathbf{n} \in I_{\mathbf{N}} \quad (\text{non-equidistant to equidistant}).$$

This is also named the type-1 NDFT [2] and it is related to the backward DFT. On purpose, we used the new variable $y_{\mathbf{n}}$ instead of $f_{\mathbf{n}}$ since the adjoint NDFT is in general not the inverse of the NDFT. This only holds true for special cases, i.e., for equidistant sampling nodes and when introducing a suitable normalization factor.

The algorithmic complexity for a direct computation of the NDFT and its adjoint is $\mathcal{O}(J|I_{\mathbf{N}}|)$ where $|I_{\mathbf{N}}| := \prod_{d=1}^D N_d$. This is because it requires two nested for loops, one over the J sampling points \mathbf{k}_j and one over the $|I_{\mathbf{N}}|$ equidistant sampling points \mathbf{n} .

2.2. NFFT. We next introduce the fast realization of the NDFT. There are different approaches for accelerating the NDFT all having in common to use some sort of approximation to enable the usage of one or several ordinary FFTs. In this work, we consider the most popular and widely used approach, which uses convolutions to grid the non-equidistant signal to an equidistant signal enabling FFT usage. The basic idea of this approach is to exploit the convolution theorem, which states that a convolution in frequency domain corresponds to a multiplication in spatial domain. To exploit this, one introduces an artificial convolution in the non-equidistant sampling domain with a window function $\hat{\varphi}$. This is corrected in the equidistant sampling domain by dividing by the window's inverse Fourier transform φ , i.e., applying a deconvolution. The key point is that the artificial convolution allows one to switch from equidistant to non-equidistant sampling points and makes it possible to first use the FFT to switch domains and then apply the resampling. Since the operation being applied is not shift-invariant in the discrete and non-equidistant setting we name it *resampling* to avoid confusing it with a regular continuous or discrete convolution. The grid on which the FFT is applied is chosen to be larger than the grid on which the equidistant input signal is sampled. To this end, a so-called oversampling factor $\sigma > 1$ is introduced and an oversampled grid of size $\tilde{\mathbf{N}} = \sigma \mathbf{N}$ is considered. σ impacts the accuracy of the NFFT and in practice it is chosen in the range [1.25, 2] with 2 being the common default value. Smaller values decrease the accuracy of the NFFT but reduce the memory requirement and are therefore commonly used for large transforms of high dimensionality.

In summary, the NFFT consists of three steps:

1. resampling correction in equidistant domain
2. fast Fourier transform
3. resampling to map from equidistant to non-equidistant domain

2.2.1. Direct NFFT. We next provide a short derivation of the NFFT and start with the D -dimensional window function $\hat{\varphi} : \mathbb{R}^D \rightarrow \mathbb{R}$. The latter is based on a one-dimensional version $\hat{\varphi}_{\text{Base}} : \mathbb{R} \rightarrow \mathbb{R}$ using the tensor product

$$(2.4) \quad \hat{\varphi}(\mathbf{k}) = \hat{\varphi}(k_1, \dots, k_D) = \prod_{d=1}^D \hat{\varphi}_{\text{Base}}\left(\frac{\tilde{N}_d}{m} k_d\right).$$

Here, $m \ll N_d$ is the second hyperparameter of the NFFT that controls the width of the window function. The window function $\hat{\varphi}$ is chosen in such a way that it is even and well localized in spatial and frequency domain. Due to the tensor product structure of $\hat{\varphi}$, its inverse Fourier transform φ can be expressed as a tensor product as well:

$$(2.5) \quad \varphi(\mathbf{n}) = \varphi(n_1, \dots, n_D) = \prod_{d=1}^D \frac{m}{\tilde{N}_d} \varphi_{\text{Base}} \left(\frac{m}{\tilde{N}_d} n_d \right).$$

Moreover, φ is real because of the even symmetry of $\hat{\varphi}$. Finally, we require $\varphi(\mathbf{n}) \neq 0$ for $\mathbf{n} \in I_{\mathbf{N}}$, which is fulfilled for the commonly used window functions.

The ansatz for deriving the NFFT is based on the inverse Fourier transform

$$(2.6) \quad \varphi(\mathbf{n}) = \int_{\mathbb{R}^D} \hat{\varphi}(\mathbf{k}) e^{2\pi i \mathbf{n} \cdot \mathbf{k}} d\mathbf{k} = \sum_{\mathbf{r} \in \mathbb{Z}^D} \int_{\mathbb{T}^D} \hat{\varphi}(\mathbf{k} + \mathbf{r}) e^{2\pi i \mathbf{n} \cdot (\mathbf{k} + \mathbf{r})} d\mathbf{k}.$$

When considering $\mathbf{n} \in \mathbb{Z}^D$ we can exploit the periodicity of the complex exponential ($e^{2\pi i \mathbf{n} \cdot (\mathbf{k} + \mathbf{r})} = e^{2\pi i \mathbf{n} \cdot \mathbf{k}}$ for $\mathbf{r} \in \mathbb{Z}^D$) and introduce the one-periodization of the window function $\tilde{\varphi}(\mathbf{k}) := \sum_{\mathbf{r} \in \mathbb{Z}^D} \hat{\varphi}(\mathbf{k} + \mathbf{r})$ yielding

$$\varphi(\mathbf{n}) = \int_{\mathbb{T}^D} \tilde{\varphi}(\mathbf{k}) e^{2\pi i \mathbf{n} \cdot \mathbf{k}} d\mathbf{k}.$$

Applying the substitution $\mathbf{k} \rightarrow \mathbf{k} - \mathbf{k}'$ results in

$$(2.7) \quad \varphi(\mathbf{n}) = \int_{\mathbb{T}^D} \tilde{\varphi}(\mathbf{k} - \mathbf{k}') e^{2\pi i \mathbf{n} \cdot (\mathbf{k} - \mathbf{k}')} d\mathbf{k}'.$$

Dividing by $\varphi(\mathbf{n})$ and $e^{2\pi i \mathbf{n} \cdot \mathbf{k}}$ yields

$$(2.8) \quad e^{-2\pi i \mathbf{n} \cdot \mathbf{k}} = \frac{1}{\varphi(\mathbf{n})} \int_{\mathbb{T}^D} \tilde{\varphi}(\mathbf{k} - \mathbf{k}') e^{-2\pi i \mathbf{n} \cdot \mathbf{k}'} d\mathbf{k}'.$$

Now we approximate the integral on the right hand side using a rectangular quadrature rule with $\tilde{\mathbf{N}} = \sigma \mathbf{N}$ sampling nodes yielding

$$(2.9) \quad e^{-2\pi i \mathbf{n} \cdot \mathbf{k}} \approx \frac{1}{|I_{\tilde{\mathbf{N}}}|\varphi(\mathbf{n})} \sum_{\mathbf{l} \in I_{\tilde{\mathbf{N}}}} \tilde{\varphi}(\mathbf{k} - \mathbf{l} \oslash \tilde{\mathbf{N}}) e^{-2\pi i \mathbf{n} \cdot (\mathbf{l} \oslash \tilde{\mathbf{N}})}$$

where \oslash denotes the element-wise division. This means, we can approximate a complex exponential sampled at any $\mathbf{k} \in \mathbb{T}^D$ using a sum of shifted complex exponentials.

Inserting the approximation (2.9) into the NDFT (2.2) yields

$$(2.10) \quad \hat{f}_j \approx \underbrace{\sum_{\mathbf{l} \in I_{\tilde{\mathbf{N}}}} \tilde{\varphi}(\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}})}_{\text{resampling}} \underbrace{\sum_{\mathbf{n} \in I_{\mathbf{N}}} \underbrace{\frac{f_{\mathbf{n}}}{|I_{\tilde{\mathbf{N}}}|\varphi(\mathbf{n})}}_{\text{resampling correction}} e^{-2\pi i \mathbf{n} \cdot (\mathbf{l} \oslash \tilde{\mathbf{N}})}}_{\text{DFT}}.$$

One can see that the resampling (last step) allows us to use the FFT (second step) since the sampling nodes $\mathbf{l} \oslash \tilde{\mathbf{N}}$ are now equidistant. To account for the resampling, the inner part first applies a resampling correction.

Until now, we have not yet made the algorithm any faster than a direct evaluation of the NDFT since the resampling requires $\mathcal{O}(J|I_{\tilde{\mathbf{N}}}|)$ operations. To do so, we need to exploit that the window function $\hat{\varphi}$ is well localized and thus close to zero for most of the evaluations performed during the resampling. To exploit this formally, we truncate $\hat{\varphi}$ at $\pm \mathbf{m} \oslash \tilde{\mathbf{N}}$ with $\mathbf{m} = (m)_{d=1}^D$ and define the truncated window function

$$(2.11) \quad \hat{\psi}(\mathbf{k}) := \begin{cases} \hat{\varphi}(\mathbf{k}) & \text{for } \mathbf{k} \in \prod_{d=1}^D [-\frac{m}{\tilde{N}_d}, \frac{m}{\tilde{N}_d}) \\ 0 & \text{otherwise} \end{cases}$$

with support $\text{supp } \hat{\psi} = \prod_{d=1}^D [-\frac{m}{\tilde{N}_d}, \frac{m}{\tilde{N}_d})$. In the same way $\hat{\psi}_{\text{Base}}$ is defined to be $\hat{\varphi}_{\text{Base}}$ truncated to $[-1, 1)$. From (2.10), it can be seen that the resampling step is based on the one-periodization $\tilde{\varphi}$. To accelerate its computation, we thus require the corresponding one-periodization $\tilde{\psi}(\mathbf{k}) := \sum_{\mathbf{r} \in \mathbb{Z}^D} \hat{\psi}(\mathbf{k} + \mathbf{r})$. Replacing $\tilde{\varphi}(\mathbf{k})$ by $\tilde{\psi}(\mathbf{k})$ in (2.10) yields

$$(2.12) \quad \hat{f}_j \approx \sum_{\mathbf{l} \in I_{\tilde{\mathbf{N}}, m}(\mathbf{k}_j)} \tilde{\psi}(\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}}) \sum_{\mathbf{n} \in I_{\tilde{\mathbf{N}}}} \frac{f_{\mathbf{n}}}{|I_{\tilde{\mathbf{N}}}| \varphi(\mathbf{n})} e^{-2\pi i \mathbf{n} \cdot (\mathbf{l} \oslash \tilde{\mathbf{N}})},$$

with the multi-index set

$$(2.13) \quad I_{\tilde{\mathbf{N}}, m}(\mathbf{k}) := \left\{ \mathbf{l} \in I_{\tilde{\mathbf{N}}} : \exists p \in \mathbb{Z} (-\mathbf{m} \oslash \tilde{\mathbf{N}} \leq \mathbf{k} - \mathbf{l} \oslash \tilde{\mathbf{N}} + p\mathbf{1} < \mathbf{m} \oslash \tilde{\mathbf{N}}) \right\},$$

where $\mathbf{1} = (1)_{d=1}^D$. This is the set of indices \mathbf{l} for which $\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}} \in \text{supp}(\tilde{\psi})$ and in turn other indices can be dropped without changing the result of the computation. It has at most $2m^D$ elements, which is much less than the $\prod_{d=1}^D \tilde{N}_d$ elements of $I_{\tilde{\mathbf{N}}}$. In particular, the number of elements is independent of $\tilde{\mathbf{N}}$. For the implementation described in section 3 we require an alternative formulation of the index set $I_{\tilde{\mathbf{N}}, m}(\mathbf{k})$ that separates the individual dimensions:

$$I_{\tilde{\mathbf{N}}, m}(\mathbf{k}) = \prod_{d=1}^D I_{\tilde{N}_d, m}(k_d) \text{ with } I_{\tilde{N}_d, m}(k) = \{\omega_{n, m}(k, \ell) : \ell \in \{1, \dots, 2m\}\}.$$

