

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. 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: implicit dealiasing, zero padding, convolution, fast Fourier transform, pseudospectral method, Hermitian symmetry, Nyquist mode, multithreading, parallelization, shared memory, correlation

*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 $\mathcal{O}(m \log m)$ [6, 8] while improving numerical accuracy [9].

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 [4] 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 [4] from being fully parallelized. In Section 4 we demonstrate that multithreaded implicit dealiasing in dimensions greater than one uses much less memory and is much faster than explicit dealiasing. The accomplishments of this work and future directions for research are summarized in Section 5. Implicitly dealiased convolution routines are publicly available in the open-source software library `FFTW++` [5], which is built on top

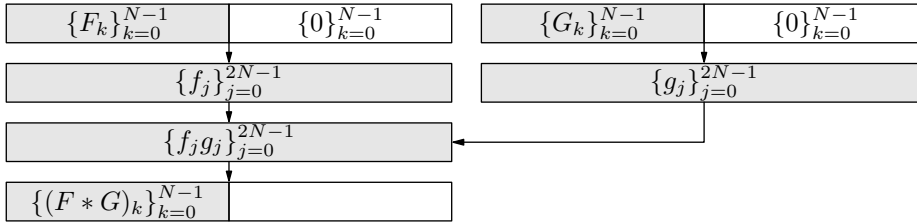


Figure 1: Computing a 1D convolution via explicit zero padding.

38 of the widely used FFTW library [7].

39 2. Dealiasing requirements for convolutions

40 To compute the standard linear convolution $\sum_{p=0}^k F_p G_{k-p}$ for input with
 41 $k \in \{0, \dots, m-1\}$, the data is padded with m zeroes for a total FFT length
 42 of $2m$. We refer to these inputs as *non-centered* and the paddings as *1/2*
 43 *padding*. If the input data is multidimensional with size $m_1 \times \dots \times m_d$, then
 44 the data must be zero padded to $2m_1 \times \dots \times 2m_d$, increasing the buffer size
 45 by a factor of 2^d .

46 For pseudospectral simulations, it is convenient to shift the zero wavenum-
 47 ber in the transformed data to the middle of the array. In this case, the inputs
 48 are $\{F_k\}_{k=-m+1}^{m-1}$ and $\{G_k\}_{k=-m+1}^{m-1}$, which we refer to as *centered*, and their
 49 convolution has components $\sum_{p=k-m+1}^{m-1} F_p G_{k-p}$ for $k = -m+1, \dots, m-1$.
 50 Convolutions on centered inputs require less padding than on non-centered
 51 inputs: data of length $2m-1$ needs to be padded only to length $3m-2$ (nor-
 52 mally extended to $3m$); this is called *2/3 padding* [11]. Explicit zero padding
 53 increases the d -dimensional buffer size in this case by a factor of $(3/2)^d$.

54 A binary convolution can be generalized to an n -ary operator $\ast(F_1, \dots, F_n)_k =$
 55 $\sum_{p_1, \dots, p_n} F_{p_1} \cdots F_{p_n} \delta_{p_1 + \dots + p_n, k}$, where δ is the Kronecker delta. For non-centered
 56 inputs, an n -ary convolution could be computed as a sequence of binary con-
 57 volutions using *1/2 padding*. However, for centered inputs with both negative
 58 and positive frequencies, each binary convolution would have to be padded
 59 further to eliminate all aliased interactions [12]. As a result, n -ary convolu-
 60 tions benefit greatly from implicit dealiasing [4].

61 We consider a generalized convolution operation that takes A inputs and
 62 produces B outputs, where the multiplication performed in the transformed
 63 space can be an arbitrary component-wise operation. In order to make use
 64 of *1/2 padding* or *2/3 padding* (for noncentered or centered inputs, respect-

65 ively), the multiplication operator must be quadratic; if the multiplication
66 operator is of higher degree, one must extend the padding to remove un-
67 desired aliases. To compute a convolution with A inputs and B outputs using
68 the convolution theorem, one performs A backward FFTs to transform the
69 inputs, applies the appropriate multiplication operation on the transformed
70 data, and then performs B forward FFTs to produce the final outputs, for a
71 total of $A + B$ FFTs.

