

Multithreaded Implicitly Dealised Convolutions

Malcolm Roberts

Computer Modelling Group Ltd, 3710 33 Street NW, Calgary, Alberta, T2L 2M1 Canada

John C. Bowman*

*Department of Mathematical and Statistical Sciences, University of Alberta, Edmonton,
Alberta T6G 2G1, Canada*

Abstract

Implicit dealiasing is a method for computing in-place linear convolutions via fast Fourier transforms that decouples work memory from input data. It offers easier memory management and, for long one-dimensional input sequences, greater efficiency than conventional zero-padding. Furthermore, for convolutions of multidimensional data, the segregation of data and work buffers can be exploited to reduce memory usage and execution time significantly. This is accomplished by processing and discarding data as it is generated, allowing work memory to be reused, for greater data locality and performance. A multithreaded implementation of implicit dealiasing that accepts an arbitrary number of input and output vectors and a general multiplication operator is presented, along with an improved one-dimensional Hermitian convolution that avoids the loop dependency inherent in previous work. An alternate data format that can accommodate a Nyquist mode and enhance cache efficiency is also proposed.

Keywords: convolution, implicit dealiasing, fast Fourier transform, multithreading, parallelization, pseudospectral method

*Corresponding author

Email addresses: malcolm.i.w.roberts@gmail.com (Malcolm Roberts),
bowman@ualberta.ca (John C. Bowman)

1. Introduction

The convolution is an important operator in a wide variety of applications ranging from statistics, signal processing, image processing, and the numerical approximation of solutions to nonlinear partial differential equations. The convolution of two sequences $\{F_k\}_{k \in \mathbb{Z}}$ and $\{G_k\}_{k \in \mathbb{Z}}$ is $\sum_{p \in \mathbb{Z}} F_p G_{k-p}$. In practical applications, the inputs $\{F_k\}_{k=0}^{m-1}$ and $\{G_k\}_{k=0}^{m-1}$ are of finite length m , yielding a linear convolution with components $\sum_{p=0}^k F_p G_{k-p}$ for $k = 0, \dots, m-1$. Computing such a convolution directly requires $\mathcal{O}(m^2)$ operations, and roundoff error is a significant problem for large m . It is therefore preferable to make use of the convolution theorem, harnessing the power of the fast Fourier transform (FFT) to map the convolution to a component-wise multiplication. This reduces the computational complexity of a convolution to $\mathcal{O}(m \log m)$ [5, 7] while improving numerical accuracy [8].

Since the FFT considers the inputs to be periodic, the direct application of the convolution theorem results in a circular convolution, due to the indices being computed modulo m . Removing these extra aliases from the periodic convolution to produce a linear convolution is called *dealiasing*.

We give a brief overview of the dealiasing requirements for different types of convolutions in Section 2. The standard method for dealiasing FFT-based convolutions is to pad the inputs with a sufficient number of zero values such that the aliased contributions are all zero, as shown in Figure 1. In Section 3, we generalize the method of implicit dealiasing [3] to handle an arbitrary number of input and output vectors, with a general spatial multiplication operator. This allows implicit dealiasing to be efficiently applied to autocorrelations and pseudospectral simulations of nonlinear partial differential equations (e.g. in hydrodynamics and magnetohydrodynamics). We also discuss key technical improvements that allow implicit dealiasing to be fully multithreaded. For an efficient in-place implementation of the centered Hermitian convolution, it was necessary to unroll the outer loop partially so that interacting wavenumbers can be simultaneously processed. This loop unrolling offers another advantage: it removes the loop interdependence that prevented Function `conv` in [3] from being fully parallelized. For the construction of 2D and 3D convolutions, discussed in Section 4, the advantage of our new 1D convolution routines (relative to those in Ref. [3]) is the better pipelining afforded by loop interdependence, not their parallelizability, as the multithreading is now done at a higher level. The higher dimensional convolutions are decomposed into a sequence of lower-dimensional convo-




Figure 1: Computing a 1D convolution via explicit zero padding.

38 lutions, each of which are run on a separate thread. We demonstrate that
 39 multithreaded implicit dealiasing in dimensions greater than one uses far less
 40 memory and is much faster than explicit dealiasing. The accomplishments
 41 of this work and future directions for research are summarized in Section 5.
 42 Implicitly dealiased convolution routines are publicly available in the open-
 43 source software library FFTW++ [4], which is built on top of the widely used
 44 FFTW library [6].

45 2. Dealiasing requirements for convolutions

46 To compute the standard linear convolution $\sum_{p=0}^k F_p G_{k-p}$ for $k \in \{0, \dots, m-$
 47 $1\}$, the data is padded with m zeroes for a total FFT length of $2m$. We refer
 48 to these inputs as *non-centered* and the paddings as *1/2 padding*. If the input
 49 data is multidimensional with size $m_1 \times \dots \times m_d$, then the data must be zero
 50 padded to $2m_1 \times \dots \times 2m_d$, increasing the buffer size by a factor of 2^d .

51 For pseudospectral simulations, it is convenient to shift the zero wavenum-
 52 ber in the transformed data to the middle of the array. In this case, the inputs
 53 are $\{F_k\}_{k=-m+1}^{m-1}$ and $\{G_k\}_{k=-m+1}^{m-1}$, which we refer to as *centered*, and their
 54 convolution has components $\sum_{p=k-m+1}^{m-1} F_p G_{k-p}$ for $k = -m + 1, \dots, m - 1$.
 55 Convolutions on centered inputs require less padding than on non-centered
 56 inputs: data of length $2m - 1$ needs to be padded only to length $3m - 2$ (nor-
 57 mally extended to $3m$); this is called *2/3 padding* [10]. Explicit zero padding
 58 increases the d -dimensional buffer size in this case by a factor of $(3/2)^d$.

59 A binary convolution can be generalized to an n -ary operator $*(F_1, \dots, F_n)_k =$
 60 $\sum_{p_1, \dots, p_n} F_{p_1} \cdots F_{p_n} \delta_{p_1 + \dots + p_n, k}$, where δ is the Kronecker delta. For non-centered
 61 inputs, an n -ary convolution could be computed as a sequence of binary con-
 62 volutions using *1/2 padding*. However, for centered inputs with both negative
 63 and positive frequencies, each binary convolution would have to be padded

64 further to eliminate all aliased interactions [11]. As a result, n -ary convolu-
 65 tions benefit greatly from implicit dealiasing [3].

66 We consider a generalized convolution operation that takes A inputs and
 67 produces B outputs, where the multiplication performed in the transformed
 68 space can be an arbitrary component-wise operation. In order to make use
 69 of 1/2 padding or 2/3 padding (for noncentered or centered inputs, respect-
 70 ively), the multiplication operator must be quadratic; if the multiplication
 71 operator is of higher degree, one must extend the padding to remove un-
 72 desired aliases. To compute a convolution with A inputs and B outputs using
 73 the convolution theorem, one performs A backward FFTs to transform the
 74 inputs, applies the appropriate multiplication operation on the transformed
 75 data, and then performs B forward FFTs to produce the final outputs, for a
 76 total of $A + B$ FFTs.

77 The choice of multiplication operator determines the type of convolution.
 78 Let $\{f_j\}$ be the inverse Fourier transform of $\{F_k\}$. An autoconvolution can
 79 be computed with just two transforms using $A = B = 1$ and the operation
 80 $f_j \rightarrow f_j^2$, while an autocorrelation would use $f_j \rightarrow f_j \bar{f}_j$, where \bar{f}_j denotes the
 81 complex conjugate of f_j . For the standard binary convolution, there are two
 82 inputs and one output, and the multiplication operation is $(f_j, g_j) \rightarrow f_j g_j$.

83 The nonlinear advective term of the 2D incompressible Navier–Stokes vorticity
 84 equation can be computed with the operation $(u_x, u_y, \partial\omega/\partial x, \partial\omega/\partial y) \rightarrow$
 85 $(u_x \partial\omega/\partial x + u_y \partial\omega/\partial y)$, where $\mathbf{u} = (u_x, u_y)$ is the 2D velocity and $\omega =$
 86 $\hat{\mathbf{z}} \cdot \nabla \times \mathbf{u}$ is the z -component of the vorticity; this requires a total of five FFTs
 87 ($A = 4$ and $B = 1$). As shown in Appendix A, it is possible to reduce the FFT
 88 count for this case to four, with $A = B = 2$. Similarly, in three dimensions,
 89 Basdevant [2] showed that the number of FFT calls can be reduced from
 90 nine to eight, with $A = 3$ and $B = 5$. For incompressible 3D magnetohydro-
 91 dynamic (MHD) flows the operation is $(\mathbf{u}, \boldsymbol{\omega}, \mathbf{B}, \mathbf{j}) \rightarrow (\mathbf{u} \times \boldsymbol{\omega} + \mathbf{j} \times \mathbf{B}, \mathbf{u} \times \mathbf{B})$,
 92 where \mathbf{u} is the velocity, $\boldsymbol{\omega} = \nabla \times \mathbf{u}$ is the vorticity, \mathbf{B} is the magnetic field,
 93 and \mathbf{j} is the current density ($A = 12$, $B = 6$) [12]. However, in Appendix A
 94 we show that Basdevant’s technique can be used to reduce the number of
 95 calls to 14 ($A = 6$, $B = 8$). For the Navier–Stokes and MHD equations, the
 96 operation is quadratic and the convolution is binary ($n = 2$), with a pad-
 97 ding ratio of 2/3 (since the Fourier modes are symmetric about the origin).
 98 In these pseudospectral applications, the physical space quantities are real
 99 valued: one can therefore use complex-to-real Fourier transforms, which are
 100 about twice as efficient as their complex counterparts.

101 **3. One-dimensional implicitly dealiased convolutions**

102 Implicit padding allows one to dealias convolutions without having to
 103 write, read, and multiply by explicit zero values. This is accomplished by
 104 implicitly incorporating the zero values into the top level of a decimated-in-
 105 frequency FFT. The extra memory previously used for padding now appears
 106 as a decoupled work buffer. One-dimensional implicitly dealiased convo-
 107 lutions therefore have the same memory requirements as explicitly padded
 108 convolutions. Although in one dimension implicit padding is only slightly
 109 more efficient than explicit zero padding on a single thread, it still has the
 110 advantage of not requiring the copying of user data to a separate enlarged
 111 zero-padded buffer before performing the FFT. We now describe the optim-
 112 ized 1D building blocks that will be used in Section 4 to construct higher-
 113 dimensional implicitly dealiased convolutions that are much more efficient
 114 and compact than their explicit counterparts.

115 *3.1. Complex convolution*

116 Dealiasing the standard convolution $\sum_{p=0}^k F_p G_{k-p}$ for $k = 0, \dots, m-1$
 117 requires extending the input data with zeros from length m to length $N \geq$
 118 $2m-1$, thus removing the beating of two modes with wavenumber $m-1$ that
 119 would otherwise contaminate mode $N = 0 \bmod N$. One generally chooses
 120 $N = 2m$ to optimize the number of small prime factors in N , resulting in
 121 improved FFT performance.