Here, the function $\omega_{\tilde{N}, m} : \mathbb{T} \times \{1, \dots, 2m\} \rightarrow I_{\tilde{N}}$ with

$$\omega_{\tilde{N}, m}(k, \ell) = ([\tilde{N}(k \bmod 1) - m + \ell - 1] \bmod \tilde{N}) - \frac{\tilde{N}}{2}$$

allows us to calculate the index set directly taking account the index wrap due to the periodization of the window function $\tilde{\psi}$.

The direct NFFT is summarized in Algorithm 2.1:

1. In the first step, the resampling correction is applied. To this end, the input data $f_{\mathbf{n}}$ is divided by the inverse Fourier transform of the window function $\varphi(\mathbf{n})$ for each $\mathbf{n} \in I_{\tilde{\mathbf{N}}}$. The output of the resampling correction is stored in a new temporary vector $\mathbf{g} = (g_{\mathbf{n}})_{\mathbf{n} \in I_{\tilde{\mathbf{N}}}}$ that is defined on the fine grid $I_{\tilde{\mathbf{N}}}$. The first step of the NFFT thus also applies the zero padding necessary for the in-place FFT in the second step of the algorithm.
2. In the second step, an ordinary D -variate FFT is applied. It operates in-place but to keep the mathematical notation sound, we introduce the output vector $\hat{\mathbf{g}} = (\hat{g}_{\mathbf{l}})_{\mathbf{l} \in I_{\tilde{\mathbf{N}}}}$.
3. In the third step, the resampling is applied.

Algorithm 2.1 Direct NFFT

Input: $D, m \in \mathbb{N}$, $\mathbf{N}, \tilde{\mathbf{N}} \in \mathbb{N}^D$, with $\tilde{\mathbf{N}} = \sigma \mathbf{N}$ and $\sigma > 1$,
non-equidistant sampling points $\mathbf{k}_j \in \mathbb{T}^D, j = 1, \dots, J$,
equidistantly sampled signal $f_{\mathbf{n}} \in \mathbb{C}, \mathbf{n} \in I_{\mathbf{N}}$

- 1: **for** $\mathbf{n} \in I_{\tilde{\mathbf{N}}}$ **do**
- 2: $g_{\mathbf{n}} = \begin{cases} \frac{f_{\mathbf{n}}}{|I_{\tilde{\mathbf{N}}}| \varphi(\mathbf{n})} & \text{for } \mathbf{n} \in I_{\mathbf{N}} \\ 0 & \text{for } \mathbf{n} \in I_{\tilde{\mathbf{N}}} \setminus I_{\mathbf{N}} \end{cases} \quad \triangleright \text{resampling correction}$
- 3: **end for**
- 4: **for** $\mathbf{l} \in I_{\tilde{\mathbf{N}}}$ **do**
- 5: $\hat{g}_{\mathbf{l}} = \sum_{\mathbf{n} \in I_{\tilde{\mathbf{N}}}} g_{\mathbf{n}} e^{-2\pi i \mathbf{n} \cdot (\mathbf{l} \oslash \tilde{\mathbf{N}})} \quad \triangleright \text{FFT}$
- 6: **end for**
- 7: **for** $j = 1, \dots, J$ **do**
- 8: $\hat{f}_j = \sum_{\mathbf{l} \in I_{\tilde{\mathbf{N}}, m}(\mathbf{k}_j)} \hat{g}_{\mathbf{l}} \tilde{\psi}(\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}}) \quad \triangleright \text{resampling}$
- 9: **end for**

Output: non-equidistantly sampled signal $\hat{f}_j \in \mathbb{C}, j = 1, \dots, J$
Complexity: $\mathcal{O}(|I_{\tilde{\mathbf{N}}}| \log |I_{\tilde{\mathbf{N}}}| + m^D J)$

2.2.2. Adjoint NFFT. The adjoint NFFT is based on the same idea as the direct NFFT. Inserting (2.9) into (2.3) and replacing $\tilde{\varphi}$ by $\tilde{\psi}$ yields

$$(2.14) \quad y_{\mathbf{n}} \approx \underbrace{\frac{1}{|I_{\tilde{\mathbf{N}}}| \varphi(\mathbf{n})} \sum_{\mathbf{l} \in I_{\tilde{\mathbf{N}}}} \underbrace{\left(\sum_{j \in I_{\tilde{\mathbf{N}}, m}^{\mathbf{T}}(\mathbf{l})} \hat{f}_j \tilde{\psi}(\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}}) \right)}_{\text{resampling}}}_{\text{adjoint DFT}}}_{\text{resampling correction}} e^{2\pi i \mathbf{n} \cdot (\mathbf{l} \oslash \tilde{\mathbf{N}})}.$$

Here, the index set $I_{\tilde{\mathbf{N}}, m}^{\mathbf{T}}(\mathbf{l})$ is defined as

$$(2.15) \quad I_{\tilde{\mathbf{N}}, m}^{\mathbf{T}}(\mathbf{l}) := \left\{ j \in \{1, \dots, J\} : \exists p \in \mathbb{Z} (-\mathbf{m} \oslash \tilde{\mathbf{N}} \leq \mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}} + p\mathbf{1} < \mathbf{m} \oslash \tilde{\mathbf{N}}) \right\}$$

and again allows one to perform the inner resampling over only a subset of the original indices. One can also see that the adjoint NFFT applies the three steps in reverse to the direct NFFT. The adjoint NFFT is summarized in [Algorithm 2.2](#):

1. In the first step, the non-equidistantly sampled signal \hat{f}_j is convolved with the one-periodic window function $\tilde{\psi}$ yielding the equidistantly sampled signal $\hat{g}_{\mathbf{l}}$. The sum is again restricted to the subset of indices $I_{\tilde{\mathbf{N}}, m}^{\mathbf{T}}(\mathbf{l})$ at which $\tilde{\psi}$ is non-zero. However, the adjoint resampling needs a different evaluation order than indicated by the summation. The reason is that it is very in-efficient to determine $I_{\tilde{\mathbf{N}}, m}^{\mathbf{T}}(\mathbf{l})$ during the summation. Instead, one uses two for loops: the inner looping over $\mathbf{l} \in I_{\tilde{\mathbf{N}}, m}(\mathbf{k}_j)$ and the outer looping over the sampling points $\mathbf{k}_j, j = 1, \dots, J$. This change of evaluation order means that the direct

Algorithm 2.2 Adjoint NFFT

Input: $D, m \in \mathbb{N}$, $\mathbf{N}, \tilde{\mathbf{N}} \in \mathbb{N}^D$, with $\tilde{\mathbf{N}} = \sigma \mathbf{N}$ and $\sigma > 1$,
non-equidistant sampling points $\mathbf{k}_j \in \mathbb{T}^D, j = 1, \dots, J$,
non-equidistantly sampled signal $\hat{f}_j \in \mathbb{C}, j = 1, \dots, J$

1: **for** $\mathbf{l} \in I_{\tilde{\mathbf{N}}}$ **do**
2: $\hat{g}_{\mathbf{l}} = \sum_{j \in I_{\tilde{\mathbf{N}}, m}^{\top}(\mathbf{l})} \hat{f}_j \tilde{\psi}(\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}})$ \triangleright *adjoint resampling*
3: **end for**
4: **for** $\mathbf{n} \in I_{\tilde{\mathbf{N}}}$ **do**
5: $g_{\mathbf{n}} = \sum_{\mathbf{l} \in I_{\tilde{\mathbf{N}}}} \hat{g}_{\mathbf{l}} e^{2\pi i \mathbf{n} \cdot (\mathbf{l} \oslash \tilde{\mathbf{N}})}$ \triangleright *adjoint FFT*
6: **end for**
7: **for** $\mathbf{n} \in I_{\mathbf{N}}$ **do**
8: $y_{\mathbf{n}} = \frac{g_{\mathbf{n}}}{|I_{\tilde{\mathbf{N}}}| \varphi(\mathbf{n})}$ \triangleright *adjoint resampling correction*
9: **end for**

Output: equidistantly sampled signal $y_{\mathbf{n}} \in \mathbb{C}$
Complexity: $\mathcal{O}(|I_{\tilde{\mathbf{N}}}| \log |I_{\tilde{\mathbf{N}}}| + m^D J)$

and the adjoint resampling have a very similar structure. The only difference is that the summation over \mathbf{l} has no data dependency for the direct transform whereas the adjoint transform needs to perform additive vector updates on $\hat{\mathbf{g}}$, which cannot be performed concurrently for different nodes \mathbf{k}_j .

2. The second step is a D -variate adjoint FFT.
3. The third step is the resampling correction. Similar to the direct transform, only the subset of \mathbf{g} on the grid $I_{\mathbf{N}}$ needs to be considered in this step.