72 The choice of multiplication operator determines the type of convolution.
73 Let $\{f_j\}$ be the inverse Fourier transform of $\{F_k\}$. An autoconvolution can
74 be computed with just two transforms using $A = B = 1$ and the operation
75 $f_j \rightarrow f_j^2$, while an autocorrelation would use $f_j \rightarrow f_j \bar{f}_j$, where \bar{f}_j denotes the
76 complex conjugate of f_j . For the standard binary convolution, there are two
77 inputs and one output, and the multiplication operation is $(f_j, g_j) \rightarrow f_j g_j$.

78 The nonlinear advective term of the 2D incompressible Navier–Stokes vorticity
79 equation can be computed with the operation $(u_x, u_y, \partial\omega/\partial x, \partial\omega/\partial y) \rightarrow$
80 $(u_x \partial\omega/\partial x + u_y \partial\omega/\partial y)$, where $\mathbf{u} = (u_x, u_y)$ is the 2D velocity and $\omega =$
81 $\hat{\mathbf{z}} \cdot \nabla \times \mathbf{u}$ is the z -component of the vorticity; this requires a total of five
82 FFTs ($A = 4$ and $B = 1$). As shown in AppendixA, it is possible to reduce
83 the FFT count for this case to four, with $A = B = 2$. Similarly, in three
84 dimensions, Basdevant [2] showed that the number of FFTs can be reduced
85 from nine to eight, with $A = 3$ and $B = 5$. For 3D magnetohydrodynamic
86 (MHD) flows the operation is $(\mathbf{u}, \boldsymbol{\omega}, \mathbf{B}, \mathbf{j}) \rightarrow (\mathbf{u} \times \boldsymbol{\omega} + \mathbf{j} \times \mathbf{B}, \mathbf{u} \times \mathbf{B})$, where
87 \mathbf{u} is the velocity, $\boldsymbol{\omega} = \nabla \times \mathbf{u}$ is the vorticity, \mathbf{B} is the magnetic field, and \mathbf{j} is
88 the current density ($A = 12$, $B = 6$) [13]. For the Navier–Stokes and MHD
89 equations, the operation is quadratic and the convolution is binary ($n = 2$),
90 with a padding ratio of $2/3$ (since the Fourier modes are symmetric about the
91 origin). In these pseudospectral applications, the physical space quantities
92 are real valued: one can therefore use real-to-complex Fourier transforms,
93 which are about twice as efficient as their complex counterparts.

94 3. One-dimensional implicitly dealiased convolutions

95 Implicit padding allows one to dealias convolutions without having to
96 write, read, and multiply by explicit zero values. This is accomplished by
97 implicitly incorporating the zero values into the top level of a decimated-in-
98 frequency FFT. The extra memory previously used for padding now appears
99 as a decoupled work buffer. One-dimensional implicitly dealiased convo-
100 lutions therefore have the same memory requirements as explicitly padded

101 convolutions. Although in one dimension implicit padding is only slightly
 102 more efficient than explicit zero padding on a single thread, it still has the
 103 advantage of not requiring the copying of user data to a separate enlarged
 104 zero-padded buffer before performing the FFT. We now describe the optimiz-
 105 ed 1D building blocks that will be used in Section 4 to construct higher-
 106 dimensional implicitly dealiased convolutions that are much more efficient
 107 and compact than their explicit counterparts.

108 3.1. Complex convolution

109 Dealiasing the standard convolution $\sum_{p=0}^k F_p G_{k-p}$ for $k = 0, \dots, m-1$
 110 requires extending the input data with zeros from length m to length $N \geq$
 111 $2m-1$, thus removing the beating of two modes with wavenumber $m-1$ that
 112 would otherwise contaminate mode $N = 0 \bmod N$. One generally chooses
 113 $N = 2m$ so that N has a large number of prime factors, resulting in improved
 114 FFT performance.