122 The backward Fourier transform $\{f_j\}_{j=0}^{N-1}$ of the zero-padded input vector
 123 $\{F_k\}_{k=0}^{N-1}$ has components $f_j = \sum_{k=0}^{N-1} \zeta_N^{jk} F_k$, where $\zeta_N = \exp(2\pi i/N)$ denotes
 124 the N^{th} root of unity. The divide-and-conquer strategy of the fast Fourier
 125 transform is based on the property $\zeta_N^r = \zeta_{N/r}$. Since $F_k = 0$ for $k \geq m$,
 126 we can compute the even- and odd-indexed terms of $\{f_j\}_{j=0}^{N-1}$ as separate
 127 subtransforms:

$$f_{2\ell} = \sum_{k=0}^{m-1} \zeta_{2m}^{2\ell k} F_k = \sum_{k=0}^{m-1} \zeta_m^{\ell k} F_k, \quad f_{2\ell+1} = \sum_{k=0}^{m-1} \zeta_{2m}^{(2\ell+1)k} F_k = \sum_{k=0}^{m-1} \zeta_m^{\ell k} \zeta_{2m}^k F_k, \quad (1)$$

128 where $\ell = 0, \dots, m-1$. That is, $\{f_j\}_{j=0}^{N-1}$ can be computed with two Fourier
 129 transforms of length m depending on the input data $\{F_k\}_{k=0}^{m-1}$, with the even-
 130 indexed and odd-indexed parts of the output stored separately. Equation (1)
 131 has a (slightly improved) computational complexity of $\mathcal{O}(N \log m)$, while
 132 avoiding the outermost bit reversal stage and the inconvenience of explicitly




Figure 2: Computing a 1D convolution via implicit dealiasing.

133 appending m extra zero values to the input data. The (scaled) inverse of
 134 Eq. (1) is given by the forward transform

$$\begin{aligned}
 2mF_k &= \sum_{j=0}^{2m-1} \zeta_{2m}^{-kj} f_j = \sum_{\ell=0}^{m-1} \zeta_{2m}^{-k2\ell} f_{2\ell} + \sum_{\ell=0}^{m-1} \zeta_{2m}^{-k(2\ell+1)} f_{2\ell+1} \\
 &= \sum_{\ell=0}^{m-1} \zeta_m^{-k\ell} f_{2\ell} + \zeta_{2m}^{-k} \sum_{\ell=0}^{m-1} \zeta_m^{-k\ell} f_{2\ell+1}, \quad k = 0, \dots, m-1, \quad (2)
 \end{aligned}$$

135 again using two Fourier transforms of length m . Equations (1) and (2) can
 136 be combined to compute a dealiased binary convolution of $\{F_k\}_{k=0}^{m-1}$ and
 137 $\{G_k\}_{k=0}^{m-1}$, as shown in Figure 2 and implemented in pseudocode in Func-
 138 tion `cconv` of Ref. [3]. For each input, two arrays of size m are used instead
 139 of one array of size $2m$. This distinction is the key to the improved effi-
 140 ciency and reduced storage requirements of the higher-dimensional implicit
 141 convolutions described in Section 4. In the 1D complex case, each of the six
 142 complex Fourier subtransforms of size m can be done out of place. Since
 143 implicit dealiasing does not compute the entire inverse Fourier transformed
 144 image at once, we included in our implementation a facility for determining
 145 the spatial coordinates of each point as it is processed. This can be used for
 146 generating an image in x space of the inverse transformed data.

147 As in our previous work [3], we calculate the ζ_N^k factors with a single
 148 complex multiply, using two short pre-computed tables $H_a = \zeta_N^{as}$ and $L_b =$
 149 ζ_N^b , where $k = as + b$ with $s = \lfloor \sqrt{m} \rfloor$, $a = 0, 1, \dots, \lceil m/s \rceil - 1$, and $b =$
 150 $0, 1, \dots, s - 1$. Since these one-dimensional tables occupy only $\mathcal{O}(\sqrt{m})$ com-
 151 plex words, we do not account for them in our storage estimates.

152 Out-of-place FFTs are often more efficient than their in-place counter-
 153 parts, and are more amenable to multithreading. It is not possible to make
 154 use of out-of-place FFTs for explicitly dealiased convolutions without allocat-
 155 ing additional memory, but the situation is different with implicitly dealiased

156 convolutions. For example, in Function `cconv` of Ref. [3], for which $A = 2$,
157 $B = 1$, all FFTs are out of place. The more general Function `cconv` in
158 this work extends implicit dealiasing to arbitrary values of A and B ; it uses
159 $A + B - 1$ out-of-place FFTs and $A + B + 1$ in-place FFTs.

160 In the special case $A > B$, the multiplication operator will free buffers
161 that can be reused, and it is possible to compute the convolution with all
162 FFTs out of place. The idea is that the input and work buffers can be
163 processed separately, and, after applying the multiplication operator, the
164 data in the last of the A work buffers is no longer needed. One can make use of
165 this buffer to perform out-of-place transforms, as shown in Function `cconvA`
166 of Appendix B, for which all $2A + 2B$ FFTs are computed out of place.
167 Likewise, when $A < B$, Function `cconvB` shows that all but one of the
168 $2A + 2B$ FFTs can be performed out of place.

169 When $A = 2$ and $B = 1$, Function `cconvA` runs a few percent faster than
170 Function `cconv` from Ref. [3], thanks to improvements in the loop structure in
171 the pre- and post-processing stages. Thus, in one dimension, as seen in Fig-
172 ure 3(a), implicit dealiasing on a single thread is now on average 12% faster
173 (wall-clock time per convolution is 12% lower) than explicit zero padding.

174 3.1.1. Multithreaded Complex 1D Binary Convolutions

175 We parallelize Functions `cconv`, `cconvA`, and `cconvB` using OpenMP in
176 our pre/post-processing phases and in the multiplication operator, while tak-
177 ing advantage of the multithreading built into the FFTW library. In Fig-
178 ure 3(a), we compare the speed of the implicit and explicit algorithms using
179 one and four threads. Using one thread, the implicit method is on average
180 1.12 times faster than the explicit method, whereas using four threads the
181 performance improvement in wall-clock time is a factor of 1.04 to 2.6 for
182 $m \geq 8192$. The reason that the explicit version benefits from parallelization
183 at smaller m values than implicit dealiasing is a simple consequence of the
184 fact that the vector sizes for explicit dealiasing are twice as large, due to
185 the memory wasted on padding. Multithreading efficiency for small vector
186 lengths is limited due to thread initialization overhead and *false sharing*, a
187 critical performance issue on symmetric multiprocessing systems, where the
188 processors share a local cache. For arrays of $m = 1048576$ double precision
189 complex numbers, each of 16 bytes, the 8MB cache boundary is exceeded
190 and the performance enhancement from multithreading is now limited by the
191 bus bandwidth to off-chip memory. The error bars in the timing figures in-
192 dicate the lower and upper one-sided standard deviations, as given in Ref. [3].

```

Input: vectors  $\{f_a\}_{a=0}^{A-1}$ 
Output: vectors  $\{f_b\}_{b=0}^{B-1}$ 
for  $a = 0$  to  $A - 1$  do
  |  $u_a \leftarrow \text{fft}^{-1}(f_a)$ 
   $\{u_b\}_{b=0}^{B-1} \leftarrow \text{mult}(\{u_a\}_{a=0}^{A-1})$ 
parallel for  $k = 0$  to  $m - 1$  do
  | for  $a = 0$  to  $A - 1$  do
  | |  $f_a[k] \leftarrow \zeta_{2m}^k f_a[k]$ 
for  $a = 0$  to  $A - 1$  do
  |  $f_a \leftarrow \text{fft}^{-1}(f_a)$ 
   $\{f_0\}_{b=0}^{B-1} \leftarrow \text{mult}(\{f_a\}_{a=0}^{A-1})$ 
   $f_0 \leftarrow \text{fft}(f_0)$ 
   $u_0 \leftarrow \text{fft}(u_0)$ 
parallel for  $k = 0$  to  $m - 1$  do
  |  $f_0[k] \leftarrow f_0[k] + \zeta_{2m}^{-k} u_0[k]$ 
for  $b = 1$  to  $B - 1$  do
  |  $f_b \leftarrow \text{fft}(f_b)$ 
  |  $u_b \leftarrow \text{fft}(u_b)$ 
  | parallel for  $k = 0$  to  $m - 1$ 
  | do
  | |  $f_b[k] \leftarrow f_b[k] + \zeta_{2m}^{-k} u_b[k]$ 
return  $\{f_b/(2m)\}_{b=0}^{B-1}$ 

```


Function `cconv` returns the in-place implicitly dealiased 1D convolution of the complex vectors $\{f_a\}_{a=0}^{A-1}$ using the multiplication operator $\text{mult} : \mathbb{C}^A \rightarrow \mathbb{C}^B$. Each of the FFT transforms is multithreaded, with $A+B-1$ out-of-place and $A+B+1$ in-place FFTs.

```


Input: vectors  $\{f_a\}_{a=0}^{A-1}$ 
Output: vectors  $\{f_b\}_{b=0}^{B-1}$ 
for  $a = 0$  to  $A - 1$  do
  |  $\text{pretransform}(f_a, u_{A-1})$ 
  |  $f_a^0 \leftarrow \text{fft}^{-1}(f_a^0)$ 
  |  $f_a^1 \leftarrow \text{fft}^{-1}(f_a^1)$ 
  |  $u_a \leftarrow \text{fft}^{-1}(u_{A-1})$ 
   $\{f_b^0\}_{b=0}^{B-1} \leftarrow \text{mult}(\{f_a^0\}_{a=0}^{A-1})$ 
   $\{f_b^1\}_{b=0}^{B-1} \leftarrow \text{mult}(\{f_a^1\}_{a=0}^{A-1})$ 
   $\{u_b\}_{b=0}^{B-1} \leftarrow \text{mult}(\{u_a\}_{a=0}^{A-1})$ 
for  $b = 0$  to  $B - 1$  do
  |  $u_b \leftarrow \text{fft}(u_b)$ 
  |  $f_b^0 \leftarrow \text{fft}(f_b^0)$ 
  |  $f_b^1 \leftarrow \text{fft}(f_b^1)$ 
  |  $\text{posttransform}(\{f_b^0\}_{k=0}^c \cup$ 
  |  $\{f_b^1\}_{k=2c+2-m}^c, f_b^1[1], u_b)$ 
return  $\{f_b/(3m)\}_{b=0}^{B-1}$ 

```

Function `conv` returns the implicitly dealiased 1D Hermitian convolution of length m ($m+1$) in the compact (noncompact) format, using the multiplication operator $\text{mult} : \mathbb{R}^A \rightarrow \mathbb{R}^B$, with $2A+2B+2$ in-place and $A+B-2$ out-of-place FFTs.



(a)



(b)

Figure 3: In-place 1D complex convolutions of length m , with $A = 2$ and $B = 1$: (a) comparison of computation times for explicit and implicit dealiasing using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency of implicit dealiasing versus number of threads. For efficiency m is chosen to be a power of two.

193 The number of samples varied from several million for small data sizes to 20
 194 for larger data sizes.

195 In Figure 3(b), we observe for $m = 2048$ to 524288 that the implicit
 196 method with four threads has a parallel efficiency of 43% to 85%, where
 197 parallel efficiency is defined as

$$\frac{\text{serial time}}{\text{wall-clock time} \times \text{number of cores}}. \quad (3)$$

198 With four cores, this corresponds to a parallel speedup factor of 1.7 to 3.4.

199 Both the FFTW-3.3.6 library and the convolution layer we built on top
 200 of it were compiled with the GCC 5.3.1 20160406 compiler. Our library was
 201 compiled with the optimizations `-fopenmp -fomit-frame-pointer -fstrict-`
 202 `aliasing -ffast-math -msse2 -mfpmath=sse -march=native` and execut-
 203 ed on a 64-bit 3.4GHz Intel i7-2600K processor with an 8MB cache. Like the
 204 FFTW library, our algorithms were vectorized with specialized SSE2 single-
 205 instruction multiple-data code.