2.2.3. Matrix-Vector Notation. The NFFT can be written in matrix-vector notation, which is helpful for conceptual understanding and can also be used for actual implementations (see discussion about GPU implementations in section 6). In matrix-vector form, the NFFT can be expressed as

$$\hat{\mathbf{f}} = \mathbf{A} \mathbf{f} = \mathbf{B} \mathbf{F} \mathbf{D} \mathbf{f}$$

where $\mathbf{A} = (e^{-2\pi i \mathbf{n} \cdot \mathbf{k}_j})_{j=1, \dots, J; \mathbf{n} \in I_{\mathbf{N}}} \in \mathbb{C}^{J \times |I_{\mathbf{N}}|}$ is the NDFT matrix, $\mathbf{f} = (f_{\mathbf{n}})_{\mathbf{n} \in I_{\mathbf{N}}} \in \mathbb{C}^{|I_{\mathbf{N}}|}$ is the input vector, and $\hat{\mathbf{f}} = (\hat{f}_j)_{j=1}^J \in \mathbb{C}^J$ is the output vector. The first NFFT step is the multiplication with the generalized diagonal matrix

$$\mathbf{D} = \left(\delta_{\mathbf{l}, \mathbf{n}} \frac{1}{|I_{\tilde{\mathbf{N}}}| \varphi(\mathbf{n})} \right)_{\mathbf{l} \in I_{\tilde{\mathbf{N}}}; \mathbf{n} \in I_{\mathbf{N}}} \in \mathbb{R}^{|I_{\tilde{\mathbf{N}}}| \times |I_{\mathbf{N}}|}$$

where δ is the Kronecker symbol. The second step is the application of the Fourier matrix

$$\mathbf{F} = \left(e^{-2\pi i \mathbf{n} \cdot (\mathbf{l} \oslash \tilde{\mathbf{N}})} \right)_{\mathbf{n} \in I_{\tilde{\mathbf{N}}}; \mathbf{l} \in I_{\tilde{\mathbf{N}}}} \in \mathbb{C}^{|I_{\tilde{\mathbf{N}}}| \times |I_{\tilde{\mathbf{N}}}|}.$$

The last step is the multiplication with the sparse matrix

$$(2.16) \quad \mathbf{B} = \left(\tilde{\psi}(\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}}) \right)_{j=1, \dots, J; \mathbf{l} \in I_{\tilde{\mathbf{N}}}} \in \mathbb{R}^{J \times |I_{\tilde{\mathbf{N}}}|}$$

having at most $(2m)^D J$ non-zero entries. The adjoint NFFT in matrix-vector form is obtained by reversing the order of the three matrices and applying the adjoint operator to each step individually:

$$\mathbf{y} = \mathbf{A}^H \hat{\mathbf{f}} = \mathbf{D}^\top \mathbf{F}^H \mathbf{B}^\top \hat{\mathbf{f}}$$

2.2.4. Algorithmic Complexity. The algorithmic complexity of the direct NFFT and its adjoint is

$$\begin{aligned} \mathcal{O}(\underbrace{|I_{\tilde{\mathbf{N}}}|}_{\text{resampling correction}} + \underbrace{|I_{\tilde{\mathbf{N}}}| \log |I_{\tilde{\mathbf{N}}}|}_{\text{FFT}} + \underbrace{(2m)^D J}_{\text{resampling}}) &= \mathcal{O}(|I_{\tilde{\mathbf{N}}}| \log |I_{\tilde{\mathbf{N}}}| + m^D J) \\ &= \mathcal{O}\left(|I_{\mathbf{N}}| \log |I_{\mathbf{N}}| + \log\left(\frac{1}{\varepsilon}\right) J\right). \end{aligned}$$

Here, we assumed in the last step that σ is constant (e.g., $\sigma = 2$), and that the accuracy ε improves exponentially with the kernel parameter m . The latter is fulfilled for all commonly used window functions.

Beside this theoretical consideration, the actual performance of the NFFT and its individual steps highly depends on the dimensionality and the size of the problem. The resampling correction step is the cheapest one and can usually be neglected in terms of computation time. In the most common setting ($J \approx |I_{\mathbf{N}}|$), the FFT is the second fastest operation and the resampling is the primary bottleneck being up to one order of magnitude slower. This is the reason why NFFT implementations usually put most efforts into optimizing the resampling as much as possible. However, in sparse sampling settings ($J \ll |I_{\mathbf{N}}|$) the FFT can become a dominant factor if the resampling is done properly. In multi-threading applications it is thus important to parallelize all parts of the NFFT. Otherwise even the resampling correction can become a bottleneck.

3. Implementation. After formulating the NFFT and its adjoint in mathematical notation we next focus on our software package NFFT.jl. Looking back to the design goals sketched in the introduction we put the highest priority on a generic, dimension-agnostic and fast/multi-threaded implementation. The second most important property is to keep the code readable and maintainable. Finally, we also want our implementation to be flexible and allow to change window functions and precomputation strategies with low effort. Reusability in different programming languages was not a top priority during the design of NFFT.jl but we sketch potential strategies for using NFFT.jl in programming languages other than Julia in [section 6](#).

3.1. Example Usage. A typical example usage of NFFT.jl for a 2D transformation is outlined below:

```

1 using NFFT
2
3 D = 2                                # dimensionality
4 J = 32*32                            # number of sampling points
5 N = (32, 32)                        # input signal size

```

```

6 k = rand(D, J) .- 0.5           # sampling points in  $[-0.5, 0.5]^D$ 
7
8 p = NFFTPlan(k, N; m=4,  $\sigma$ =2) # create the NFFT plan
9
10 f = randn(Complex{Float64}, N) # signal to be transformed
11 fHat = p * f                    # compute direct NFFT
12 y = adjoint(p) * fHat          # compute adjoint NFFT

```

For simplicity, the input signal and the sampling nodes are initialized with random numbers. Based on the signal size N (line 5) and the sampling nodes \mathbf{k} (line 6), an `NFFTPlan` object is created in line 8. The constructor takes care of allocating all necessary memory for temporary arrays and performing precomputations to make later transformations as fast as possible. This precomputation approach is very common in scientific programming. The actual NFFT's are applied in line 11 (direct) and 12 (adjoint). Since the NFFT can be interpreted as a matrix-vector multiplication, `NFFT.jl` uses the method `*` to express the transformation. The method `adjoint` creates a lazy wrapper type that allows `*` to call the adjoint transformation. In this way the syntax is very close to the mathematical notation. The chosen interface also matches the common naming scheme and coding pattern used for linear transformations in the Julia ecosystem.

In our example, the output vector was allocated within the `*` method. To avoid this allocation one can use the interface

```

1 mul!(fHat, p, f)
2 mul!(y, adjoint(p), fHat)

```

which allows one to pass the output vector as the first argument. Internally, `*` is implemented as a small wrapper around `mul!`. Again, `mul!` is a standard method in Julia to express in-place linear transformations. Finally, there is also a high-level interface

```

1 nfft(k, f; m=4,  $\sigma$ =2)
2 nfft_adjoint(k, N, fHat; m=4,  $\sigma$ =2)

```

that automatically creates a plan before calling the low-level NFFT functions. This is convenient if the NFFT is applied only once.

3.2. Memory Management. `NFFTPlan` is a struct holding several temporary arrays required by the NFFT. In particular it holds:

- the temporary vector $\hat{\mathbf{g}}$.
- the forward and backward plan of the inner FFT, i.e., data structures formed during precomputation to facilitate the computation of FFTs.
- different index and data vectors related to the resampling (see [subsection 3.7](#)) and the resampling correction (see [subsection 3.5](#)).

The concrete number of data that the plan holds depends on the precomputation strategy (see [subsection 3.7.4](#)). Furthermore, the block-partitioning strategy discussed in [section subsection 3.7.2](#) also needs additional memory.

3.3. Generic Types. One important design goal of `NFFT.jl` is the ability to use different number types. In Julia, this can be done by introducing a type parameter T that can take on any floating point number type. We restrict all real data types (like

the input nodes k and the precomputed window function entries) to be of type T while complex values like the input and output vectors of the NFFT are of type `Complex{T}`. This forcing of a common number type is important for optimal performance, as the need for type promotion is removed. The type T is encoded in the `NFFTPlan` type and allows the JIT compiler to generate dedicated machine code for each number type. Index types are all stored as 64-bit integers, which is the native integer number type on most modern CPUs. In principle a flexibilization of the index type is straight forward but this will only be considered if the current restriction to 64-bit integers poses a real problem.

3.4. Dimension-Agnostic Implementation. The second important design goal of NFFT.jl is to be dimension-agnostic. This is achieved by an integer parameter D , which allows us to encode the dimension in the `NFFTPlan`. This ensures the generation of dedicated machine code for each dimension. For instance, the signal size N is stored as a fixed size tuple `N::NTuple{D,Int64}`. Besides these storage aspects, the real challenge is to implement the actual calculations in a dimension-agnostic way. Two possible solutions to tackle this challenge are:

1. Write dedicated implementations for each dimension D .
2. Implement multi-dimensional loops using iterators that treat D as a runtime parameter.

The first solution ensures maximum performance but leads to code duplication. This is problematic since code changes need to be done in all copies of the code thus increasing the maintainability cost. The second solution can be implemented with less code and without any duplication but it has the downside of being much slower because the compiler cannot emit fast machine code. Because of these downsides we considered neither of these two solutions. Instead, we exploit that Julia's architecture enables the implementation of fast and generic dimension-agnostic Cartesian iterators. During the development of Julia, two different Cartesian iterator types have been established. The first option named `CartesianIndices` is iterator-based and allows one to iterate over a D -dimensional array A with

```
1 for i in CartesianIndices(A)
2     A[i] = ...
3 end
```

Since the dimension D is a parameter of the array A , the JIT compiler is able to generate efficient code for `CartesianIndices`-based for loops. The second option, which was developed much earlier, is a macro-based solution available in the `Base.Cartesian` module. A dimension-agnostic loop using `Base.Cartesian` is formulated as

```
1 @nloops $D i A begin
2     @nref $D A i = ...
3 end
```

and generates the following code for $D = 3$:

```
1 for i_3 = axes(A, 3)
2     for i_2 = axes(A, 2)
3         for i_1 = axes(A, 1)
4             A[i_1, i_2, i_3] = ...
5         end
2     end
1 end
```

```

6   end
7   end

```

This looks exactly like a hand-written implementation for a given dimensionality with the important difference, that the code is automatically specialized for the array dimension D . In NFFT.jl, we mainly use the **CartesianIndices**-based solution, since it is easier to read and maintain. However, in cases where more flexibility and control is needed, we use the **Base.Cartesian** macros. In particular, we use it in situations where the inner and the outer for loop are handled differently. This allows us to apply multi-threading only to the outer for loop.

3.5. Resampling Correction. The resampling correction (and its adjoint) is a simple element-wise operation. Due to oversampling, the input and output array do not have the same size but the iteration only runs over $I_{\mathbf{N}}$. To this end, we use **Base.Cartesian** and multi-thread the most outer for loop.