115 The backward Fourier transform $\{f_j\}_{j=0}^{N-1}$ of the zero-padded input vector
 116 $\{F_k\}_{k=0}^{N-1}$ has components $f_j = \sum_{k=0}^{N-1} \zeta_N^{jk} F_k$, where $\zeta_N = \exp(2\pi i/N)$ denotes
 117 the N^{th} root of unity. The divide-and-conquer strategy of the fast Fourier
 118 transform is based on the property $\zeta_N^r = \zeta_{N/r}$. Since $F_k = 0$ for $k \geq m$,
 119 we can compute the even- and odd-indexed terms of $\{f_j\}_{j=0}^{N-1}$ as separate
 120 subtransforms:

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

121 where $\ell = 0, \dots, m-1$. That is, $\{f_j\}_{j=0}^{N-1}$ can be computed with two Fourier
 122 transforms of length m depending on the input data $\{F_k\}_{k=0}^{m-1}$, with the even-
 123 indexed and odd-indexed parts of the output stored separately. Equation (1)
 124 has a (slightly improved) computational complexity of $\mathcal{O}(N \log m)$, while
 125 avoiding the outermost bit reversal stage and the inconvenience of explicitly
 126 appending m extra zero values to the input data. The (scaled) inverse of
 127 Eq. (1) is given by the forward transform

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

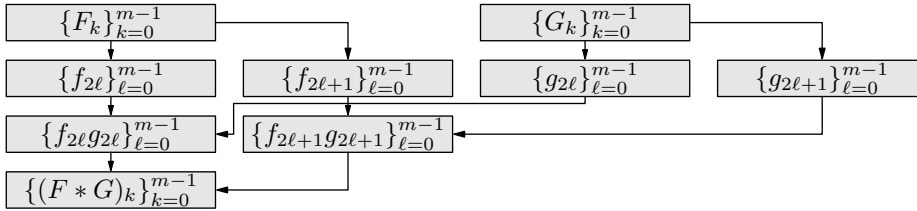


Figure 2: Computing a 1D convolution via implicit dealiasing.

128 again using two Fourier transforms of length m . Equations (1) and (2) can
 129 be combined to compute a dealiasied binary convolution of $\{F_k\}_{k=0}^{m-1}$ and
 130 $\{G_k\}_{k=0}^{m-1}$, as shown in Figure 2 and implemented in pseudocode in Func-
 131 tion `cconv` of Ref. [4]. For each input, two arrays of size m are used instead
 132 of one array of size $2m$. This distinction is the key to the improved effi-
 133 ciency and reduced storage requirements of the higher-dimensional implicit
 134 convolutions described in Section 4. In the 1D complex case, each of the six
 135 complex Fourier subtransforms of size m can be done out of place. Since
 136 implicit dealiasing does not compute the entire inverse Fourier transformed
 137 image at once, we included in our implementation a facility for determining
 138 the spatial coordinates of each point as it is processed. This can be used for
 139 generating an image in x space of the inverse transformed data.

140 As in our previous work [4], we calculate the ζ_N^k factors with a single
 141 complex multiply, using two short pre-computed tables $H_a = \zeta_N^{as}$ and $L_b =$
 142 ζ_N^b , where $k = as + b$ with $s = \lfloor \sqrt{m} \rfloor$, $a = 0, 1, \dots, \lceil m/s \rceil - 1$, and $b =$
 143 $0, 1, \dots, s - 1$. Since these one-dimensional tables occupy only $\mathcal{O}(\sqrt{m})$
 144 complex words, we do not account for them in our storage estimates.

145 Out-of-place FFTs are often more efficient than their in-place counter-
 146 parts, and are more amenable to multithreading. It is not possible to make
 147 use of out-of-place FFTs for explicitly dealiasied convolutions without allocat-
 148 ing additional memory, but the situation is different with implicitly dealiasied
 149 convolutions. For example, in Function `cconv` of Ref. [4], all FFTs are out
 150 of place. Our general function `cconv` in this work has $A + B - 1$ out-of-place
 151 FFTs and $A + B + 1$ in-place FFTs.

152 For $A > B$, the multiplication operator will free buffers that can be re-
 153 used, and it is possible to compute the convolution with all FFTs out of
 154 place. The idea is that the input and work buffers can be processed separ-
 155 ately, and, after applying the multiplication operator, the data in the last of
 156 the A work buffers is no longer needed. We make use of this buffer to imple-

157 ment out-of-place transforms. In Function `cconvA`, we present an algorithm
 158 for an implicitly dealiased convolution with $A > B$ in which all $2A + 2B$
 159 FFTs of length m are out of place. Likewise, when $A < B$, Function `cconvB`
 160 shows that all but one of the $2A + 2B$ FFTs can be performed out-of-place.