206 3.2. Centered data formats

207 In this work, we extend the treatment of centered Fourier input data
 208 $\{F_{-m+1}, \dots, F_{m-1}\}$ for even m from Ref. [3], to all natural numbers m . We
 209 also implement an optional new data layout $\{F_{-m}, \dots, F_{m-1}\}$. In addition to
 210 handling convolutions of Fourier transformed real-space data of even length,
 211 this extended format can yield significant performance improvements, even if
 212 the additional mode F_{-m} is simply set to zero. We refer to $\{F_{-m+1}, \dots, F_{m-1}\}$
 213 as the *compact* format and $\{F_{-m}, \dots, F_{m-1}\}$ as the *noncompact* format. The
 214 noncompact format is consistent with the output of a real-to-complex FFT
 215 and allows for a multithreaded implementation (Procedure `fft1padBackward`)
 216 since it doesn't have the loop dependency seen in Procedure `fft0padBackward`
 217 in Appendix B.

218 Although the compact format has slightly smaller storage requirements,
 219 on some architectures with more than one (typically a power of two) memory
 220 banks, stride resonances can significantly hurt performance if successive mul-
 221 tidimensional array accesses fall on the same memory bank. A similar effect
 222 can occur on modern architectures due to cache associativity. It is therefore
 223 useful in the centered case to allow the user to choose between the two data
 224 formats.

225 In the compact (noncompact) format, it is convenient to shift the Fourier
 226 origin so that the $k = 0$ mode is indexed as array element $m - 1$ (m). This

227 shift, which can be built into implicitly dealiased convolution algorithms at
 228 no extra cost, allows for more convenient coding of wavenumber loops since
 229 the high-wavenumber cutoff is naturally aligned with the array boundaries.

230 In the compact case, where $N = 2m - 1$, one needs to pad to $N \geq 3m - 2$
 231 to prevent modes with wavenumber $m - 1$ from beating together to contaminate
 232 the mode with wavenumber $-m + 1$. The ratio of the number of physical
 233 to total modes, $(2m - 1)/(3m - 2)$, is then asymptotic to $2/3$ for large m
 234 [10]. With explicit padding, for efficiency reasons one normally chooses the
 235 padded vector length N to be a power of 2, with $m = \lfloor (N + 2)/3 \rfloor$, while for
 236 implicit padding, it is advantageous to choose the subtransform length m to
 237 be a power of 2. Moreover, it is convenient to pad implicitly slightly beyond
 238 $3m - 2$, to $N = 3m$, to support a radix 3 subdivision at the outer level.

239 In the case of an even number $2m - 2$ of spatial data points, one must use
 240 the noncompact data format, with modes running from $-m$ to $m - 1$. We
 241 now describe a noncompact implicitly dealiased centered Fourier transform;
 242 the compact case is obtained by setting $F_{-m} = 0$. Suppose then that $F_k = 0$
 243 for $k > m$. On decomposing $j = (3\ell + r) \bmod N$, where $r \in \{-1, 0, 1\}$, the
 244 substitution $k' = m + k$ allows us to write the backward transform as

$$f_{3\ell+r} = \sum_{k=-m}^{m-1} \zeta_m^{\ell k} \zeta_{3m}^{rk} F_k = \sum_{k'=0}^{m-1} \zeta_m^{\ell k'} \zeta_{3m}^{r(k'-m)} F_{k'-m} + \sum_{k=0}^{m-1} \zeta_m^{\ell k} \zeta_{3m}^{rk} F_k = \sum_{k=0}^{m-1} \zeta_m^{\ell k} w_{k,r}, \quad (4)$$

245 where for $0 \leq k \leq m - 1$,

$$w_{k,r} \doteq \zeta_{3m}^{rk} (F_k + \zeta_3^{-r} F_{k-m}). \quad (5)$$

246 Here \doteq is used to emphasize a definition. The forward transform is then

$$3mF_k = \sum_{r=-1}^1 \zeta_{3m}^{-rk} \sum_{\ell=0}^{m-1} \zeta_m^{-\ell k} f_{3\ell+r}, \quad k = -m + 1, \dots, m - 1. \quad (6)$$

247 The use of the remainder $r = -1$ instead of $r = 2$ allows us to exploit
 248 the optimization $\zeta_{3m}^{-k} = \overline{\zeta_{3m}^k}$ in Eqs. (5) and (6). The number of complex
 249 multiplies needed to evaluate Eq. (5) for $r = \pm 1$ can be reduced by com-
 250 puting the intermediate complex quantities $A_k \doteq \zeta_{3m}^k (\operatorname{Re} F_k + \zeta_3^{-1} \operatorname{Re} F_{k-m})$
 251 and $B_k \doteq i \zeta_{3m}^k (\operatorname{Im} F_k + \zeta_3^{-1} \operatorname{Im} F_{k-m})$, where $\zeta_3^{-1} = (-\frac{1}{2}, -\frac{\sqrt{3}}{2})$, so that for
 252 $k > 0$, $w_{k,1} = A_k + B_k$ and $w_{k,-1} = \overline{A_k - B_k}$. The resulting noncompact
 253 backward transform is given in Procedure `fft1padBackward`; its inverse is

254 given in Procedure `fft1padForward`. The compact versions of these routines
 255 are Procedure `fft0padBackward` in Appendix B and its inverse, Procedure
 256 `fftpad0Forward` from Ref. [3].

Input: vector \mathbf{f}
Output: vector \mathbf{f} , vector \mathbf{u}
parallel for $k = 0$ **to** $m - 1$ **do**

$\mathbf{A} \leftarrow \zeta_{3m}^k \left[\operatorname{Re} \mathbf{f}[m+k] + \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \operatorname{Re} \mathbf{f}[k] \right]$
$\mathbf{B} \leftarrow i\zeta_{3m}^k \left[\operatorname{Im} \mathbf{f}[m+k] + \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \operatorname{Im} \mathbf{f}[k] \right]$
$\mathbf{f}[k] \leftarrow \mathbf{f}[k] + \mathbf{f}[m+k]$
$\mathbf{f}[m+k] \leftarrow \mathbf{A} + \mathbf{B}$
$\mathbf{u}[k] \leftarrow \overline{\mathbf{A} - \mathbf{B}}$

$\mathbf{f}[0, \dots, m-1] \leftarrow \mathbf{fft}^{-1}(\mathbf{f}[0, \dots, m-1])$
 $\mathbf{f}[m, \dots, 2m-1] \leftarrow \mathbf{fft}^{-1}(\mathbf{f}[m, \dots, 2m-1])$
 $\mathbf{u}[0, \dots, m-1] \leftarrow \mathbf{fft}^{-1}(\mathbf{u}[0, \dots, m-1])$

Procedure `fft1padBackward`(\mathbf{f}, \mathbf{u}) stores the shuffled $3m$ -padded centered backward Fourier transform values of a noncompact-format vector \mathbf{f} of length $2m$ in \mathbf{f} and an auxiliary vector \mathbf{u} of length m . The Fourier origin corresponds to array position m .

257 3.2.1. Centered Hermitian Implicitly Padded 1D FFT

258 The Fourier transform of real data satisfies the Hermitian symmetry
 259 $F_{-k} = \overline{F_k}$. This implies that the Fourier coefficient corresponding to $k = 0$ is
 260 real. There is a further consequence of this symmetry when the length N of
 261 the discrete transform $\sum_{j=0}^{N-1} \zeta_N^{-kj} f_j$ is even. Due to the periodicity of the discrete
 262 transform in N , the highest frequency (Nyquist) mode must also be real:
 263 $F_{N/2} = \overline{F_{-N/2}} = \overline{F_{N/2}}$. Letting $m = \lfloor N/2 \rfloor + 1$, in the case where N is even,
 264 the $2m$ modes can therefore be indexed as $\{F_{-m+1}, \dots, F_m\}$ where F_0 and F_m
 265 are real. An odd number $2m - 1$ of modes is indexed as $\{F_{-m+1}, \dots, F_{m-1}\}$.

266 Hermitian symmetry can be used to reduce the computational complexity
 267 and storage requirements of complex-to-real and real-to-complex Fourier
 268 transforms by a factor of about two. A one-dimensional convolution of Hermitian
 269 data only requires the data corresponding to non-negative wavenumbers.
 270 In the compact case, with modes in $\{-m+1, \dots, m-1\}$, the un-
 271 symmetrized physical data needs to be padded with at least $m - 1$ zeros,

```

Input: vector  $f$ , vector  $u$ 
Output: vector  $f$ 
 $f[0, \dots, m-1] \leftarrow \text{fft}(f[0, \dots, m-1])$ 
 $f[m, \dots, 2m-1] \leftarrow \text{fft}(f[m, \dots, 2m-1])$ 
 $u[0, \dots, m-1] \leftarrow \text{fft}(u[0, \dots, m-1])$ 
 $f[m] \leftarrow f[0] + f[m] + u[0]$ 
 $f[0] \leftarrow 0$ 
parallel for  $k = 1$  to  $m-1$  do
     $A \leftarrow f[k]$ 
     $f[k] \leftarrow A + \left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right) \zeta_{3m}^{-k} f[m+k] + \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \zeta_{3m}^k u[k]$ 
     $f[m+k] \leftarrow A + \zeta_{3m}^{-k} f[m+k] + \zeta_{3m}^k u[k]$ 
return  $f/(3m)$ 

```

Procedure `fft1padForward(f,u)` returns the inverse of `fft1padBackward(f,u)`, with $f[0]$ (the Nyquist mode for spatial data of even length) set to zero.

272 just as in Section 3.2. Hermitian symmetry thus necessitates padding the m
 273 non-negative wavenumbers with at least $c \doteq \lfloor m/2 \rfloor$ zeros. The resulting 2/3
 274 padding ratio (for even m) turns out to work particularly well for develop-
 275 ing implicitly dealiased centered Hermitian convolutions. As in the centered
 276 case, we again choose the Fourier size to be $N = 3m$.

277 When $N = 3m$, the most negative (Nyquist) wavenumber $-m$ can con-
 278 structively beat with itself, producing an alias in mode $-m + (-m) = -2m =$
 279 $m \bmod 3m$, which is equivalent to itself modulo $2m$. To remove this alias we
 280 set the Nyquist mode to zero at the end of the convolution, after account-
 281 ing for its effects on the other modes. We note that there are no aliases in
 282 $\{-m, \dots, 2m-1\}$ arising from the interaction of mode $-m$ with any of the
 283 other modes. Those interactions can (and in fact, should) be retained.

284 In the noncompact case, it is sufficient to retain the modes $\{0, \dots, m\}$.
 285 One could of course treat this as compact data of size $m+1$, but in the
 286 frequently occurring case where $m = 2^p$ (for $p \in \mathbb{Z}$), this would require
 287 the computation of subtransforms of length $2^p + 1$ instead of 2^p . The most
 288 efficient available FFT algorithms are typically those of size 2^p .

289 Fortunately, the choice $N = 3m$ also works for the noncompact case,
 290 provided that the entry for the Nyquist mode, which must be real, is zeroed
 291 at the end of the convolution. For example, direct autoconvolution of the

292 Hermitian data $\{(1, 0), (2, 3), (4, 0)\}$ yields $\{(59, 0), (20, -18), (3, 12)\}$. With
 293 $m = 3$, the compact implicitly dealiased convolution in Function `conv` pro-
 294 duces identical results. For $m = 2$, the noncompact version yields the correct
 295 values $\{(59, 0), (20, -18)\}$ for the first m elements only if the data $(3, 12)$ for
 296 the Nyquist mode at $k = m$ is taken into account. That is, in order to enforce
 297 Hermitian symmetry and reality of the corresponding spatial field, we split
 298 the coefficient for $k = -m$ equally between F_{-m} and its Hermitian conjugate
 299 F_m .