One important implementation detail of the resampling correction is to take the indices of the later FFT into account. The FFT is usually defined with a sum running from 0 to $\tilde{N}_d - 1$ while the FFT within the NFFT needs to be performed on the grid $I_{\tilde{\mathbf{N}}}$ with indices running from $-\frac{\tilde{N}_d}{2}$ to $\frac{\tilde{N}_d}{2} - 1$. Accordingly, an **fftshift** needs to be performed. To avoid having to perform an extra step, we integrate the index mapping into the resampling correction step and thus only need to touch each point once.

The primary bottleneck of the resampling correction is the evaluation of the function φ . While this can often be neglected because the resampling is much more expensive, there are situations, such as 1D transforms with $J \ll |I_{\mathbf{N}}|$, where the resampling correction can become a bottleneck. Therefore we perform two optimizations:

1. We cache φ so that it can be reused when applying many NFFTs to different data. Here we can exploit the tensor product structure of φ and cache the vectors

$$\beta_d^{\text{tensor}} = (\beta_{n_d,d}^{\text{tensor}})_{n_d \in I_{N_d}}, \quad \beta_{n_d,d}^{\text{tensor}} = \frac{1}{m \varphi_{\text{Base}}\left(\frac{m}{\tilde{N}_d} n_d\right)}, \quad d = 1, \dots, D.$$

During the resampling correction we can then calculate $\frac{1}{|I_{\tilde{\mathbf{N}}}| \varphi(\mathbf{n})}$ on the fly using:

```

1: for  $n_D \in I_{N_D}$  do
2:    $\gamma_D \leftarrow \beta_{n_D,D}^{\text{tensor}}$ 
3:   for  $n_{D-1} \in I_{N_{D-1}}$  do
4:      $\gamma_{D-1} \leftarrow \gamma_D \beta_{n_{D-1},D-1}^{\text{tensor}}$ 
5:     ...
6:     for  $n_1 \in I_{N_1}$  do
7:        $\gamma_1 \leftarrow \gamma_2 \beta_{n_1,1}^{\text{tensor}}$ 
8:        $g_{\mathbf{n}} \leftarrow f_{\mathbf{n}} \gamma_1$ 
9:     end for
10:    ...
11:   end for
12: end for

```

Here, we used the assignment notation \leftarrow to indicate that the temporary variables γ_d are updated during the loop. This form of precomputation requires only $\mathcal{O}(\sum_{d=1}^D N_d)$ memory (and evaluations of φ), whereas a full precom-

putation has a requirement of $\mathcal{O}(\prod_{d=1}^D N_d)$. During the actual resampling correction, there is no real downside, since there is just one additional multiplication in the inner loop, which in practice is not necessarily slower because the bottleneck is the load and store operation acting on $g_{\mathbf{n}}$ and $f_{\mathbf{n}}$ in line 8. For completeness NFFT.jl also has an option for full precomputation, which is useful in the GPU prototype that we sketch in section [section 6](#).

2. While the number of evaluations of φ is greatly reduced in the multi-dimensional case, we observed a measurable influence for 1D transformations. We therefore optimize the precomputation of $\frac{1}{|I_{\tilde{\mathbf{N}}}| \varphi(\mathbf{n})}$ by approximating φ with a Chebyshev polynomial, which is much faster to evaluate than our default window (Kaiser-Bessel).

3.6. Fast Fourier Transform. The FFT is usually not part of the NFFT implementation itself. Instead one uses an existing high-performance FFT library. We chose the FFTW.jl package, which is a wrapper around the popular FFTW library [9]. An alternative is to use the binary compatible implementation from Intel’s Math Kernel Library (MKL), which can be changed via a runtime switch. Right now, NFFT.jl is hardcoded to use FFTW/MKL, while in principle it is possible to make this step exchangable using the interface package AbstractFFTs.jl.

The FFTW package provides a planner interface to split off precomputations and preallocations. In NFFT.jl, the `FFTWPlan` is integrated into the `NFFTPlan`, thus separating planning from computation. Since FFTW requires a dedicated plan for forward and backward transformation, both plans are precomputed and stored in the `NFFTPlan`. The number of FFTW threads is matched to the number of threads used for the resampling and its correction within the NFFT. FFTW.jl allows one to use the thread-pool of the Julia runtime, which enables nested parallelism, i.e., if multiple NFFT/FFT are called in parallel from different threads, there will be no over-commitment breaking down the performance.

3.7. Resampling. The most important and most challenging operation of any NFFT implementation is the resampling. While the mathematical formula looks straight forward, there can be slowdowns of more than one order of magnitude between textbook implementations and optimized ones. We next outline the most important aspects that should be taken into account.

3.7.1. Loop Optimization. To discuss the details of the implementation, we recapitulate the resampling steps, which are given by

$$(3.1) \quad \text{direct resampling: } \hat{f}_j = \sum_{\mathbf{l} \in I_{\tilde{\mathbf{N}},m}(\mathbf{k}_j)} \hat{g}_{\mathbf{l}} \tilde{\psi}(\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}}), \quad j = 1, \dots, J$$

$$(3.2) \quad \text{adjoint resampling: } \hat{g}_{\mathbf{l}} = \sum_{j \in I_{\tilde{\mathbf{N}},m}^{\top}(\mathbf{l})} \hat{f}_j \tilde{\psi}(\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}}), \quad \mathbf{l} \in I_{\tilde{\mathbf{N}}}$$

Since the direct and adjoint transform have a very similar structure, we can share most of the concepts and code discussed in this section. The operations (3.1) and (3.2) can either be memory bound because of the access to $\hat{g}_{\mathbf{l}}$, \hat{f}_j and \mathbf{k}_j or computational bound because of the calculation of $\tilde{\psi}(\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}})$. In [subsection 3.7.2](#) and [subsection 3.7.3](#) we discuss how to prevent cache misses, which substantially accelerates memory access time. In [subsection 3.7.4](#) we discuss different precomputation strategies for accelerating the computation of $\tilde{\psi}(\mathbf{k}_j - \mathbf{l} \oslash \tilde{\mathbf{N}})$.

When both aspects are carefully taken into account one can perform additional loop optimizations. While these are often applied automatically, appropriately written code is required for the compiler to do so. The first optimization is the usage of an (immutable) fixed-size array. This allows the compiler to allocate arrays on the stack and keep them in CPU registers. Furthermore, it enables loop unrolling and allows the compiler to use SIMD (single instruction multiple data) instructions for parallel processing on a single core. In Julia it is possible to create fixed-size arrays during runtime using tuple types. We therefore create D fixed-size vectors

$$\begin{aligned}\tilde{\psi}_{d,j}^{\text{local}} &:= \left(\tilde{\psi}_{\ell,d,j}^{\text{local}} \right)_{\ell=1}^{2m}, \quad d = 1, \dots, D \\ \tilde{\psi}_{\ell,d,j}^{\text{local}} &:= \tilde{\psi}_{\text{Base}} \left(\frac{1}{m} \left(\tilde{N}_d k_{d,j} - \omega_{\tilde{N}_d, m}(k_{d,j}, \ell) \right) \right)\end{aligned}$$

and calculate (3.1) by

$$(3.3) \quad \hat{f}_j = \sum_{\ell_D=1}^{2m} \tilde{\psi}_{\ell_D, D, j}^{\text{local}} \sum_{\ell_{D-1}=1}^{2m} \tilde{\psi}_{\ell_{D-1}, D-1, j}^{\text{local}} \cdots \sum_{\ell_1=1}^{2m} \tilde{\psi}_{\ell_1, 1, j}^{\text{local}} \hat{g}_{\mathbf{l}}, \quad \mathbf{l} = \left(\omega_{\tilde{N}_d, m}(k_{d,j}, \ell_d) \right)_{d=1}^D.$$

The nested product calculation in (3.3) exploits the tensor product structure of the window function and results in a very simple inner loop (over ℓ_1) that can be fully optimized by the compiler.

Remark 3.1. Local caching of the window function using the vectors $\tilde{\psi}_{d,j}^{\text{local}}$ should not be confused with the precomputation of $\tilde{\psi}$ discussed in subsection 3.7.4. This is because the local vector $\tilde{\psi}_{d,j}^{\text{local}}$ is calculated right before the resampling. Hence it is a local computation trick that helps the compiler to generate very efficient machine code. On the other hand, this caching can make precomputation unnecessary for larger dimensions. This is because the calculation of all non-zero entries of the \mathbf{B} matrix – which might require an expensive window evaluation – has an arithmetic complexity of $\mathcal{O}(mD)$ while the sum in (3.3) requires $\mathcal{O}(m^D)$ operations. With increasing D , the caching operation becomes negligible even for expensive window functions.

3.7.2. Block-Partitioning Motivation. The non-equidistant nature of the sampling points \mathbf{k}_j implies that there is no natural ordering of the sampling points in multiple dimensions. This means that subsequent sampling points $\mathbf{k}_j, \mathbf{k}_{j+1}$ can have a large distance in \mathbb{T}^D . If this is the case, they interact with largely separated regions in the vector $\hat{\mathbf{g}}$. Computational-wise, this leads to cache misses and degrades performance. This is especially a problem when using multiple threads, where cache misses can be avoided if closely located points are processed on the same CPU core.

A second and much more severe issue comes up within the adjoint resampling step. The direct resampling accesses $\hat{\mathbf{g}}$ only for reading, which can be done concurrently without changing the result of the computation. The adjoint, however, performs in-place additions acting on $\hat{\mathbf{g}}$ and results in race conditions when carried out in parallel. Locking the access to $\hat{\mathbf{g}}$ using a mutex could solve this problem but basically results in a serial execution since the access to $\hat{\mathbf{g}}$ is a large portion of the computation time. In practice, it even degrades the performance, since the CPU pipeline cannot be optimally used because of the unpredictable locking of the access to $\hat{\mathbf{g}}$.

Both of the aforementioned issues have been tackled in NFFT3 [27] and FINUFFT [2]. Ref. [27] proposed node sorting and 1D block-partitioning as independent concepts for the NFFT3 library. Sorting is done with respect to the index \mathbf{l} that is

obtained by flooring $\widetilde{\mathbf{N}} \odot \mathbf{k}_j - m\mathbf{1}$ to $I_{\widetilde{\mathbf{N}}}$. Block-partitioning is performed by splitting $\widehat{\mathbf{g}}$ in its 1D in-memory representation into T regions, where T is the number of threads. Then, all points acting on a specific block are determined and used during the actual computation. This approach achieves close to linear speedup with the number of threads when the density of the sampling points is close to uniform. It, however, yields sub-optimal scaling for non-uniform sampling density since the load is not well-balanced. This issue was solved in [2] where a more general block-partitioning strategy was proposed. The first idea is to use multi-dimensional blocks to improve the data locality. The second idea is to use a fixed number of points per block yielding better load balancing. Decoupling the number of blocks from the number of threads (i.e., use more blocks than threads) ensures that a single long-lasting thread does not slow down the entire computation. We note that a block partitioning scheme to exploit data locality has already been proposed in [12], although no focus has yet been placed on parallel implementation in that work.