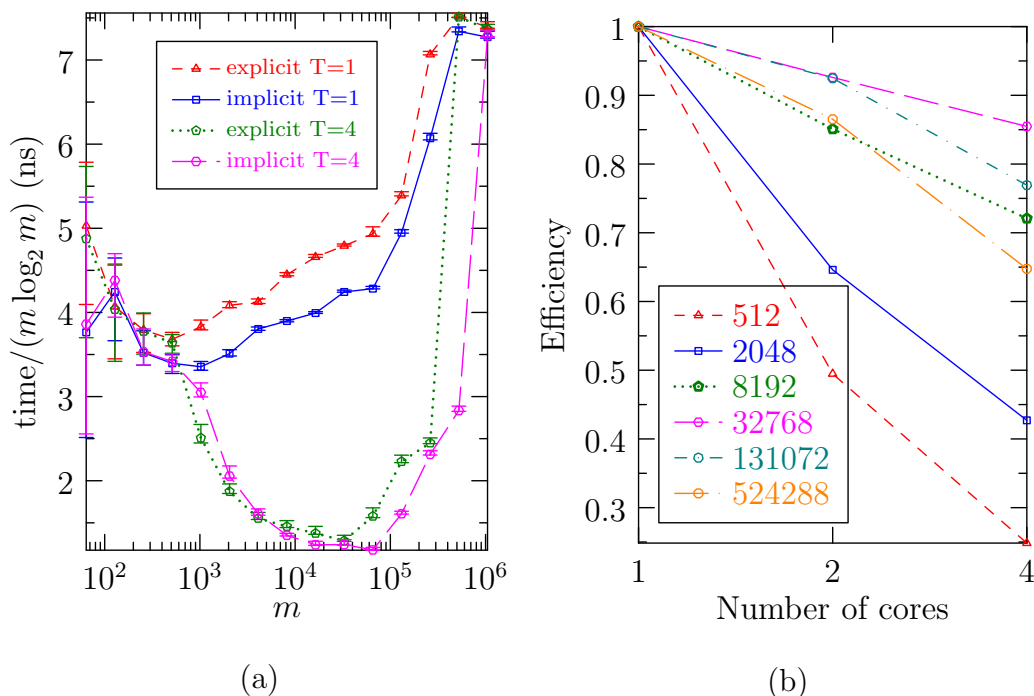


Figure 3: In-place 1D complex convolutions of length m : (a) comparison of computation times for explicit and implicit dealiasing using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency of implicit dealiasing versus number of threads. For efficiency m is chosen to be a power of two.

161 When $A = 2$ and $B = 1$, Function `cconvA` runs a few percent faster than
 162 Function `cconv` from Ref. [4], thanks to improvements in the loop structure
 163 in the pre- and post-processing stages. Thus, in one dimension, as seen in
 164 Figure 3(a), implicit dealiasing on a single thread is now on average 12%
 165 faster than explicit zero padding.

166 3.1.1. Multithreaded Complex 1D Binary Convolutions

167 We parallelize Functions `cconv`, `cconvA`, and `cconvB` using OpenMP in
 168 our pre/post-processing phases and in the multiplication operator, while tak-
 169 ing advantage of the multithreading built into the FFTW library. In Fig-
 170 ure 3(a), we compare the speed of the implicit and explicit algorithms using

171 one and four threads. Using one thread, the implicit method is on average
 172 1.12 times faster than the explicit method, whereas using four threads the
 173 performance improvement is a factor of 1.04 to 2.6 for $m \geq 8192$. The reason
 174 that the explicit version benefits from parellization at smaller m values than
 175 implicit dealiasing is a simple consequence of the fact that the vector sizes for
 176 explicit dealiasing are twice as large, due to the memory wasted on padding.
 177 The error bars in the timing figures indicate the lower and upper one-sided
 178 standard deviations, as given in Ref. [4]. The number of samples varied
 179 from several million for small data sizes to 20 for larger data sizes. In Fig-
 180 ure 3(b), we observe for $m = 2048$ to 524288 that the implicit method with
 181 four threads has a parallel efficiency of 43% to 85%, giving a speedup of a
 182 factor of 1.7 to 3.4.