300 Let us now describe a noncompact implicitly padded Hermitian FFT; the
 301 compact case can then be obtained by setting $F_{-m} = F_m = 0$. We find

$$f_{3\ell+r} = \sum_{k=-m}^m \zeta_m^{\ell k} \zeta_{3m}^{rk} F_k = \sum_{k'=0}^{m-1} \zeta_m^{\ell k'} \zeta_{3m}^{r(k'-m)} F_{k'-m} + \sum_{k=0}^m \zeta_m^{\ell k} \zeta_{3m}^{rk} F_k = \sum_{k=0}^{m-1} \zeta_m^{\ell k} w_{k,r}, \quad (7)$$

302 Given that $F_k = 0$ for $k > m$, the backward (complex-to-real) transform
 303 appears as Eq. (7), but now with

$$w_{k,r} \doteq \begin{cases} F_0 + F_m \operatorname{Re} \zeta_3^{-r} & \text{if } k = 0, \\ \zeta_{3m}^{rk} (F_k + \zeta_3^{-r} \overline{F_{m-k}}) & \text{if } 1 \leq k \leq m-1. \end{cases} \quad (8)$$

304 We note for $k > 0$ that $w_{k,r}$ obeys the Hermitian symmetry $w_{k,r} = \overline{w_{m-k,r}}$, so
 305 that the Fourier transform $\sum_{k=0}^{m-1} \zeta_m^{\ell k} w_{k,r}$ in Eq. (7) will indeed be real valued.
 306 This allows us to build a backward implicitly dealiased centered Hermitian
 307 transform using three complex-to-real Fourier transforms of the first $c + 1$
 308 components of $w_{k,r}$ (one for each $r \in \{-1, 0, 1\}$). The forward transform is
 309 given by $3mF_k = \sum_{r=-1}^1 \zeta_{3m}^{-rk} \sum_{\ell=0}^{m-1} \zeta_m^{-\ell k} f_{3\ell+r}$ for $k = 0, \dots, m-1$. Since
 310 $f_{3\ell+r}$ is real, a real-to-complex transform can be used to compute the first
 311 $c + 1$ frequencies of $\sum_{\ell=0}^{m-1} \zeta_m^{-\ell k} f_{3\ell+r}$; the remaining $m - c - 1$ frequencies are
 312 then computed using Hermitian symmetry.

313 3.2.2. Multithreaded Hermitian 1D Binary Convolution

314 An in-place implicitly padded Hermitian convolution was previously de-
 315 scribed in Function `conv` of Ref. [3] for the case of $2M$ inputs and one output,
 316 where the multiplication operator was restricted to a dot product. How-
 317 ever, that algorithm cannot be efficiently applied to the autoconvolution
 318 case (with just one input and one output), to pseudospectral simulations of
 319 3D Navier–Stokes and magnetohydrodynamic flows, or to the reduced-FFT
 320 scheme of Basdevant for 2D turbulence (see Appendix A). Furthermore,

321 interloop dependencies at the outermost level prevent it from being multithreaded. Moreover, to facilitate an in-place implementation, the transformed values for $r = 1$ were awkwardly stored in reverse order in the upper half of the input vector, exploiting the quadratic nature of the real-space multiplication operator. By unrolling the outer loop of the in-place Hermitian 1D convolution, these deficiencies can be eliminated, resulting in the fully multithreaded implementation in Function `conv`, generalized to handle A inputs and B outputs and an arbitrary quadratic multiplication operator $\text{mult} : \mathbb{R}^A \rightarrow \mathbb{R}^B$.

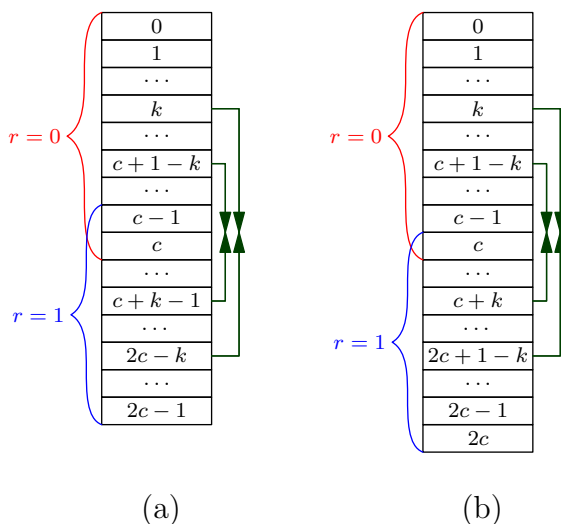


Figure 4: Loop unrolling for the Hermitian 1D convolution when (a) $m = 2c$ and (b) $m = 2c + 1$.

330 To multithread Procedure `pretransform` in Appendix B, we unrolled
 331 two iterations of the loop from Procedure `build` in Ref. [3] to read, process,
 332 and write the entries for the elements indexed by k , $m - k$, $c + 1 - k$, and
 333 $m - c - 1 + k$ simultaneously, for $k = 1, \dots, \lceil c/2 \rceil$, as shown in Figure 4.
 334 The implicitly padded transformed data for remainders $r = 0$ and $r = 1$
 335 is stored in the input data vector f , whereas the data for remainder $r = -1$
 336 is stored in the auxiliary vector u . In the frequently encountered case where
 337 $m = 2c$, the values at position $c - 1$ and c overlap for the remainders $r = 0$
 338 and $r = 1$, whereas when $m = 2c + 1$ there is only one overlapping value,
 339 at index c . In the noncompact case, any Nyquist inputs $f[m]$ and $g[m]$ are
 340 properly accounted for, but on output $f[m]$ is set to zero, for consistency

341 with Hermitian symmetry. A similar loop unrolling is used in a revised
342 implementation of the post-processing phase (see Procedure `posttransform`
343 in Appendix B) to allow for an arbitrary number of inputs A and outputs B .

344 In the pseudocode, the portions of the arrays corresponding to remainders
345 $r = 0$ and $r = 1$ are shown in Figure 4 and distinguished by the superscripts
346 0 and 1. Since these portions overlap when written to the array f , additional
347 code is required to save and restore the overlapping elements in the actual
348 in-place implementation.

349 In our general Function `conv`, only $A + B - 2$ of the $3A + 3B$ FFTs
350 can be performed out of place. When $A = B$ this is the best that one
351 can do: for example, for an autoconvolution ($A = B = 1$), there is no free
352 buffer available that would enable the use of out-of-place transforms. Nev-
353 ertheless when $A > B$ or $B > A$ the optimized Function `convA` or `convB`,
354 respectively, in Appendix B performs one FFT in place and the remaining
355 $3A + 3B - 1$ FFTs out of place. Even with one thread, for $A = 2$ and
356 $B = 1$, Functions `conv` and `convA` both have slightly better performance
357 than Function `conv` in Ref. [3], primarily due to the removal of loop interde-
358 pendence. Most importantly, `conv`, `convA`, and `convB` have fully parallelized
359 pre- and post-processing phases and use FFTW’s built-in parallel FFTs, which
360 are typically much more efficient when the transforms are out of place. The
361 new routines accept either compact or noncompact inputs and can therefore
362 also benefit from the performance advantage of the noncompact data format
363 discussed in Section 3.2.

364 In the case of a binary convolution with two input vectors and one output
365 vector, a fully in-place convolution requires a total of nine Hermitian Four-
366 ier transforms of size m , for an overall computational scaling of $\frac{9}{2}Km \log_2 m$
367 operations, where $K = 34/9$ [3], in agreement with the leading-order scaling
368 of an explicitly padded centered Hermitian convolution. In our new imple-
369 mentation, eight of the nine Fourier transforms can now be performed out
370 of place, using the same amount of memory ($6c + 2$ words in the compact
371 case) as required to compute a centered Hermitian convolution with explicit
372 padding.

373 As seen in Fig. 5, the efficiency of the resulting implicitly dealiased
374 centered Hermitian convolution is comparable to an explicit implementa-
375 tion. For each algorithm, we benchmark only those vector lengths that yield
376 optimal performance. The optimal values of m for the explicit version are
377 $\lfloor (2^p + 2)/3 \rfloor$ for natural numbers p , whereas for the implicit version the op-
378 timal values are powers of two, so direct comparison of the methods using

379 optimal problem sizes is not possible. Instead, we compare the two methods
 380 using a linear interpolation (with respect to $\log m$) of the execution time res-
 381 caled by the computational complexity of the algorithm. With one thread,
 382 the implicit version runs between 5% and 23% faster for $m \geq 2048$; with four
 383 threads, the implicit version is between 12% and 105% faster for $m \geq 65536$,
 384 as shown in Fig. 5(a). We demonstrate the parallel efficiency of the implicit
 385 routine in Fig. 5(b) using one, two, and four threads. For $m \geq 8192$, the
 386 parallel efficiency of the implicit method with four threads is between 45%
 387 and 68%, giving a speedup of a factor of 1.8 to 2.7.

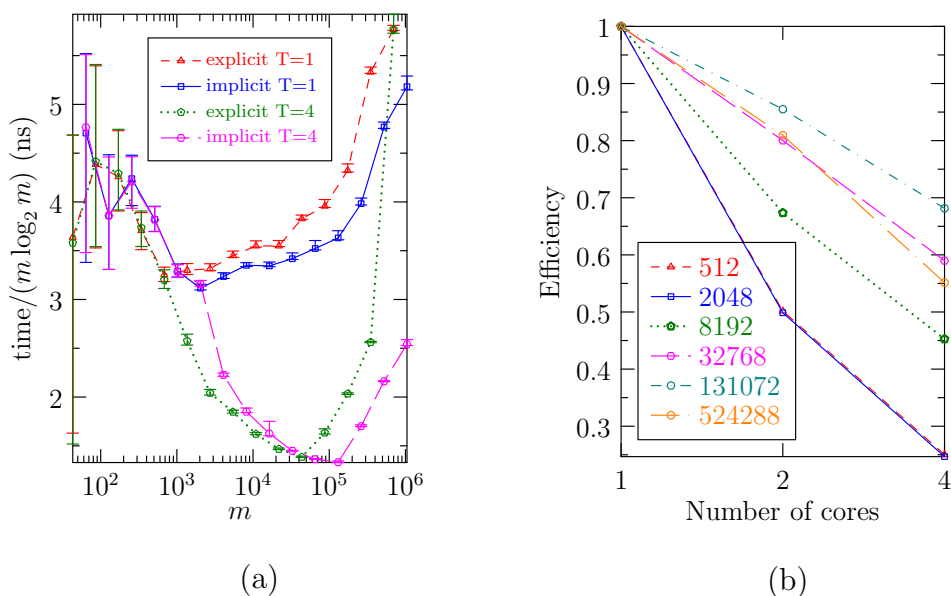


Figure 5: In-place 1D Hermitian convolutions of length m : (a) comparison of computation times for explicit and implicit dealiasing using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency of implicit dealiasing versus number of threads. For implicit convolutions, m is chosen to be a power of two, while for explicit convolutions $m = \lfloor (N + 2)/3 \rfloor$, where N is a power of two.

388 4. Higher-dimensional convolutions

389 A d -dimensional convolution can be computed by performing an inverse
 390 FFT of size $m_1 \times \dots \times m_d$, applying the appropriate multiplication on the
 391 transformed data, followed by an FFT back to the original space. Equival-
 392 ently, one can perform $\prod_{i=2}^d m_i$ inverse FFTs in the first dimension, followed

393 by m_1 convolutions of dimension $d - 1$, and finally $\prod_{i=2}^d m_i$ FFTs in the first
 394 dimension. The innermost operation of a recursive multidimensional convo-
 395 lution thus reduces to a 1D convolution. Using this decomposition, one can
 396 reuse the work buffer for each implicitly dealiased subconvolution, thereby re-
 397 ducing the total memory demand relative to the explicit d -dimensional deali-
 398 asing requirement. For multithreaded implicitly dealiased convolutions, the
 399 initial inverse FFT can be parallelized by dividing the $\prod_{i=2}^d m_i$ 1D FFTs and
 400 m_1 subconvolutions between the T threads. Since each implicitly dealiased
 401 subconvolution requires a work buffer, the total memory requirement grows
 402 with the number of threads, but is still much lower than that required for
 403 explicit multidimensional convolutions when $T \ll m_1$. When $T \leq m_1$, we
 404 compute T subconvolutions at a time, using one inner thread per subconvo-
 405 lution to avoid over-subscription. Otherwise, if $T > m_1$, we parallelize only
 406 the inner subconvolution (over all T threads).