In NFFT.jl we followed the approach proposed in [2] with the main difference that our blocks are smaller and all have the same size. For reference, we also implement a regular resampling without block-partitioning, which can be used to investigate the performance gains of block-partitioning. To the best of our knowledge, both NFFT3 and FINUFFT use no block-partitioning in the direct resampling but only in the adjoint resampling. We apply the concept to both operations.

3.7.3. Block-Partitioning Implementation. We next describe block-partitioning formally. The basic idea is to split the domain \mathbb{T}^D into $\mathbf{P} \in 2\mathbb{N}^D$ equally sized blocks

$$(3.4) \quad R_{\mathbf{p}} := \prod_{d=1}^D \left[\frac{p_d}{P_d}, \frac{p_d+1}{P_d} \right), \quad \mathbf{p} \in I_{\mathbf{P}}$$

such that $\bigcup_{\mathbf{p} \in I_{\mathbf{P}}} R_{\mathbf{p}} = \mathbb{T}^D$. Based on this we can collect the indices $j = 1, \dots, J$ of the nodes \mathbf{k}_j in a block $R_{\mathbf{p}}$ by

$$(3.5) \quad \Gamma_{\mathbf{p}} := \{j \in \{1, \dots, J\} : \mathbf{k}_j \in R_{\mathbf{p}}\}.$$

Instead of iterating over $j = 1, \dots, J$ with a single for loop, we can now use two nested for loops:

1. an outer for loop over the blocks $\mathbf{p} \in I_{\mathbf{P}}$.
2. an inner for loop over the node indices $j \in \Gamma_{\mathbf{p}}$.

We next partition the index set $I_{\widetilde{\mathbf{N}}}$ into \mathbf{P} blocks of size $\mathbf{Q} = \lceil \widetilde{\mathbf{N}} \oslash \mathbf{P} \rceil$, where the ceiling is performed in an element-wise fashion. Then, we calculate those indices $\mathbf{l} \in I_{\widetilde{\mathbf{N}}}$ that could be necessary in the resampling for the nodes in the block \mathbf{p} . This is the index set

$$I_{\mathbf{p},m}^{\text{block}} := \bigcup_{\mathbf{k} \in R_{\mathbf{p}}} I_{\widetilde{\mathbf{N}},m}(\mathbf{k}),$$

which has $|\mathbf{Q} + 2m|$ entries.

Having defined the index set $I_{\mathbf{p},m}^{\text{block}}$ we have everything in place to define the block-partitioned resampling. We start with the direct resampling, which is summarized in Algorithm 3.1. The outer loop of the algorithm runs over the blocks (line 1). For each block \mathbf{p} , a local cache $\widehat{\mathbf{q}}_{\mathbf{p}} = (\widehat{\mathbf{q}}_{\mathbf{n},\mathbf{p}})_{\mathbf{n} \in I_{\mathbf{p},m}^{\text{block}}} \in \mathbb{C}^{\mathbf{Q}+2m}$ is created and the data from $\widehat{\mathbf{g}}$ is copied to $\widehat{\mathbf{q}}$ (lines 2–4). Then, the inner loop over all nodes within $\Gamma_{\mathbf{p}}$ is carried out.

This loop initializes \hat{f}_j with zero (line 6) and then performs the inner resampling over $I_{\tilde{N},m}(\mathbf{k}_j)$ where now the local cache $\hat{\mathbf{q}}$ is used instead of $\hat{\mathbf{g}}$ (lines 7–9). Multi-threading the direct block-partitioned resampling is straight forward. The outer for loop over the blocks is run in parallel (indicated by **parfor**) and does not need to take data dependencies into account.

Algorithm 3.1 Block-Partitioned Resampling

```

1: parfor  $p \in I_P$  do                                      $\triangleright$  for each block
2:   for  $l \in I_{p,m}^{\text{block}}$  do
3:      $\hat{\mathbf{q}}_{l,p} \leftarrow \hat{\mathbf{g}}_n$                                 $\triangleright$  init block
4:   end for
5:   for  $j \in \Gamma_p$  do                                        $\triangleright$  for each node in block
6:      $\hat{f}_j \leftarrow 0$ 
7:     for  $l \in I_{\tilde{N},m}(\mathbf{k}_j)$  do                                $\triangleright$  inner resampling
8:        $\hat{f}_j \leftarrow \hat{f}_j + \hat{\mathbf{q}}_{l,p} \tilde{\psi}(\mathbf{k}_j - l \oslash \tilde{N})$ 
9:     end for
10:  end for
11: end parfor

```

The block-partitioned implementation of the adjoint resampling is summarized in [Algorithm 3.2](#). In the first step (lines 1–3), $\hat{\mathbf{g}}$ is initialized with zeros. Then, the main outer loop over the blocks p is initiated. It first initializes the local caches $\hat{\mathbf{q}}$ with zero (lines 5–7). Then, for each node index j in Γ_p all summands acting on the local block $\hat{\mathbf{q}}$ are added at the appropriate location (lines 9–11). Afterwards, the local cache needs to be added to the global vector $\hat{\mathbf{g}}$ (lines 14–16). To multi-thread the block-partitioned adjoint resampling, first the initial loop initializing $\hat{\mathbf{g}}$ with zero is run in parallel. Then, the main outer for loop over the blocks is multi-threaded. As a matter of fact, the data dependency caused by the in-place addition now needs to be taken into account. The first loop in lines 8–12 is unproblematic, since $\hat{\mathbf{q}}$ is thread-local and the inner loop is run in series. However, the second loop (lines 14–16), needs to be locked with a mutex, since otherwise different threads would simultaneously update values in $\hat{\mathbf{g}}$. While this locking could potentially slowdown computation, we note that in practice, the first loop (lines 8–12) has a much higher workload and therefore observed only a small scaling penalty. This is in full agreement with the observations made in [\[2\]](#).

Remark 3.2. In practice, the sets $I_{p,m}^{\text{block}}$ and $I_{\tilde{N},m}(\mathbf{k}_j)$ can be calculated efficiently. Both represent structured blocks within $I_{\tilde{N}}$ that can be represented by just start/end indices in each dimension. In some cases, an index wrap needs to be taken into account. This wrapping can be implemented efficiently by precalculating the wrapping indices for each block. In the actual implementation, all indices are shifted to positive indices with an additional index offset that needs to be taken into account during the calculation. The loop over $I_{\tilde{N},m}(\mathbf{k}_j)$ can be implemented without wrapping since $\hat{\mathbf{q}}$ contains extra padding, i.e., wrapping only needs to be taken into account in the for loop in lines 2–4 of the direct resampling and lines 14–16 of the adjoint resampling.

3.7.4. Window precomputation. Next we discuss different strategies to avoid large computational cost of the window function evaluation. In the literature (e.g., [\[19\]](#) and [\[2\]](#)), we found the following options:

1. **No precomputation**

Algorithm 3.2 Block-Partitioned Adjoint Resampling

```

1: parfor  $l \in I_{\tilde{N}}$  do
2:    $\hat{g}_l \leftarrow 0$  ▷ init output array
3: end parfor
4: parfor  $p \in I_P$  do ▷ for each block
5:   for  $l \in I_{p,m}^{\text{block}}$  do
6:      $\hat{q}_{l,p} \leftarrow 0$  ▷ init block
7:   end for
8:   for  $j \in \Gamma_p$  do ▷ for each node in block
9:     for  $l \in I_{\tilde{N},m}(k_j)$  do ▷ inner resampling
10:       $\hat{q}_{l,p} \leftarrow \hat{q}_{l,p} + \hat{f}_j \tilde{\psi}(k_j - l \oslash \tilde{N})$ 
11:    end for
12:   end for
13:   lock ▷ critical section
14:     for  $l \in I_{p,m}^{\text{block}}$  do
15:        $\hat{g}_l \leftarrow \hat{g}_l + \hat{q}_{l,p}$  ▷ add to output array
16:     end for
17:   end lock
18: end parfor

```

One can directly evaluate the window during the application of the resampling. Depending on the window, this can lead to very slow runtimes if the window evaluation is expensive.

2. Full precomputation

Another common strategy is to compute all window entries prior to the transformation. There are different variants for this, either one can keep the indexing logic the same and just use the cached window entry if needed. Alternatively one can go one step further and represent the entire resampling as a multiplication with a sparse matrix B , see (2.16). This requires storing the indices of the non-zero entries in B .

3. Tensor product based precomputation

The tensor product approach is similar to full computation but calculates the window function in each direction separately. This is basically the same as using the local caches $\tilde{\psi}_{d,j}^{\text{local}}$ with the exception that the caches are stored globally for all nodes k_j . Pulling all window entries into the inner sum, as full precomputation would do, is in many situations slower, since it increases the memory bandwidth, while the local tensor caching approach (with and without precomputation) is CPU bound.

4. Linear interpolation based precomputation

A common pattern to speed up the evaluation of the window function is linear interpolation. To this end, during precomputation a lookup table

$$(3.6) \quad \tilde{\psi}^{\text{linear}} := \left(\tilde{\psi}_s^{\text{linear}} \right)_{s=0}^{S-1} \quad \text{with} \quad \tilde{\psi}_s^{\text{linear}} := \tilde{\psi}_{\text{Base}} \left(\frac{s}{S-2} \right)$$

is created where S is the number sampling points. During resampling one then needs to perform the lookup

$$(3.7) \quad \tilde{\psi}_{\text{Base}} \left(\frac{1}{m} \left(\tilde{N}_d k_{d,j} - l_d \right) \right) \approx \tilde{\psi}_{\tilde{s}}^{\text{linear}} + \alpha (\tilde{\psi}_{\tilde{s}+1}^{\text{linear}} - \tilde{\psi}_{\tilde{s}}^{\text{linear}})$$

where

$$\kappa = \left\lfloor \frac{S-2}{m} \left(\tilde{N}_d k_{d,j} - l_d \right) \right\rfloor, \quad \tilde{s} = \lfloor \kappa \rfloor, \quad \alpha = \kappa - \lfloor \kappa \rfloor.$$

The sampling point $s = S - 1$ laying outside the support of the window function is required since for $|\tilde{N}_d k_{d,j} - l_d| = m$, we have $\tilde{s} = S - 2$ and in turn the lookup table $\tilde{\psi}^{\text{linear}}$ is accessed at index $S - 1$ in (3.7). When choosing $S - 2$ to be a multiple of m (i.e., $S = mY + 2$, $Y \in \mathbb{N}$), it is possible to perform the calculation of $\kappa, \tilde{s}, \alpha$ only once and increase \tilde{s} with a constant offset when iterating l_d through the index set $I_{\tilde{N}_d, m}$. The error of the linear interpolation depends on S , i.e., the larger S , the smaller is the error. NFFT3 provides experimentally determined values for S , which keep the window approximation error for specific m below the entire NFFT error. For instance, for the Kaiser-Bessel window they can be written as

$$S = m2^{c_{\min}(m,8)} + 2, \quad \mathbf{C} = (c_m)_{m=1}^8 = (3 \quad 7 \quad 9 \quad 14 \quad 17 \quad 20 \quad 23 \quad 24).$$

Performance-wise it turns out that this method works well for small $m < 5$ but larger m require very large lookup tables. For instance, $m = 8$ and the Kaiser-Bessel window requires 134217730 entries (1 GB in double precision). This large amount of memory makes the operation much slower for large m since the lookup table does not fit into cache anymore and the resampling basically needs to run over the entire look-up table with a spacing of $\frac{S-2}{m}$ in $2m$ steps.