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

190 3.2. Centered data formats

191 In this work, we extend the treatment of centered Fourier input data
 192 $\{F_{-m+1}, \dots, F_{m-1}\}$ for even m from Ref. [4], to all natural numbers m . We
 193 also implement an optional new data layout $\{F_{-m}, \dots, F_{m-1}\}$. In addition to
 194 handling convolutions of Fourier transformed real-space data of even length,
 195 this extended format can yield significant performance improvements, even if
 196 the additional mode F_{-m} is simply set to zero. We refer to $\{F_{-m+1}, \dots, F_{m-1}\}$
 197 as the *compact* format and $\{F_{-m}, \dots, F_{m-1}\}$ as the *noncompact* format. In
 198 particular, in the noncompact case, a fully multithreaded implementation
 199 (Procedure `fft1padBackward`) is possible since it doesn't have the loop de-
 200 pendency seen in Procedure `fft0padBackward`.² The noncompact format is
 201 consistent with the output of a real-to-complex FFT.

²We correct here a sequencing error in the pseudocode for Procedure `fft0padBackward` in Ref. [4]. Minor 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 $2m_z - 1$).

202

```

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
  |  $f_b \leftarrow \text{fft}(f_b)$ 
  |  $u_0 \leftarrow \text{fft}(u_b)$ 
  | 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 `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.

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

```

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

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

209 formats.

210 In the compact (noncompact) format, it is convenient to shift the Fourier
 211 origin so that the $k = 0$ mode is indexed as array element $m - 1$ (m). This
 212 shift, which can be built into implicitly dealiased convolution algorithms at
 213 no extra cost, allows for more convenient coding of wavenumber loops since
 214 the high-wavenumber cutoff is naturally aligned with the array boundaries.

215 In the compact case, where $N = 2m - 1$, one needs to pad to $N \geq 3m - 2$
 216 to prevent modes with wavenumber $m - 1$ from beating together to contaminate
 217 the mode with wavenumber $-m + 1$. The ratio of the number of physical
 218 to total modes, $(2m - 1)/(3m - 2)$, is then asymptotic to $2/3$ for large m
 219 [11]. With explicit padding, for efficiency reasons one normally chooses the
 220 padded vector length N to be a power of 2, with $m = \lfloor (N + 2)/3 \rfloor$, while for
 221 implicit padding, it is advantageous to choose the subtransform length m to
 222 be a power of 2. Moreover, it is convenient to pad implicitly slightly beyond
 223 $3m - 2$, to $N = 3m$, to support a radix 3 subdivision at the outer level.

224 In the case of an even number $2m - 2$ of spatial data points, one must
 225 use the noncompact data format, with modes running from $-m$ to $m - 1$.
 226 When $N = 3m$, the most negative (Nyquist) wavenumber $-m$ can constructively
 227 beat with itself, producing an alias in mode $-m + (-m) = -2m =$
 228 $m \bmod 3m$, which is equivalent to itself modulo $2m$. To remove this alias we
 229 set the Nyquist mode to zero at the end of the convolution, after account-
 230 ing for its effects on the other modes. We note that there are no aliases in
 231 $\{-m, \dots, 2m - 1\}$ arising from the interaction of mode $-m$ with any of the
 232 other modes. As we will see in Section 3.2.1, those interactions can (and
 233 in fact, should) be retained. We split the coefficient for $k = -m$ equally
 234 between F_{-m} and its equivalent F_m (modulo $2m$); in Section 3.2.1, we will
 235 see that this is important for maintaining Hermitian symmetry and the cor-
 236 responding reality of the spatial field.