407 A single-threaded 2D implicitly 1/2-padded complex convolution is shown
 408 in Figure 6. Each input buffer is implicitly padded and inverse Fourier trans-
 409 formed in the x direction to produce the data shown in the square boxes. An
 410 implicitly padded inverse FFT is then performed in the y direction, column-
 411 by-column, using a one-dimensional work buffer, to produce a single column
 412 of the Fourier transformed image, depicted in yellow. The Fourier trans-
 413 formed columns of two inputs F and G are then multiplied pointwise and
 414 stored back into the F column. At this point, the y forward-padded FFT can
 415 then be performed, with the result stored in the lower-half of the column,
 416 next to the previously processed data shown in red. The process is repeated
 417 on the remaining columns, shifting and reusing the work buffer. Once all
 418 the columns have been processed, a forward-padded FFT in the x direction
 419 produces the final convolution in the left-hand half of the F buffer.

420 The reuse of subconvolution work memory allows the convolution to be
 421 computed using less total memory: for 1/2 padded convolutions, the memory
 422 requirement per input is only about twice what would be required if deali-
 423 asing was outright ignored. This represents a memory savings of a factor
 424 of 2^{d-1} as compared to explicit padding; for 2/3 padded convolutions, the
 425 memory savings factor is $(3/2)^{d-1}$. In addition to having reduced memory
 426 requirements, implicitly dealiased multidimensional convolutions are signific-
 427 antly faster than their explicit counterparts, due to better data locality and
 428 cache management, along with the fact that transforms of data known to be
 429 zero are automatically avoided.

430 In the following subsections, we show that the algorithms developed in




Figure 6: The reuse of memory in a 2D complex implicitly dealiased convolution: after applying a 1D y convolution to the yellow column, the upper half is reused for the next column.

431 Section 3 can be used as building blocks to construct efficient implicitly
 432 padded higher-dimensional convolutions.

433 4.1. Complex 2D convolution

434 Pseudocode for the implicitly padded transforms described by Eqs. (1)–
 435 (2) was given in Ref. [3] as Procedures `fftpadBackward` and `fftpadForward`.
 436 In order to compute a 2D convolution in parallel, the loops in these proced-
 437 ures were parallelized, and parallel FFTs were used. Since the input and
 438 output of these routines are multidimensional and the required FFT is one-
 439 dimensional, we use the FFTW multiple 1D FFT routine. A multithreaded
 440 version of this routine is available in the FFTW library, but we found that
 441 its parallel performance was sometimes lacking. This was somewhat surpris-
 442 ing, as there exists a simple algorithm to parallelize such problems: if one
 443 wishes to perform M FFTs using T threads, one can simply divide the M
 444 FFTs among the T threads, with any remaining r FFTs distributed among
 445 the first r threads. At run time, we automatically test for the possibility
 446 that this decomposition is faster than FFTW’s parallel multiple FFT and use
 447 whichever algorithm runs faster. This yielded a significant improvement in
 448 the parallel performance of our convolutions.

449 As shown in Fig. 7(a), the resulting implicit 2D algorithm dramatically
 450 outperforms the explicit version: using one thread, the mean speedup is
 451 a factor of 1.5, with a maximum speedup of 1.8. Using four threads, the
 452 mean speedup over the parallel explicit version is approximately 2.6, with
 453 a maximum speedup factor of 4.5. Fig. 7(b) shows the parallel efficiency

454 of the 2D implicitly dealiased complex convolution for a variety of problem
 455 sizes. The parallel efficiency for the implicit routine ranges from 58% to 92%
 456 with four threads, for a speedup of 2.3 to 3.7 relative to one thread. The
 457 explicit routine has a parallel efficiency between 25% and 90%. Notably, the
 458 2D explicit version has poor parallel performance for problem sizes of 512^2
 459 and above using FFTW’s built-in multithreading.

460 Because the same temporary array u is used for each column of the convo-
 461 lution, the memory requirement is $2Cm_xm_y + TCm_y$ complex words using T
 462 threads, where $C = \max\{A, B\}$. Assuming that $T < 2m_x$, this is far less than
 463 the $4Cm_xm_y$ complex words needed for an explicitly padded convolution.

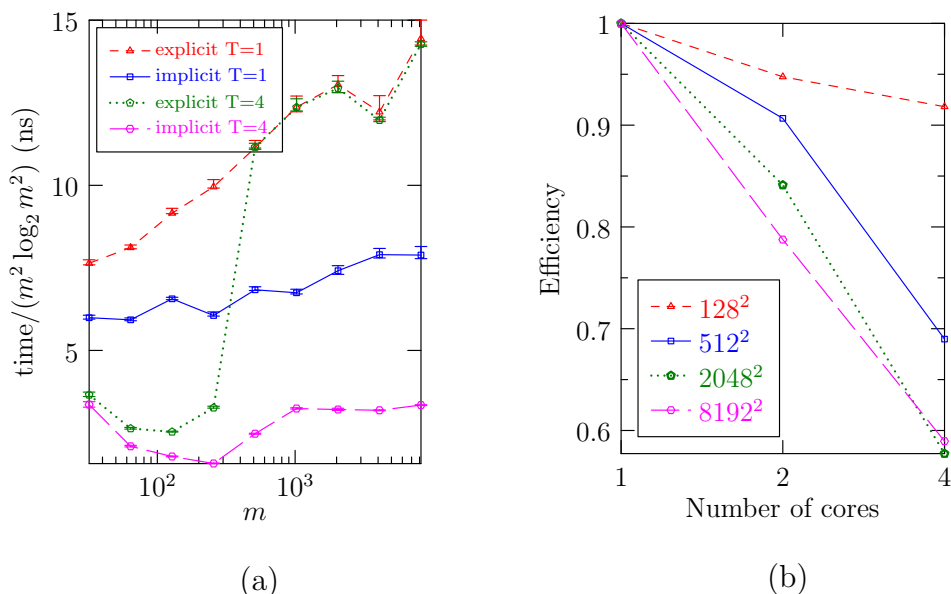


Figure 7: In-place 2D complex convolutions of size $m \times m$: (a) comparison of computation times for explicit and implicit dealiasing using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency of implicit dealiasing versus number of threads. Here m is chosen to be a power of two.

464 4.2. Centered Hermitian 2D convolution

465 In two dimensions, the Fourier-centered Hermitian symmetry appears as
 466 $F_{-k,-\ell} = \overline{F_{k,\ell}}$. This symmetry is exploited in the centered Hermitian convo-
 467 lution algorithm shown for the noncompact case in Function `conv2`. As with
 468 the 1D Hermitian convolution, one has the option to use a compact or non-
 469 compact data format. For the compact data format, the array has dimensions

470 $\{-m_x + 1, \dots, m_x - 1\} \times \{0, \dots, m_y - 1\}$, whereas the noncompact version
471 has dimensions $\{-m_x, \dots, m_x - 1\} \times \{0, \dots, m_y\}$. One can also perform con-
472 volutions on data that is compact in one direction and noncompact in the
473 other. For serial computations, the best performance typically is achieved
474 when the x direction is compact and the y direction is noncompact, so that
475 each dimension is odd, to reduce cache associativity issues. However, when
476 running on more than one thread, the noncompact format should be used
477 in the x direction since Procedures `fft1padBackward` and `fft1padForward` are
478 fully multithreaded.

479 While the noncompact case requires slightly more memory than the com-
480 pact case, one advantage of the noncompact version is that the output of
481 the Fourier transform of $(2m_x - 2) \times (2m_y - 2)$ real values corresponds to
482 the modes $\{-m_x, \dots, m_x - 1\} \times \{0, \dots, m_y\}$, where the Fourier origin has
483 been shifted in the x direction to the middle of the array. Moreover, one
484 is able to use the extra memory in the x direction for temporary storage,
485 and having $m_y + 1$ variables in the y direction avoids latency issues with
486 cache associativity when m_y is a power of two. These factors combine to
487 give the noncompact format a performance advantage over the compact one:
488 the noncompact case is typically slightly faster than the compact case when
489 using one thread and 25% faster on average when using four threads.

```

Input: matrix  $\{f_a\}_{a=0}^{A-1}$ 
Output: matrix  $\{f_b\}_{b=0}^{B-1}$ 
for  $a = 0$  to  $A - 1$  do
  parallel for  $j = 0$  to  $m_y - 1$  do
    | fft1padBackward( $f_a^T[j], U_a^T[j]$ )
  parallel for  $i = 0$  to  $m_x - 1$  do
    | cconv( $\{f_a[i]\}_{a=0}^{A-1}$ )
    | cconv( $\{U_a[i]\}_{a=0}^{A-1}$ )
  for  $b = 0$  to  $B - 1$  do
    parallel for  $j = 0$  to  $m_y - 1$  do
      | fft1padForward( $f_b^T[j], U_b^T[j]$ )
  return  $f$ 

```

490

Function `cconv2` returns the in-place implicitly dealiased convolution of $m_x \times m_y$ matrices $\{f_a\}_{a=0}^{A-1}$ in $\{f_b\}_{b=0}^{B-1}$, using A temporary $m_x \times m_y$ matrices $\{U_a\}_{a=0}^{A-1}$.

```

Input: matrix  $\{f_a\}_{a=0}^{A-1}$ 
Output: matrix  $\{f_b\}_{b=0}^{B-1}$ 
for  $a = 0$  to  $A - 1$  do
  parallel for  $j = 0$  to  $m_y$  do
    | fft1padBackward( $f_a^T[j], U_a^T[j]$ )
  parallel for  $i = 0$  to  $2m_x - 1$  do
    | conv( $\{f_a[i]\}_{a=0}^{A-1}$ )
  parallel for  $i = 0$  to  $m_x - 1$  do
    | conv( $\{U_a[i]\}_{a=0}^{A-1}$ )
  for  $b = 0$  to  $B - 1$  do
    parallel for  $j = 0$  to  $m_y$  do
      | fft1padForward( $f_b^T[j], U_b^T[j]$ )
  return  $f$ 

```

Function `conv2` returns the in-place implicitly dealiased centered Hermitian convolution of $2m_x \times (m_y + 1)$ matrices $\{f_a\}_{a=0}^{A-1}$ in the noncompact data format, using A temporary $m_x \times (m_y + 1)$ matrices $\{U_a\}_{a=0}^{A-1}$.

491 The explicit version requires storage for $9Cm_x(m_y + 1)/2$ complex words,
 492 where $C = \max\{A, B\}$. For the noncompact case, the implicit version using
 493 T threads requires storage for $3Cm_x(m_y + 1) + TC(\lfloor m_y/2 \rfloor + 1)$ complex
 494 words, which is much less than the explicit case when $m_x \geq T$. As shown
 495 in Fig. 8(a), implicit padding again yields a dramatic improvement in speed:
 496 the implicit version is on average 1.36 times faster than the explicit version
 497 when using one thread, and 2.93 times faster than the explicit version when
 498 using four threads. In Fig. 8(b) we show that the parallel efficiency of the
 499 implicit version is between 73% and 87% efficiency when using four threads,
 500 giving a speedup of a factor of 2.9 to 3.5. As in the 2D complex case, the
 501 explicit version does not parallelize well for large problem sizes.