5. Polynomial approximation based precomputation

An alternative window approximation was proposed in [2]. It uses piecewise polynomial approximation with a high polynomial degree. The idea is to split the window into $2m$ parts and perform polynomial approximation independently on each interval $[-1 + (\ell - 1)/m, -1 + \ell/m]$, $\ell = 1, \dots, 2m$ by setting up the Vandermonde matrix \mathbf{V} and the corresponding sampling vector \mathbf{b} :

$$\mathbf{V} = \left(\left(-\frac{1}{2} + \frac{\vartheta}{\Theta - 1} \right)^z \right)_{\vartheta=0, \dots, \Theta-1; z=0, \dots, Z-1},$$

$$\mathbf{b}_\ell = \left(\tilde{\psi}_{\text{Base}} \left(-1 + \frac{1}{m} \left(\ell - 1 + \frac{\vartheta}{\Theta - 1} \right) \right) \right)_{\vartheta=0}^{\Theta-1}.$$

The polynomial coefficients $\boldsymbol{\mu}_\ell = (\mu_{z,\ell})_{z=0}^{Z-1}$ are then calculated by solving the least-squares problem $\|\mathbf{V}\boldsymbol{\mu}_\ell - \mathbf{b}_\ell\|_2 \xrightarrow{\mu_\ell} \min$. The reason to consider the least-squares problem is that we apply an oversampling and choose more sampling points than the polynomial degree $\Theta = 2Z$ to achieve smaller errors. We thus are in the approximation setting and not the interpolation setting. During resampling the polynomial approximation uses the lookup

$$\tilde{\psi}_{\text{Base}} \left(\frac{1}{m} \left(\tilde{N}_d k_{d,j} - \omega_{\tilde{N}_d, m}(k_{d,j}, \ell) \right) \right) \approx \sum_{z=0}^{Z-1} \mu_{z,\ell} \zeta^z, \quad \ell = 1, \dots, 2m$$

where $\zeta = \tilde{N}_d k_{d,j} - \left\lfloor \tilde{N}_d k_{d,j} \right\rfloor - \frac{1}{2}$. By using $2m$ intervals, the position ζ is the same for all intervals and therefore only the weights $\mu_{z,\ell}$ need to change when iterating through $\ell = 1, \dots, 2m$.

TABLE 1
Memory requirement for different window function precomputation strategies.

Precomputation	Memory [# values]
No	-
Full	$(2m)^D J$
Tensor	$2mDJ$
Linear	$S = mY + 2$
Polynomial	$2mZ$

Our polynomial approximation implementation is slightly different then the one implemented in FINUFFT. In particular, we sample the window function equidistantly on the real line, whereas FINUFFT uses a sampling of a square in the complex space that is tailored to the specific window function used in FINUFFT. In our experiments a polynomial degree of $Z - 1 = 2m$ is sufficient to keep the window function approximation error below the NFFT error, while [2] reported $2m + 3$ to be necessary. This might be due to the different sampling when setting up the approximation problem or the different window function being used. In contrast to FINUFFT, which hardcodes the polynomial coefficients for a specific set of m/σ in a C file generated by a Matlab script, NFFT.jl sets up the coefficients $\mu_{z,\ell}$ up on the fly during precomputation.

Table 1 summarizes the memory requirement for the different precomputation options. The least memory is required for no precomputation followed by polynomial approximation, which effectively always requires less than or equal to 272 entries taking into account that machine precision is reached for $m = 8$. Full and tensor product precomputation are more expensive since they have a per node cost. For $D > 1$ tensor precomputation requires much less memory then full precomputation. Linear interpolation is difficult to classify, since it can require less memory than full/tensor precomputation for small m and large J but also require more memory for large m and small J .

We implemented all precomputation strategies except for no precomputation in NFFT.jl. They can be changed by passing the `precompute` option to the `NFFTPlan` and setting it to one of the enum values `FULL`, `TENSOR`, `LINEAR`, or `POLYNOMIAL`. The reason for not implementing no precomputation is that an efficient implementation of this option with exchangeable window function requires making the window function a type parameter of the `NFFTPlan`. Since direct evaluation is known to be slower than polynomial interpolation even for cheap windows (see the discussion in [2]), we refrained from implementing no precomputation to avoid an increase in software complexity. Tensor product precomputation is combined with polynomial approximation in NFFT.jl, i.e., the precomputation is accelerated with polynomial approximation when using the `TENSOR` option.

Having a look at the two reference libraries, NFFT3 supports the first four precomputation strategies, while FINUFFT supports only polynomial approximation.

3.7.5. Window Size Considerations. We defined the truncated window function $\hat{\psi}$ to have a support of $[-\frac{m}{N_d}, \frac{m}{N_d})$ along dimension d . This means that the sum within the resampling has $(2m)^D$ non-zero summands. When considering the support $[-\frac{m}{N_d}, \frac{m}{N_d}]$, as commonly done in the literature, one needs to consider the special case

that a node \mathbf{k}_j lies exactly on the grid specified by $\mathbf{l} \oslash \widetilde{\mathbf{N}}$ and in turn the sum has $(2m+1)^D$ summands. In principle one could switch between $(2m+1)^D$ and $(2m)^D$ by checking if \mathbf{k}_j lies on the grid. As a matter of fact, this would actually harm the performance, since an `if` statement in a hot loop would need to dynamically switch between $(2m)^D$ and $(2m+1)^D$, which could lead to CPU pipeline flushes.

Our solution mitigates the problem with the potential downside of a slightly higher error. This, however, is very unlikely since \mathbf{k}_j would need to match a grid point exactly in floating point precision. An alternative to our approach is implemented in NFFT3. It considers $(2m+2)^D$ points and uses $\hat{\varphi}$ instead of $\hat{\psi}$ during window evaluation. This has the advantage that no unnecessary multiplications with zero are carried out, which would happen if $\hat{\psi}$ was sampled outside of its support. In turn, this method can lower the approximation error of the NFFT. In [section 5](#), we investigate how much the choice of $(2m+2)^D$ points improves the accuracy and whether it is better to instead use $(2m)^D$ points and increment m , which results in the same number of operations but a different window shape.

4. Materials and Methods. In the remainder of the manuscript an extensive evaluation of NFFT.jl is carried out. Both accuracy and performance are investigated and compared to NFFT3 and FINUFFT.

In all examples $J = |I_N|$ random sampling points $\mathbf{k}_j \in \mathbb{T}^D$ are chosen. This is a typical yet challenging use case and in particular the random nodes can be considered to be the worst-case scenario an NFFT implementation needs to tackle. We use 1D–3D datasets with $N_{1D} = 2^{18} = 262144$, $\mathbf{N}_{2D} = (512, 512)^\top$, $\mathbf{N}_{3D} = (64, 64, 64)^\top$. The oversampling parameter is chosen to be $\sigma = 2$ while the kernel size parameter m is chosen between 3 and 10. For NFFT.jl and NFFT3, the Kaiser-Bessel window (Kaiser-Bessel function for φ and the corresponding Fourier transform for $\hat{\varphi}$) is chosen. FINUFFT uses the *exponential of semicircle* kernel.

All computations were performed on a computer with 1024 GB of main memory and an AMD EPYC 7702 CPU. NFFT.jl was used in version 0.13. FINUFFT (v2.1.0) and NFFT3 (v3.5.2) were applied using the Julia wrapper packages FINUFFT.jl and NFFT3.jl. Julia was used in version 1.8.2. FFTW is used by all libraries and the same option (`FFTW_MEASURE`) was used such that the FFT computation time is the same. This has been verified by inspecting the inner NFFT timings, which can be accessed for all libraries. Benchmarks are performed by repeating the same computation several times (with a 120 seconds threshold) and using the minimum time of all runs. This can be viewed as an estimate for the lower bound that is reached with hot CPU caches and no other workload affecting the computation. The accuracy of the NFFT and its adjoint are determined by calculating the relative error $\frac{\|\mathbf{f}_{\text{NDFT}} - \mathbf{f}_{\text{NFFT}}\|_\infty}{\|\mathbf{f}_{\text{NDFT}}\|_\infty}$ where \mathbf{f}_{NDFT} is the result vector when applying the NDFT and \mathbf{f}_{NFFT} is the result vector of the NFFT.

The code for the entire analysis can be accessed in the `benchmark/paper` directory of the NFFT.jl GitHub repository, which can be found at: <https://github.com/JuliaMath/NFFT.jl>.

5. Results.

5.1. Accuracy Analysis. We start the evaluation by looking at the accuracy of the NFFT.jl for the 2D dataset. The relative error of the direct and adjoint transform is shown for different precomputation strategies of the window function, an oversampling factor of $\sigma = 2$ and $m = 3, \dots, 10$ in [Figure 1](#). First of all one can see that the error decreases exponentially with the kernel parameter m . For the Kaiser-Bessel




FIG. 1. Relative error of NFFT.jl for different precomputation strategies, an oversampling factor of $\sigma = 2$ and window size parameters $m = 3, \dots, 10$. Left shows the error of the direct transform while right shows the error of the adjoint transform.