237 We now describe a noncompact implicitly dealiased centered Fourier trans-
 238 form; the compact case is obtained by setting $F_{-m} = F_m = 0$. Suppose then
 239 that $F_k = 0$ for $k > m$. On decomposing $j = (3\ell + r) \bmod N$, where
 240 $r \in \{-1, 0, 1\}$, the substitution $k' = m + k$ allows us to write the backward
 241 transform as

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

242 where, on recombining F_m into F_{-m} ,

$$w_{k,r} \doteq \begin{cases} F_0 + \operatorname{Re} \zeta_3^{-r} F_{-m} & \text{if } k = 0, \\ \zeta_{3m}^{rk} (F_k + \zeta_3^{-r} F_{k-m}) & \text{if } 1 \leq k \leq m-1. \end{cases} \quad (4)$$

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

$$3mF_k = \sum_{r=-1}^1 \zeta_{3m}^{-rk} \sum_{\ell=0}^{m-1} \zeta_m^{-\ell k} f_{3\ell+r}, \quad k = -m+1, \dots, m-1, \quad (5)$$

244 with the Nyquist mode F_{-m} set to zero. The use of the remainder $r = -1$
 245 instead of $r = 2$ allows us to exploit the optimization $\zeta_{3m}^{-k} = \overline{\zeta_{3m}^k}$ in Eqs. (4)
 246 and (5). The number of complex multiplies needed to evaluate Eq. (4) for
 247 $r = \pm 1$ can be reduced by computing the intermediate complex quantities
 248 $A_k \doteq \zeta_{3m}^k (\operatorname{Re} F_k + \zeta_3^{-1} \operatorname{Re} F_{k-m})$ and $B_k \doteq i \zeta_{3m}^k (\operatorname{Im} F_k + \zeta_3^{-1} \operatorname{Im} F_{k-m})$,
 249 where $\zeta_3^{-1} = (-\frac{1}{2}, -\frac{\sqrt{3}}{2})$, so that for $k > 0$, $w_{k,1} = A_k + B_k$ and $w_{k,-1} =$
 250 $\overline{A_k - B_k}$. The resulting noncompact backward transform is given in Procedure
 251 `fft1padBackward`; its inverse is given in Procedure `fft1padForward`.
 252 The compact versions of these routines are Procedure `fft0padBackward` and
 253 its inverse, Procedure `fft0padForward` from Ref. [4].

254 3.2.1. Centered Hermitian Implicitly Padded 1D FFT

255 The Fourier transform of real data satisfies the Hermitian symmetry
 256 $F_{-k} = \overline{F_k}$. This implies that the Fourier coefficient corresponding to $k = 0$ is
 257 real. There is a further consequence of this symmetry when the length N of
 258 the discrete transform $\sum_{j=0}^{N-1} \zeta_N^{-kj} f_j$ is even. Due to the periodicity of the discrete
 259 transform in N , the highest frequency (Nyquist) mode must also be real:
 260 $F_{N/2} = \overline{F_{-N/2}} = \overline{F_{N/2}}$. Letting $m = \lfloor N/2 \rfloor + 1$, in the case where N is even,
 261 the $2m$ modes can therefore be indexed as $\{F_{-m+1}, \dots, F_m\}$ where F_0 and F_m
 262 are real. An odd number $2m - 1$ of modes is indexed as $\{F_{-m+1}, \dots, F_{m-1}\}$.

263 Hermitian symmetry can be used to reduce the computational complexity
 264 and storage requirements of real-to-complex and complex-to-real Fourier
 265 transforms by a factor of about two. A one-dimensional convolution of Hermitian
 266 data only requires the data corresponding to non-negative wavenumbers.
 267 In the compact case, with modes in $\{-m+1, \dots, m-1\}$, the unsymmetrized
 268 physical data needs to be padded with at least $m-1$ zeros,
 269 just as in Section 3.2. Hermitian symmetry thus necessitates padding the m
 270 non-negative wavenumbers with at least $c \doteq \lfloor m/2 \rfloor$ zeros. The resulting $2/3$

```

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