502 4.3. Complex 3D convolution

503 The decoupling of the 2D work arrays in Function `cconv2` facilitates the
 504 construction of an efficient 3D implicit complex convolution, as described
 505 in Function `cconv3`. For A inputs with dimensions $m_x \times m_y \times m_z$ and
 506 B outputs, the explicit version requires $8Cm_xm_y m_z$ complex words, where

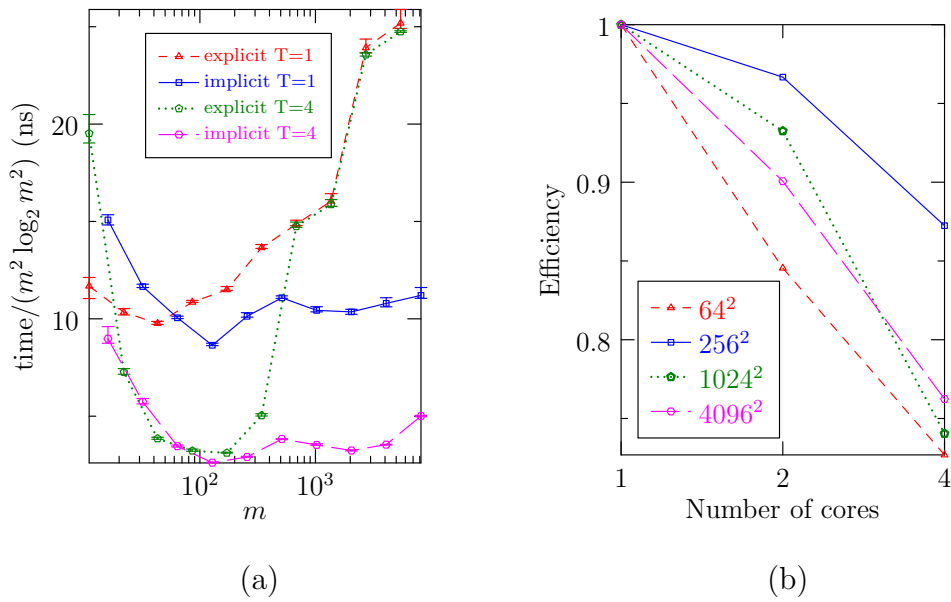


Figure 8: In-place 2D Hermitian convolutions of size $2m \times (m + 1)$: (a) comparison of computation times for explicit and implicit dealiasing using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency of implicit dealiasing versus number of threads. For implicit convolutions, m is chosen to be a power of two, while for explicit convolutions $m = \lfloor (N + 2)/3 \rfloor$, where N is a power of two.

507 $C = \max\{A, B\}$. In contrast, the implicit version with T threads requires
 508 $2Cm_xm_y m_z + TCm_y m_z + TCm_z$ complex words, approximately one quarter
 509 the storage requirements for the explicit version when $m_x \gg T$. As shown
 510 in Fig. 9(a), implicitly dealiased convolutions are consequently much faster
 511 than their explicit counterparts. For a single thread, the implicit version
 512 is on average 1.9 times as fast as the explicit version and 4.3 times faster
 513 on average when comparing execution times over four threads. Fig. 9(b)
 514 shows the parallel efficiency of the implicit version, which is between 65%
 515 and 86% efficient when using four threads, giving a speedup of a factor of 2.6
 516 to 3.4 over one thread. The explicit version has reasonable parallel efficiency
 517 for small problem sizes, but this drops to roughly 25% on four threads for
 518 problem size $m \geq 64$.

519 4.4. Centered Hermitian 3D convolution

520 As with the 1D and 2D cases, we offer compact and noncompact ver-
 521 sions of a 3D Hermitian convolution, and users can choose formats that are
 522 compact/noncompact in each direction separately. For serial computations,
 523 the best performance typically is achieved when the x and y directions are
 524 compact and the z direction is noncompact, so that each dimension is odd,
 525 in the interest of cache associativity. However, just as for 2D Hermitian con-
 526 volutions, when running on more than one thread, the x direction should be
 527 made noncompact to obtain optimal multithreading efficiency.

```

Input:  $\{f_a\}_{a=0}^{A-1}$ 
Output:  $\{f_b\}_{b=0}^{B-1}$ 
 $R \leftarrow$ 
 $\{0, \dots, m_y - 1\} \times \{0, \dots, m_z - 1\}$ 
for  $a = 0$  to  $A - 1$  do
  parallel foreach  $(j, k) \in R$  do
    | fftpadBackward( $f_a^T[k][j], U_a^T[k][j]$ )
  parallel for  $i = 0$  to  $m_x - 1$  do
    | cconv2( $\{f_a[i]\}_{a=0}^{A-1}$ )
    | cconv2( $\{U_a[i]\}_{a=0}^{A-1}$ )
  for  $b = 0$  to  $B - 1$  do
    | parallel foreach  $(j, k) \in R$  do
      | fftpadForward( $f_b^T[k][j], U_b^T[k][j]$ )
  return f

```

Function `cconv3` returns the in-place implicitly dealiased complex convolution of $m_x \times m_y \times m_z$ matrices $\{f_a\}_{a=0}^{A-1}$, using A temporary $m_x \times m_y \times m_z$ matrices $\{U_a\}_{a=0}^{A-1}$.

Pseudocode for the noncompact algorithm is given in Function `conv3`. The noncompact version again offers a performance advantage over the compact version, with the single-threaded compact and noncompact cases roughly equal in execution time on a single thread, and the noncompact case offering between a 1% and 10% performance advantage when parallelized over four threads.

In the noncompact format, the memory requirements for an explicit 3D Hermitian convolution with A inputs and B outputs is $\frac{27}{2}Cm_xm_y(m_z + 1)$ complex words, whereas the implicit version requires only $6Cm_xm_y(m_z + 1) + TCm_y(m_z + 1) + TC(\lfloor m_z/2 \rfloor + 1)$ complex words using T threads, where $C = \max\{A, B\}$. We did not implement a high-performance version of the explicit routine, so instead we show the execution time of the implicit routine using one and four threads in Fig. 10 (a). The parallel efficiency is shown in Fig. 10(b) and ranges between 65% and 92%, which translates to a speedup of a factor of 2.6 to 3.7 using four threads instead of one.

```

Input:  $\{f_a\}_{a=0}^{A-1}$ 
Output:  $\{f_b\}_{b=0}^{B-1}$ 
 $R \leftarrow \{0, \dots, 2m_y\} \times \{0, \dots, m_z\}$ 
for  $a = 0$  to  $A - 1$  do
  parallel foreach  $(j, k) \in R$  do
    | fft1padBackward( $f_a^T[k][j], U_a^T[k][j]$ )
  parallel for  $i = 0$  to  $2m_x - 1$  do
    | conv2( $\{f_a[i]\}_{a=0}^{A-1}$ )
  parallel for  $i = 0$  to  $m_x - 1$  do
    | conv2( $\{U_a[i]\}_{a=0}^{A-1}$ )
  for  $b = 0$  to  $B - 1$  do
    | parallel foreach  $(j, k) \in R$  do
      | fft1padForward( $f_b^T[k][j], U_b^T[k][j]$ )
  return f

```

Function `conv3` returns the in-place implicitly dealiased Hermitian convolution of $2m_x \times 2m_y \times (m_z + 1)$ matrices $\{f_a\}_{a=0}^{A-1}$, using A temporary $m_x \times m_y \times (m_z + 1)$ matrices $\{U_a\}_{a=0}^{A-1}$.




Figure 9: In-place 3D complex convolutions of size $m \times m \times m$: (a) comparison of computation times for explicit and implicit dealiasing using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency of implicit dealiasing versus number of threads. Here m is chosen to be a power of two.




Figure 10: Implicitly dealiased in-place 3D Hermitian convolutions of size $(2m-1) \times (2m-1) \times (m+1)$ for $T = 1$ and $2m \times (2m-1) \times (m+1)$ for $T = 4$: (a) computation times using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency versus number of threads. Here m is chosen to be a power of two.

544 5. Concluding remarks

545 In this work we developed an efficient method for computing implicitly
 546 dealiased convolutions parallelized over multiple threads. Methods were de-
 547 veloped for noncentered complex data and centered Hermitian-symmetric
 548 data with inputs in one, two, and three dimensions. We showed how more
 549 general multiplication operators can be supported, allowing for the efficient
 550 computation of autoconvolutions, correlations (which are identical to convo-
 551 lutions for Hermitian-symmetric data), and general nonlinearities in pseudo-
 552 spectral simulations.

553 Implicitly dealiased convolutions require less memory, are faster, and have
 554 greater parallel efficiency than their explicitly dealiased counterparts. Spe-
 555 cifically, in d dimensions the memory savings for $1/2$ padding is a factor
 556 of 2^{d-1} ; for $2/3$ padding the savings factor is $(3/2)^{d-1}$. The decoupling of
 557 temporary storage and user data means that even in one dimension, users
 558 can save memory by not having to copy their data to a separate buffer. In
 559 higher dimensions, this decoupling allows one to reuse work memory. By
 560 avoiding the need to compute the entire Fourier image at once, one obtains

561 a dramatic reduction in total memory use. Moreover, the resulting increased
 562 data locality significantly enhances performance, particularly under parallel-
 563 ization. For example, a 3D implicitly dealiased complex convolution runs
 564 about twice as fast as an explicitly dealiased convolution on one thread, and
 565 over four times faster than the explicit method when both are parallelized
 566 over four threads. For large problem sizes, an implicit complex convolution
 567 requires one-half of the memory needed for a zero-padded convolution in two
 568 dimensions and one-quarter in three dimensions. In the centered Hermitian
 569 case, the memory use in two dimensions is 2/3 of the amount used for an ex-
 570 plicit convolution and 4/9 of the corresponding storage requirement in three
 571 dimensions.

572 An upcoming paper will discuss the implementation of implicit deali-
 573 asing on distributed-memory architectures, using hybrid MPI/OpenMP. Im-
 574 plicit dealiasing of higher-dimensional convolutions over distributed memory
 575 benefits significantly from the reduction of communication costs associated
 576 with the smaller memory footprint. It also provides a natural way of overlap-
 577 ping communication (during the transpose phase) with FFT computation.

578 In future work, we wish to develop specialized implicit convolutions of
 579 real data for applications in signal processing, such as computing cross cor-
 580 relations and autocorrelations of time series. We are also exploring novel ap-
 581 plications of implicitly dealiased convolutions for computing sparse Fourier
 582 transforms [9], fractional phase Fourier (chirp-z) transforms [1], and partial
 583 Fourier transforms [13?].

584 **Acknowledgment**

585 We thank Prof. Michael Jolly for making us aware of the Basdevant re-
 586 duction. Financial support for this work was provided by Discovery Grant
 587 RES0020458 from the Natural Sciences and Engineering Research Council of
 588 Canada.