FIG. 2. Relative error of NFFT.jl compared to NFFT3 and FINUFFT for an oversampling factor of $\sigma = 2$ and window size parameters $m = 3, \dots, 10$. Left shows the error of the direct transform while right shows the error of the adjoint transform.

window the optimum accuracy is reached at $m = 8$ at which point the error saturates. When comparing the different precomputation strategies one can observe almost no difference. This is remarkable since **FULL** does not apply a window approximation, while **TENSOR** (in NFFT.jl), **LINEAR** and **POLYNOMIAL** do. This shows that by a proper selection of the window approximation parameters it is possible to keep this approximation error below the other approximation errors of the NFFT. Accuracy-wise, it thus does not matter which precomputation strategy is used.

We next compare the accuracy between NFFT.jl, NFFT3 and FINUFFT. FINUFFT does not allow to directly set the parameters σ and m but instead uses a tolerance parameter ε to derive σ and m . By reading the source code and enabling debugging information, it was possible to select ε such that, for a fixed value $\sigma = 2$, the desired value m was obtained. Due to an upper limit of $m = 8$, larger m were omitted for FINUFFT. The accuracy comparison is shown in Figure 2. In a first qualitative comparison, one can see that the errors are very similar with NFFT3 being slightly more accurate than NFFT.jl for the direct NFFT, and NFFT.jl being more accurate than FINUFFT for both transformations. When looking at the quantitative values, the mean relative error ratio between NFFT.jl and NFFT3 is 1.767 for the direct transform and 1.106 for the adjoint transform. The mean relative error ratio between FINUFFT and NFFT.jl is 2.722 for the direct transform and 2.242 for the adjoint transform. In all cases, only $m = 3, \dots, 7$ have been used for the calculation of the mean ratio. The differences in accuracy can be explained as follows. FINUFFT uses a different window function that is not as accurate as the Kaiser-Bessel window used in the other two implementations. The difference between NFFT.jl and NFFT3 is caused by NFFT.jl using a window of size $2m$ while NFFT3 uses a window of size $2m + 2$ leading to a slightly lower window truncation error. However, the accuracy

gain by increasing m by one – resulting in a window size of $2(m+1) = 2m+2$ – is much greater than the gain when using $2m+2$ points without changing the shape of the window. This justifies the choice of $2m$ points taken by NFFT.jl and FINUFFT.

5.2. Block-Size Investigation. We next take a closer look at the resampling and investigate the block-partitioning performance within NFFT.jl. Figure 3 shows the runtime performance of the direct and the adjoint NFFT for different block sizes, different dimensionalities (1D–3D), POLYNOMIAL precomputation, and different numbers of threads (1–8). The block sizes are chosen in such a way, that the largest value (1D: 2^{19} , 2D: $(1024, 1024)^T$, 3D: $(128, 128, 128)^T$) corresponds to the usage of one block only. This edge case results in no parallelization and thus leads to sub-optimal performance in the multi-threaded cases. When choosing smaller block sizes also the single-threaded performance is improved for all dimensionalities since this allows for better usage of CPU caches. For all dimensionalities, one can see that the performance degrades when choosing the blocks too small since in that case, the administration overhead for the blocks becomes too high. The optimum block size differs for different numbers of threads and the direct and adjoint transform. It ranges from 10^3 to 10^5 in 1D, from 64 to 128 in 2D, and from 16 to 32 in 3D. Based on that we set the default value to 4096 (1D), $(64, 64)^T$ (2D), and $(16, 16, 16)^T$ (3D) in NFFT.jl. For higher dimensions the block size is set to one.

In addition to the block-partitioning performance (solid), the figure also shows the performance of the regular resampling (dashed). For the adjoint transform the latter is only available for single-threading. One can see that the block-partitioned resampling outperforms the regular resampling in all cases when considering an optimal block size. This shows that block-partitioning should not only be used for the adjoint but also for the direct NFFT.

5.3. Performance Analysis. A performance comparison of NFFT.jl, NFFT3, and FINUFFT for 1D–3D and running on a single thread is shown in Figure 4. In addition to the runtime performance (columns 1–2) also the precomputation time (column 3) is shown. All times are calculated for $m = 3, \dots, 8$ and plotted versus the relative error. We use this representation, since it allows us to make comparisons even if the accuracy of different implementations differs.

The runtime of NFFT.jl is shown for both POLYNOMIAL and TENSOR pre-computation. One can see that the runtime performance is usually a little bit higher for TENSOR with the downside of a larger precomputation time. This trade-off is more prominent for the 1D and 2D transform whereas in 3D the runtimes are very close while the precomputation times are almost negligibly small compared to the transformation time. As a result, we decided to make POLYNOMIAL the default option since it provides good performance and uses less memory.

When comparing the different libraries one can see that the performance varies for different dimensionality and different accuracy. In all cases, FINUFFT and NFFT.jl are faster than NFFT3. In 1D, NFFT.jl with POLYNOMIAL precomputation is as fast as FINUFFT while TENSOR precomputation is fastest. In 2D, TENSOR precomputation is still a little bit faster. In that case, FINUFFT is slightly slower in the direct transform while for the adjoint transform NFFT.jl and FINUFFT are almost on the same level with FINUFFT being slightly faster for high accuracy. The results in 3D are similar, with the only exception that the NFFT.jl with POLYNOMIAL precomputation is at the same level as NFFT.jl with TENSOR precomputation. This is expected because the local window caching discussed in section 3.7.1 exploits the tensor product structure, resulting in a lower time fraction for the window precomputation than for the actual




FIG. 3. Runtime performance of NFFT.jl for block-partitioned (solid) and regular (dashed) resampling, 1D-3D transformations, and 1-8 threads. Left shows the runtime of the direct transformation while right shows the runtime of the adjoint transformation.

summation in the resampling step.

When looking at the precomputation time one can see that FINUFFT is fastest for all dimensionalities. This is most important for 1D transforms where the precomputation time is larger or in the same order as the actual transformation time.

The multi-threading performance of NFFT.jl, NFFT3, and FINUFFT is compared in Figure 5 for the 2D data ($m = 4$ and $\sigma = 2$). The upper two plots show the runtime for $t = 1, \dots, 8$ threads. One can see that all libraries speedup computation by adding more threads although the parallel efficiency (lower two plots) drops with increasing threads. The parallel efficiency is defined as the ratio between actual and theoretically possible speedup. It is in a similar range for all three libraries with a slightly higher value for NFFT3. This might be due to the lock required in NFFT.jl's and FINUFFT's block-partitioning implementation, which is not required in NFFT3's implementation.

6. Discussion. The aim of this work was to introduce the software package NFFT.jl, which is written in pure Julia and combines high-performance with a flexible architecture. The use of Julia is not an implementation detail but, on the contrary, enables us to make less compromises in the implementation, which would neither be

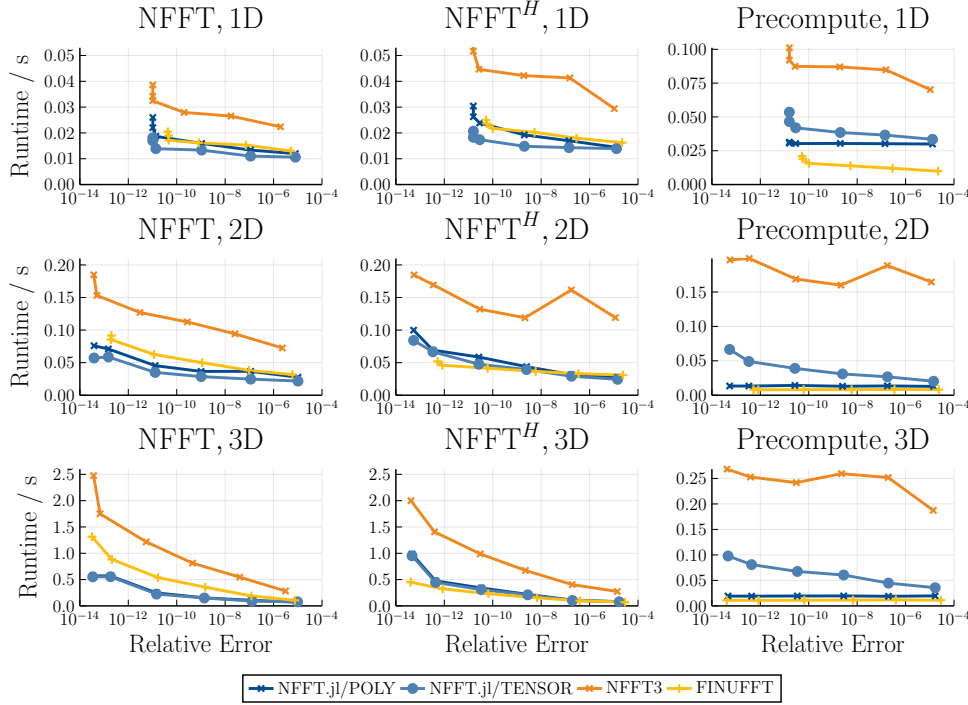


FIG. 4. Single-threading performance of NFFT.jl with the TENSOR and the POLYNOMIAL precomputation strategy compared to NFFT3 (with TENSOR) and FINUFFT.

possible in high-level languages like Matlab and Python, nor in low-level languages like C/C++. Along these lines, our work can also be understood as a showcase of how Julia enables new possibilities in scientific computing that are challenging to achieve with classical tools. The highlight of NFFT.jl is a completely generic implementation that is both number-type and dimension-agnostic.

A comparison of code complexity is not really possible since it also depends on the programmers experience in a certain programming language, whether a code base is found to be complex. Even a comparison of code size is difficult, since the number of characters also depends on the verbosity of keywords used in a given language. Keeping that in mind we counted memory used by the source code of the core NFFT algorithm for all three libraries. It is 184 KB for the `kernel/nfft` directory of NFFT3 and 246 KB for the `src` directory of FINUFFT. The latter, however, contains 106 KB of generated code, which we would not count as source code yielding 140 KB effectively. In contrast, the `src` directory of NFFT.jl has 74 KB of source code. This shows that Julia allows for using less code for achieving the same or even more functionality than its C/C++ counterparts.

The performance of NFFT.jl was found to be the same or better than NFFT3 and FINUFFT. NFFT.jl was faster in both the 1D and the 2D examples and was at the same speed as FINUFFT for the 3D example. Regarding precomputation time, NFFT.jl was only slightly slower than FINUFFT and much faster than NFFT3 for the POLYNOMIAL precomputation. NFFT.jl with TENSOR precomputation was slower than POLYNOMIAL but still much faster than NFFT3, which is due to our polynomial window approximation being applied during precomputation.