589 **Appendix A. Basdevant formulation**

590 *Appendix A.1. 3D incompressible Navier–Stokes equation*

591 A naive implementation of the pseudospectral method for the 3D incom-
 592 pressible Navier–Stokes equation,

$$\frac{\partial u_i}{\partial t} + \frac{\partial D_{ij}}{\partial x_j} = -\frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j^2} + F_i, \quad (\text{A.1})$$

593 where $D_{ij} = \overline{u_i u_j}$, requires three backward FFTs to compute the velocity
594 components from their spectral representations and six forward FFTs of the
595 independent components of the symmetric tensor D_{ij} , for a total of nine FFTs
596 per integration stage. However Basdevant [2] showed that this number can
597 be reduced to eight, by subtracting the divergence of the symmetric matrix
598 $S_{ij} = \delta_{ij} \text{tr } D/3$ from both sides of Eq. (A.1):

$$\frac{\partial u_i}{\partial t} + \frac{\partial(D_{ij} - S_{ij})}{\partial x_j} = -\frac{\partial(p\delta_{ij} + S_{ij})}{\partial x_j} + \nu \frac{\partial^2 u_i}{\partial x_j^2} + F_i. \quad (\text{A.2})$$

599 Since the symmetric matrix $D_{ij} - S_{ij}$ is traceless, it has just five independent
600 components. Together with the three backward FFTs required for the velo-
601 city components u_i , we see that only eight FFTs are required per integration
602 stage. The effective pressure $p\delta_{ij} + S_{ij}$ is solved as usual from the inverse
603 Laplacian of the force minus the nonlinearity.

604 *Appendix A.2. 2D incompressible Navier–Stokes equation*

The vorticity $\mathbf{w} = \nabla \times \mathbf{u}$ evolves according to

$$\frac{\partial \mathbf{w}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{w} = (\mathbf{w} \cdot \nabla) \mathbf{u} + \nu \nabla^2 \mathbf{w} + \nabla \times \mathbf{F},$$

where in two dimensions the vortex stretching term $(\mathbf{w} \cdot \nabla) \mathbf{u}$ vanishes and \mathbf{w} is normal to the plane of motion. For C^2 velocity fields, the curl of the nonlinearity can be written as

$$\frac{\partial}{\partial x_1} \frac{\partial}{\partial x_j} D_{2j} - \frac{\partial}{\partial x_2} \frac{\partial}{\partial x_j} D_{1j} = \left(\frac{\partial^2}{\partial x_1^2} - \frac{\partial^2}{\partial x_2^2} \right) D_{12} + \frac{\partial}{\partial x_1} \frac{\partial}{\partial x_2} (D_{22} - D_{11}),$$

so that the scalar vorticity ω evolves as

$$\frac{\partial \omega}{\partial t} + \left(\frac{\partial^2}{\partial x_1^2} - \frac{\partial^2}{\partial x_2^2} \right) (u_1 u_2) + \frac{\partial^2}{\partial x_1 \partial x_2} (u_2^2 - u_1^2) = \nu \nabla^2 \omega + \frac{\partial F_2}{\partial x_1} - \frac{\partial F_1}{\partial x_2}.$$

605 Two backward FFTs are required to compute u_1 and u_2 in physical space,
606 from which the quantities $u_1 u_2$ and $u_2^2 - u_1^2$ can be calculated and then trans-
607 formed to Fourier space with two additional forward FFTs. The advective
608 term in 2D can thus be calculated with just four FFTs.

609 *Appendix A.3. 3D incompressible MHD equation*

In a similar manner, the incompressible MHD equations can be written

$$\begin{aligned} \frac{\partial u_i}{\partial t} + \frac{\partial(D_{ij} - S_{ij})}{\partial x_j} &= -\frac{\partial(p\delta_{ij} + S_{ij})}{\partial x_j} + \nu \frac{\partial^2 u_i}{\partial x_j^2}, \\ \frac{\partial B_i}{\partial t} + \frac{\partial G_{ij}}{\partial x_j} &= \eta \frac{\partial^2 B_i}{\partial x_j^2}, \end{aligned} \tag{A.3}$$

610 where $D_{ij} = u_i u_j - B_i B_j$, $S_{ij} = \delta_{ij} \text{tr } D/3$, and $G_{ij} = B_i u_j - u_i B_j$. The trace-
 611 less symmetric matrix $D_{ij} - S_{ij}$ has five independent components, whereas
 612 the antisymmetric matrix G_{ij} has only three. Since an additional six FFT
 613 calls are required to compute the components of \mathbf{u} and \mathbf{B} in x space, the
 614 MHD nonlinearity can be computed with 14 FFT calls ($A = 6$, $B = 8$).

615 **Appendix B. Pseudocode**

616 Here we correct a sequencing error in Procedure `fft0padBackward` from
 617 Ref. [3].²

618 We also list Function `cconvA`, which implements optimized implicitly
 619 dealiased convolutions, for the case $A > B$, and Function `cconvB`, for the
 620 case $A < B$.

621 Finally, we document the routines `pretransform` and `posttransform`,
 622 which are used by the 1D Hermitian convolution `conv`, along with the op-
 623 timized routines `convA` and `convB` discussed in the text.

²Minor typographical errors also appeared on page 388 ($0 \dots N$ should be $0 \dots N - 1$),
 page 391 (n should be N), and on p. 400 ($2m_1 - 1$ should be $2m_z - 1$).

```

Input: vector f
Output: vector f, vector u
if noncompact then u[0] ← f[0] - f[m]
else u[0] ← f[0]
if m = 2c then F ← f[c]
parallel for k = 1 to d do
  a ←  $\zeta_{3m}^{-k} \left[ \operatorname{Re} f[k] + \left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right) \operatorname{Re} f[m - k] \right]$ 
  b ←  $-i\zeta_{3m}^{-k} \left[ \operatorname{Im} f[k] + \left(\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \operatorname{Im} f[m - k] \right]$ 
  A ←  $\zeta_{3m}^{k-c-1} \left[ \operatorname{Re} f[c + 1 - k] + \left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right) \operatorname{Re} f[m - c - 1 + k] \right]$ 
  B ←  $-i\zeta_{3m}^{k-c-1} \left[ \operatorname{Im} f[c + 1 - k] + \left(\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \operatorname{Im} f[m - c - 1 + k] \right]$ 
  u[k] ← a - b
  u[c + 1 - k] ← A - B
  f[k] ← f[k] +  $\overline{f[m - k]}$ 
  f[c + 1 - k] ← f[c + 1 - k] +  $\overline{f[m - c - 1 + k]}$ 
  f[m - c - 1 + k] ←  $\overline{a + b}$ 
  f[m - k] ←  $\overline{A + B}$ 
if m = 2c then
  u[c] ← Re F +  $\sqrt{3}$  Im F
  f[c] ← 2 Re F

```

Procedure `pretransform(f,u)` prepares the arrays to be Fourier transformed in Function `conv` from an unpadded vector `f` of m ($m + 1$) values in the compact (noncompact) format, and an auxiliary vector `u` of length $c + 1$, where $c = \lfloor m/2 \rfloor$ and $d = \lfloor (c + 1)/2 \rfloor$. The Fourier origin corresponds to array position 0.

Input: vector \mathbf{f} , real w , vector \mathbf{u}

Output: vector \mathbf{f}

if noncompact then $\mathbf{f}[m] \leftarrow 0$

if $m = 2c$ **and** $m > 2$ **then**

$$\mathbf{a} \leftarrow \mathbf{f}[1] + \zeta_{3m}^{-k} \mathbf{w} + \zeta_{3m}^k \mathbf{u}[1]$$

$$\mathbf{b} \leftarrow \overline{\mathbf{f}[1]} + \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \zeta_{3m}^k \overline{\mathbf{w}} + \left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right) \zeta_{3m}^{-k} \overline{\mathbf{u}[k]}$$

$$\mathbf{A} \leftarrow \mathbf{f}[c] + \zeta_{3m}^{k-c-1} \mathbf{f}[m-1] + \zeta_{3m}^{c+1-k} \mathbf{u}[c]$$

$$\mathbf{f}[1] \leftarrow \mathbf{a}$$

$$\mathbf{f}[c] \leftarrow \mathbf{A}$$

$$\mathbf{f}[m-1] \leftarrow \mathbf{b}$$

parallel for $k = 2c + 2 - m$ **to** $c - d$ **do**

$$\mathbf{a} \leftarrow \mathbf{f}[k] + \zeta_{3m}^{-k} \mathbf{f}[m-c-1+k] + \zeta_{3m}^k \mathbf{u}[k]$$

$$\mathbf{b} \leftarrow \overline{\mathbf{f}[k]} + \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \zeta_{3m}^k \overline{\mathbf{f}[m-c-1+k]} + \left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right) \zeta_{3m}^{-k} \overline{\mathbf{u}[k]}$$

$$\mathbf{A} \leftarrow \mathbf{f}[c+1-k] + \zeta_{3m}^{k-c-1} \mathbf{f}[m-k] + \zeta_{3m}^{c+1-k} \mathbf{u}[c+1-k]$$

$$\mathbf{B} \leftarrow \overline{\mathbf{f}[c+1-k]} + \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \zeta_{3m}^{m-c-1-k} \overline{\mathbf{f}[m-k]} +$$

$$\left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right) \zeta_{3m}^{c+1-m-k} \overline{\mathbf{u}[c+1-k]}$$

$$\mathbf{f}[k] \leftarrow \mathbf{a}$$

$$\mathbf{f}[c+1-k] \leftarrow \mathbf{A}$$

$$\mathbf{f}[m-c-1+k] \leftarrow \mathbf{B}$$

$$\mathbf{f}[m-k] \leftarrow \mathbf{b}$$

if $c+1 = 2d$ **then**

if $d > 1$ **or** $m = 2c+1$ **then** $\mathbf{w} = \mathbf{f}[m-d]$

$$\mathbf{a} \leftarrow \mathbf{f}[d] + \zeta_{3m}^{-k} \mathbf{w} + \zeta_{3m}^k \mathbf{u}[d]$$

$$\mathbf{b} \leftarrow \overline{\mathbf{f}[d]} + \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \zeta_{3m}^k \overline{\mathbf{w}} + \left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right) \zeta_{3m}^{-k} \overline{\mathbf{u}[d]}$$

$$\mathbf{f}[d] \leftarrow \mathbf{a}$$

$$\mathbf{f}[m-d] \leftarrow \mathbf{b}$$

Procedure `posttransform(f,w,u)` is called by Function `conv` to combine the contributions for $r = 0, 1$, and -1 into an implicitly-dealiased Hermitian convolution. The vector \mathbf{f} has length $m(m+1)$ in the compact (noncompact) format, and the auxiliary vector \mathbf{u} has length $c+1$, where $c = \lfloor m/2 \rfloor$ and $d = \lfloor (c+1)/2 \rfloor$. When $m = 2c$, the scalar w contains the overlapped value for $r = 1$ and $k = 1$.

```

Input: vector  $\mathbf{f}$ 
Output: vector  $\mathbf{f}$ , vector  $\mathbf{u}$ 
 $\mathbf{u}[0] \leftarrow \mathbf{f}[m-1]$ 
for  $k = 1$  to  $m-1$  do
   $\mathbf{A} \leftarrow \zeta_{3m}^k [\operatorname{Re} \mathbf{f}[m-1+k] + \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \operatorname{Re} \mathbf{f}[0]]$ 
   $\mathbf{B} \leftarrow i\zeta_{3m}^k [\operatorname{Im} \mathbf{f}[m-1+k] + \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right) \operatorname{Im} \mathbf{f}[0]]$ 
   $\mathbf{C} \leftarrow \mathbf{f}[m-1+k] + \mathbf{f}[0]$ 
   $\mathbf{f}[0] \leftarrow \mathbf{f}[k]$ 
   $\mathbf{f}[k] \leftarrow \mathbf{C}$ 
   $\mathbf{f}[m-1+k] \leftarrow \mathbf{A} + \mathbf{B}$ 
   $\mathbf{u}[k] \leftarrow \overline{\mathbf{A} - \mathbf{B}}$ 
 $\mathbf{f}[0, \dots, m-1] \leftarrow \operatorname{fft}^{-1}(\mathbf{f}[0, \dots, m-1])$ 
 $\mathbf{u}[m] \leftarrow \mathbf{f}[m-1]$ 
 $\mathbf{f}[m-1] \leftarrow \mathbf{u}[0]$ 
 $\mathbf{f}[m-1, \dots, 2m-2] \leftarrow \operatorname{fft}^{-1}(\mathbf{f}[m-1, \dots, 2m-2])$ 
 $\mathbf{u}[0, \dots, m-1] \leftarrow \operatorname{fft}^{-1}(\mathbf{u}[0, \dots, m-1])$ 

```

Procedure `fft0padBackward(f,u)` stores the shuffled $3m$ -padded centered backward Fourier transform values of a compact-format vector of length $2m-1$ in \mathbf{f} and an auxiliary vector \mathbf{u} of length $m+1$.