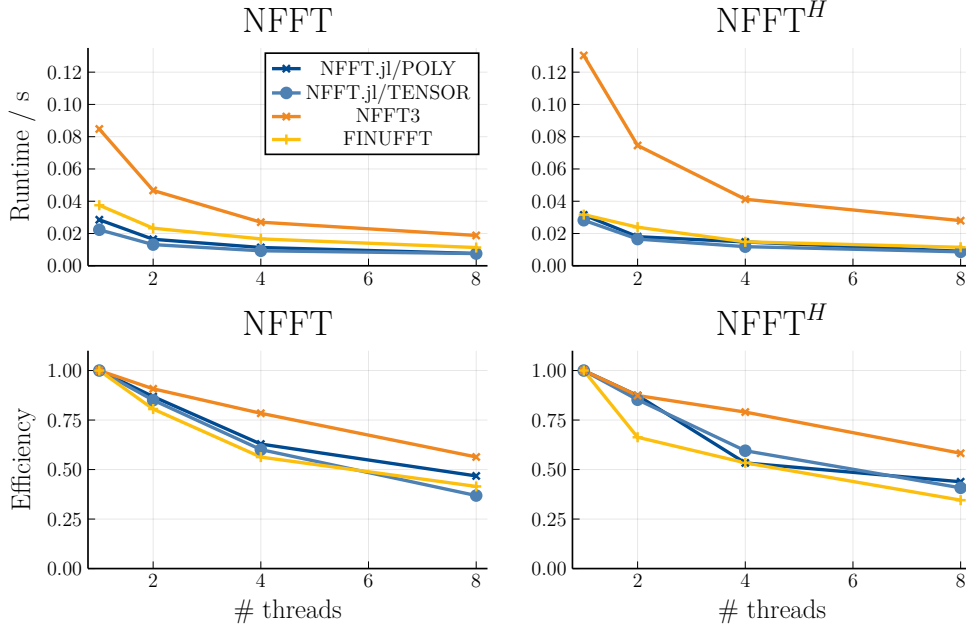


FIG. 5. Multi-threading performance of NFFT.jl with the *TENSOR* and the *POLYNOMIAL* precomputation strategy compared to NFFT3 (with *TENSOR*) and FINUFFT. Shown is the speedup and the parallel efficiency for $t = 1, 2, 4, 8$ threads.

Beside the benchmark between different NFFT software libraries we also made some general findings on NFFT implementation strategies. First of all we compared the use of $2m$ and $2m + 2$ sampling points for the window function and found that the gain in accuracy by the two additional points sampled outside the support of $\hat{\psi}$ are not worth the additional accuracy, since it is more effective to instead increase the window parameter m itself, which increases the width of the window. Moreover, we investigated different block sizes and found dimension-dependent value ranges for which the NFFT reaches highest performance. Since the optimum block-size is highly dependent on the transform size, the number of used threads, and the CPU, we plan to develop an optimization mode, where the optimum block size is automatically chosen based on online-benchmarks, similar to the `FFTW_MEASURE` option in FFTW.

One principle downside of NFFT.jl compared to its C/C++ counterparts is that the latter are currently more binding-friendly than the former. We note, however, that Julia can be embedded in C and therefore it is possible to integrate NFFT.jl in every programming language that allows to call C code. While this still requires a full Julia installation there is also the possibility to statically compile Julia code into a shared library using `PackageCompiler.jl`. We note, however, that static compilation of Julia code is still actively being worked on and will likely evolve in future versions of the language.

An important future step for the NFFT.jl project is to exploit GPU implementations written in Julia. Right now, there is a prototype implementation named `CuNFFT.jl`, which is fully functional and required only 125 lines of code. This is achieved by using the `FULL` precomputation operation and uploading the entire sparse matrix onto the GPU. For optimum performance, one however, would need to write custom kernels and use a similar block-partitioning strategy as we did in the CPU

implementation.

NFFT.jl is already a feature-rich NFFT library that is used in different contexts [23, 11, 16]. During the write-up of this paper, we focused on fixing performance bottlenecks and streamlining the interface of the core algorithm. This does not imply, that NFFT.jl is feature complete yet. One important extension would be the implementation of the NNFFT (also named type-3 NFFT), which has non-equidistant sampling points in both domains. Another missing feature are the fast versions of the non-equidistant cosine and sine transforms (NFCT and NFST). Application oriented tools – like quadrature weights for the non-equidistant sampling points – are also in the scope of NFFT.jl. Those are implemented in the package NFFTTools.jl, which is one layer lower in the package graph and can in principle be also used in combination with other NFFT Julia packages such as NFFT3.jl and FINUFFT.jl.

7. Conclusion. In conclusion, this work has outlined how the scientific programming language Julia can be used to implement a very flexible software package for NFFT computation. Our implementation is completely type- and dimension-agnostic and still uses less code than established packages. We implemented state-of-the-art acceleration techniques taken from two established NFFT libraries and showcased their performance characteristic in 1D–3D examples.

Acknowledgments. We would like to thank all developers and users who have contributed to NFFT.jl. Furthermore, we thank Daniel Potts, Alexander Barnett, Toni Volkmer, Michael Quellmalz and Michael Schmischke for discussions and hints on the use of NFFT3 and FINUFFT. Finally, we thank Jeffrey A. Fessler for fruitful discussions on NFFT terminology.

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] 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.
- [3] G. BEYLKIN, *On the fast Fourier transform of functions with singularities*, Applied and Computational Harmonic Analysis, 2 (1995), pp. 363–381.
- [4] J. BEZANSON, A. EDELMAN, S. KARPINSKI, AND V. B. SHAH, *Julia: A fresh approach to numerical computing*, SIAM review, 59 (2017), pp. 65–98.
- [5] 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.
- [6] A. DUTT AND V. ROKHLIN, *Fast Fourier transforms for nonequispaced data*, SIAM Journal on Scientific computing, 14 (1993), pp. 1368–1393.
- [7] J. A. FESSLER, *On nufft-based gridding for non-cartesian MRI*, Journal of magnetic resonance, 188 (2007), pp. 191–195.
- [8] J. A. FESSLER AND B. P. SUTTON, *Nonuniform fast fourier transforms using min-max interpolation*, IEEE transactions on signal processing, 51 (2003), pp. 560–574.
- [9] M. FRIGO AND S. G. JOHNSON, *The design and implementation of FFTW3*, Proceedings of the IEEE, 93 (2005), pp. 216–231. Special issue on “Program Generation, Optimization, and Platform Adaptation”.
- [10] D. HILLMANN, G. HUTTMANN, AND P. KOCH, *Using nonequispaced fast Fourier transformation to process optical coherence tomography signals*, in European Conference on Biomedical Optics, Optical Society of America, 2009, p. 7372.0R.
- [11] R. D. JACOBSEN, M. NIELSEN, AND M. G. RASMUSSEN, *Generalized sampling in julia*, arXiv preprint arXiv:1607.04091, (2016).
- [12] K. O. JOHNSON AND J. G. PIPE, *Convolution kernel design and efficient algorithm for sampling density correction*, Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine, 61 (2009), pp. 439–447.

- [13] J. KEINER, S. KUNIS, AND D. POTTS, *Using NFFT 3—a software library for various nonequispaced fast fourier transforms*, ACM Transactions on Mathematical Software (TOMS), 36 (2009), pp. 1–30.
- [14] F. KNOLL, A. SCHWARZL, C. DIWOKY, AND D. K. SODICKSON, *gpuNUFFT - an open source GPU library for 3D regridding with direct Matlab interface*, in International Society for Magnetic Resonance in Medicine: Scientific Meeting & Exhibition, 2014, pp. 4297–4297.
- [15] T. KNOPP, *Experimental multi-threading support for the julia programming language*, in Proceedings of the First Workshop for High Performance Technical Computing in Dynamic Languages, IEEE Press, 2014, pp. 1–5.
- [16] T. KNOPP AND M. GROSSER, *MRIReco.jl: An MRI reconstruction framework written in julia*, Magnetic resonance in medicine, 86 (2021), pp. 1633–1646.
- [17] T. KNOPP, S. KUNIS, AND D. POTTS, *A note on the iterative MRI reconstruction from nonuniform k-space data*, International journal of biomedical imaging, 2007 (2007).
- [18] S. KUNIS AND S. KUNIS, *The nonequispaced FFT on graphics processing units*, Pamm, 12 (2012), pp. 7–10.
- [19] S. KUNIS AND D. POTTS, *Time and memory requirements of the nonequispaced fft*, Sampling Theory in Signal and Image Processing, 7 (2008), pp. 77–100.
- [20] D. POTTS, G. STEIDL, AND M. TASCHÉ, *Fast Fourier transforms for nonequispaced data: A tutorial*, Modern sampling theory, (2001), pp. 247–270.
- [21] D. POTTS AND M. TASCHÉ, *Continuous window functions for NFFT*, Advances in Computational Mathematics, 47 (2021), pp. 1–34.
- [22] D. POTTS AND M. TASCHÉ, *Uniform error estimates for nonequispaced fast Fourier transforms*, Sampling Theory, Signal Processing, and Data Analysis, 19 (2021), pp. 1–42.
- [23] A. SCHUTZ, A. FERRARI, D. MARY, É. THIÉBAUT, AND F. SOULEZ, *Large scale 3D image reconstruction in optical interferometry*, in 2015 23rd European Signal Processing Conference (EUSIPCO), IEEE, 2015, pp. 474–478.
- [24] Y.-H. SHIH, G. WRIGHT, J. ANDÉN, J. BLASCHKE, AND A. H. BARNETT, *cuFINUFFT: a load-balanced GPU library for general-purpose nonuniform FFTs*, in 2021 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), IEEE, 2021, pp. 688–697.
- [25] T. S. SØRENSEN, T. SCHAEFFTER, K. Ø. NOE, AND M. S. HANSEN, *Accelerating the nonequispaced fast Fourier transform on commodity graphics hardware*, IEEE Transactions on Medical Imaging, 27 (2008), pp. 538–547.
- [26] G. STEIDL, *A note on fast Fourier transforms for nonequispaced grids*, Advances in computational mathematics, 9 (1998), pp. 337–352.
- [27] T. VOLKMER, *OpenMP parallelization in the NFFT software library*, Preprint, (2012), <https://www-user.tu-chemnitz.de/~potts/paper/openmpNFFT.pdf>.
- [28] A. F. WARE, *Fast approximate Fourier transforms for irregularly spaced data*, SIAM review, 40 (1998), pp. 838–856.
- [29] S.-C. YANG, H.-J. QIAN, AND Z.-Y. LU, *A new theoretical derivation of NFFT and its implementation on GPU*, Applied and Computational Harmonic Analysis, 44 (2018), pp. 273–293.