624 References

- 625 [1] Bailey, D.H., Swarztrauber, P.N., 1991. The fractional Fourier transform
626 and applications. *SIAM review* 33, 389–404.
- 627 [2] Basdevant, C., 1983. Technical improvements for direct numerical sim-
628 ulation of homogeneous three-dimensional turbulence. *Journal of Com-
629 putational Physics* 50, 209–214.
- 630 [3] Bowman, J.C., Roberts, M., 2011. Efficient dealiased convolutions
631 without padding. *SIAM J. Sci. Comput.* 33, 386–406.
- 632 [4] Bowman, J.C., Roberts, M., May 6, 2010. `FFTW++`: A fast Fourier
633 transform `C++` header class for the `FFTW3` library. [http://fftwpp.
634 sourceforge.net](http://fftwpp.sourceforge.net).

```

Input: vectors  $\{f_a\}_{a=0}^{A-1}$ 
Output: vectors  $\{f_b\}_{b=0}^{B-1}$ 
for  $a = 0$  to  $A - 1$  do
|  $u_a \leftarrow \text{fft}^{-1}(f_a)$ 
 $\{u_b\}_{b=0}^{B-1} \leftarrow \text{mult}(\{u_a\}_{a=0}^{A-1})$ 
parallel for  $k = 0$  to  $m - 1$  do
| for  $a = 0$  to  $A - 1$  do
| |  $f_a[k] \leftarrow \zeta_{2m}^k f_a[k]$ 
 $u_{A-1} \leftarrow \text{fft}^{-1}(f_{A-1})$ 
for  $a = A - 2$  to  $0$  do
|  $f_{a+1} \leftarrow \text{fft}^{-1}(f_a)$ 
 $\{f_b\}_{b=1}^B \leftarrow$ 
 $\text{mult}(\{f_a\}_{a=1}^{A-1} \cup \{u_{A-1}\})$ 
for  $b = 0$  to  $B - 1$  do
|  $f_b \leftarrow \text{fft}(f_{b+1})$ 
 $u_{A-1} \leftarrow \text{fft}(u_b)$ 
parallel for  $k = 0$  to  $m - 1$  do
| |  $f_b[k] \leftarrow f_b[k] + \zeta_{2m}^{-k} u_{A-1}[k]$ 
return  $\{f_b/(2m)\}_{b=0}^{B-1}$ 

```

Function `cconvA` returns the in-place implicit dealiased 1D convolution of the complex vectors $\{f_a\}_{a=0}^{A-1}$ using the multiplication operator $\text{mult} : \mathbb{C}^A \rightarrow \mathbb{C}^B$, with $A > B$. All $2A + 2B$ FFTs are out of place.

```

Input: vectors  $\{f_a\}_{a=0}^{A-1}$ 
Output: vectors  $\{f_b\}_{b=0}^{B-1}$ 
for  $a = A - 1$  to  $0$  do
|  $u_a \leftarrow \text{fft}^{-1}(f_a)$ 
parallel for  $k = 0$  to  $m - 1$  do
| for  $a = A - 1$  to  $0$  do
| |  $f_a[k] \leftarrow \zeta_{2m}^k f_a[k]$ 
for  $a = A - 1$  to  $0$  do
|  $f_{a+1} \leftarrow \text{fft}^{-1}(f_a)$ 
 $\{f_b\}_{b=1}^{B-1} \cup \{u_{B-1}\} \leftarrow$ 
 $\text{mult}(\{f_a\}_{a=1}^A)$ 
for  $b = 0$  to  $B - 2$  do
|  $f_b \leftarrow \text{fft}(f_{b+1})$ 
 $f_{B-1} \leftarrow \text{fft}(u_{B-1})$ 
 $\{u_b\}_{b=0}^{B-1} \leftarrow \text{mult}(\{u_a\}_{a=0}^{A-1})$ 
 $u_0 \leftarrow \text{fft}(u_0)$ 
parallel for  $k = 0$  to  $m - 1$  do
|  $f_0[k] \leftarrow f_0[k] + \zeta_{2m}^{-k} u_0[k]$ 
for  $b = 1$  to  $B - 1$  do
|  $u_0 \leftarrow \text{fft}(u_b)$ 
parallel for  $k = 0$  to  $m - 1$ 
do
| |  $f_b[k] \leftarrow f_b[k] + \zeta_{2m}^{-k} u_0[k]$ 
return  $\{f_b/(2m)\}_{b=0}^{B-1}$ 

```

Function `cconvB` returns the in-place implicit dealiased 1D convolution of the complex vectors $\{f_a\}_{a=0}^{A-1}$ using the multiplication operator $\text{mult} : \mathbb{C}^A \rightarrow \mathbb{C}^B$, with $B > A$, with $2A + 2B - 1$ out-of-place and 1 in-place FFTs.

Input: vectors $\{f_a\}_{a=0}^{A-1}$
Output: vectors $\{f_b\}_{b=0}^{B-1}$

```

for a = 0 to A - 1 do
|pretransform( $f_a, u_{A-1}$ )
| $u_a \leftarrow \text{fft}^{-1}(u_{A-1})$ 
 $\{u_b\}_{b=0}^{B-1} \leftarrow \text{mult}(\{u_a\}_{a=0}^{A-1})$ 

 $u_{A-1} \leftarrow \text{fft}^{-1}(f_{A-1}^0)$ 
for a = A - 2 to 0 do
| $f_{a+1}^0 \leftarrow \text{fft}^{-1}(f_a^0)$ 
 $\{f_b^0\}_{b=1}^B \leftarrow$ 
 $\text{mult}(\{f_a^0\}_{a=1}^{A-1} \cup \{u_{A-1}\})$ 

 $u_{A-1} \leftarrow \text{fft}^{-1}(f_{A-1}^1)$ 
for a = A - 2 to 0 do
| $f_{a+1}^1 \leftarrow \text{fft}^{-1}(f_a^1)$ 
 $\{f_b^1\}_{b=1}^B \leftarrow$ 
 $\text{mult}(\{f_a^1\}_{a=1}^{A-1} \cup \{u_{A-1}\})$ 

for b = 0 to B - 1 do
| $u_{A-1} \leftarrow \text{fft}(u_b)$ 
| $f_b^0 \leftarrow \text{fft}(f_{b+1}^0)$ 
| $f_b^1 \leftarrow \text{fft}(f_{b+1}^1)$ 
|posttransform( $\{f_b^0\}_{k=0}^c \cup$ 
| $\{f_b^1\}_{k=2c+2-m}^c, f_b^1[1], u_{A-1}$ )

return  $\{f_b/(3m)\}_{b=0}^{B-1}$ 

```

Function `convA` returns the implicitly dealiased 1D Hermitian convolution of length m ($m + 1$) in the compact (noncompact) format, for $A > B$, with 1 in-place and $3A + 3B - 1$ out-of-place FFTs.

Input: vectors $\{f_a\}_{a=0}^{A-1}$
Output: vectors $\{f_b\}_{b=0}^{B-1}$

```

for a = A - 1 to 0 do
|pretransform( $f_a, u_a$ )
| $u_{a+1} \leftarrow \text{fft}^{-1}(u_a)$ 
| $f_{a+1}^0 \leftarrow \text{fft}^{-1}(f_a^0)$ 
| $f_{a+1}^1 \leftarrow \text{fft}^{-1}(f_a^1)$ 

 $\{f_b^0\}_{b=1}^{B-1} \cup \{u_0\} \leftarrow \text{mult}(\{f_a^0\}_{a=1}^A)$ 
for b = 0 to B - 2 do
| $f_b^0 \leftarrow \text{fft}(f_{b+1}^0)$ 
 $f_{B-1}^0 \leftarrow \text{fft}(u_0)$ 

 $\{f_b^1\}_{b=1}^{B-1} \cup \{u_0\} \leftarrow \text{mult}(\{f_a^1\}_{a=1}^A)$ 
for b = 0 to B - 2 do
| $f_b^1 \leftarrow \text{fft}(f_{b+1}^1)$ 
 $f_{B-1}^1 \leftarrow \text{fft}(u_0)$ 

 $\{u_b\}_{b=1}^{B-1} \cup \{u_0\} \leftarrow \text{mult}(\{u_a\}_{a=1}^A)$ 
 $u_0 \leftarrow \text{fft}(u_0)$ 
posttransform( $\{f_{B-1}^0\}_{k=0}^c \cup$ 
 $\{f_{B-1}^1\}_{k=2c+2-m}^c, f_{B-1}^1[1], u_0$ )
for b = 0 to B - 2 do
| $u_0 \leftarrow \text{fft}(u_{b+1})$ 
|posttransform( $\{f_b^0\}_{k=0}^c \cup$ 
| $\{f_b^1\}_{k=2c+2-m}^c, f_b^1[1], u_0$ )

return  $\{f_b/(3m)\}_{b=0}^{B-1}$ 

```

Function `convB` returns the implicitly dealiased 1D Hermitian convolution of length m ($m + 1$) in the compact (noncompact) format, for $B > A$, with 1 in-place and $3A + 3B - 1$ out-of-place FFTs.

- 635 [5] Cooley, J.W., Tukey, J.W., 1965. An algorithm for the machine cal-
636 culation of complex Fourier series. *Mathematics of Computation* 19,
637 297–301.
- 638 [6] Frigo, M., Johnson, S.G., 2005. The design and implementation of
639 **FFTW3**. *Proceedings of the IEEE* 93, 216–231.
- 640 [7] Gauss, C.F., 1866. *Nachlass: Theoria interpolationis methodo nova*
641 *tractata*, in: *Carl Friedrich Gauss Werke*. Königliche Gesellschaft der
642 *Wissenschaften*, Göttingen. volume 3, pp. 265–327.
- 643 [8] Gottlieb, D., Orszag, S.A., 1977. *Numerical Analysis of Spectral Meth-*
644 *ods: Theory and Applications*. Society for Industrial and Applied Math-
645 *ematics*, Philadelphia.
- 646 [9] Hassanieh, H., Indyk, P., Katabi, D., Price, E., 2012. Simple and
647 practical algorithm for sparse Fourier transform, in: *Proceedings of the*
648 *Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*,
649 *SIAM*. pp. 1183–1194.
- 650 [10] Orszag, S.A., 1971. Elimination of aliasing in finite-difference schemes
651 by filtering high-wavenumber components. *Journal of the Atmospheric*
652 *Sciences* 28, 1074.
- 653 [11] Roberts, M., 2011. *Multispectral Reduction of Two-Dimensional Tur-*
654 *bulence*. Ph.D. thesis. University of Alberta. Edmonton, AB, Canada.
655 http://www.math.ualberta.ca/~bowman/group/roberts_phd.pdf.
- 656 [12] Roberts, M., Leroy, M., Morales, J., Bos, W., Schneider, K., 2014. Self-
657 organization of helically forced MHD flow in confined cylindrical geo-
658 metries. *Fluid Dynamics Research* 46, 061422. URL: <http://stacks.iop.org/1873-7005/46/i=6/a=061422>.
- 660 [13] Ying, L., Fomel, S., 2009. Fast computation of partial Fourier trans-
661 forms. *Multiscale Modeling and Simulation* 8, 110–124.