```

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

```

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.

padding ratio (for even m) turns out to work particularly well for developing implicitly dealiased centered Hermitian convolutions. As in the centered case, we again choose the Fourier size to be $N = 3m$.

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

Fortunately, the choice $N = 3m$ also works for the noncompact case, provided the entry for the Nyquist mode, which must be real, is zeroed at the end of the convolution. For example, direct autoconvolution of the Hermitian data $\{(1, 0), (2, 3), (4, 0)\}$ yields $\{(59, 0), (20, -18), (3, 12)\}$. With $m = 3$, the compact implicitly dealiased convolution in Function `conv` produces identical results. For $m = 2$, the noncompact version yields the correct values $\{(59, 0), (20, -18)\}$ for the first m elements only if the data $(3, 12)$ for the Nyquist mode at $k = m$ is taken into account.

Let us now describe a noncompact implicitly padded Hermitian FFT; the compact case can then be obtained by setting $F_m = 0$. Given that $F_k = 0$ for $k > m$, the backward (complex-to-real) transform appears as Eq. (3), but

290 now with

$$w_{k,r} \doteq \begin{cases} F_0 + \operatorname{Re} \zeta_3^{-r} F_m & \text{if } k = 0, \\ \zeta_{3m}^{rk} (F_k + \zeta_3^{-r} \overline{F_{m-k}}) & \text{if } 1 \leq k \leq m-1. \end{cases} \quad (6)$$

291 We note for $k > 0$ that $w_{k,r}$ obeys the Hermitian symmetry $w_{k,r} = \overline{w_{m-k,r}}$, so
 292 that the Fourier transform $\sum_{k=0}^{m-1} \zeta_m^{\ell k} w_{k,r}$ in Eq. (3) will indeed be real valued.
 293 This allows us to build a backward implicitly dealiased centered Hermitian
 294 transform using three complex-to-real Fourier transforms of the first $c+1$
 295 components of $w_{k,r}$ (one for each $r \in \{-1, 0, 1\}$). The forward transform is
 296 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
 297 $f_{3\ell+r}$ is real, a real-to-complex transform can be used to compute the first
 298 $c+1$ frequencies of $\sum_{\ell=0}^{m-1} \zeta_m^{-\ell k} f_{3\ell+r}$; the remaining $m-c-1$ frequencies are
 299 then computed using Hermitian symmetry.

300 3.2.2. Multithreaded Hermitian 1D Binary Convolution

301 An in-place implicitly padded Hermitian convolution was previously de-
 302 scribed in Function `conv` of Ref. [4] for the case of $2M$ inputs and one output,
 303 where the multiplication operator was restricted to a dot product. How-
 304 ever, that algorithm cannot be efficiently applied to the autoconvolution
 305 case (with just one input and one output), to pseudospectral simulations of
 306 3D Navier–Stokes and magnetohydrodynamic flows, or to the reduced-FFT
 307 scheme of Basdevant for 2D turbulence (see AppendixA). Furthermore, in-
 308 terloop dependencies at the outermost level prevent it from being multith-
 309 readed. Moreover, to facilitate an in-place implementation, the transformed
 310 values for $r = 1$ were awkwardly stored in reverse order in the upper half
 311 of the input vector, exploiting the quadratic nature of the real-space mul-
 312 tiplication operator. By unrolling the outer loop of the in-place Hermitian
 313 1D convolution, these deficiencies can be eliminated, resulting in the fully
 314 multithreaded implementation in Function `conv`, generalized to handle A
 315 inputs and B outputs and an arbitrary quadratic multiplication operator
 316 `mult` : $\mathbb{R}^A \rightarrow \mathbb{R}^B$.

```

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

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

for  $b = 0$  to  $B - 1$  do
   $u_0 \leftarrow \text{fft}(u_b)$ 
   $f_b^0 \leftarrow \text{fft}(f_b^0)$ 
   $f_b^1 \leftarrow \text{fft}(f_b^1)$ 
  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 `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.

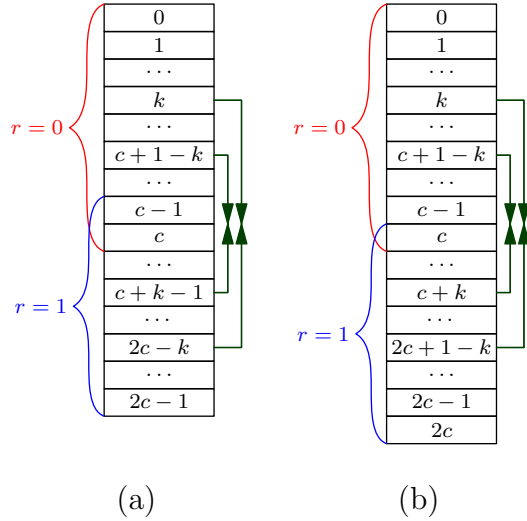


Figure 4: Loop unrolling for the Hermitian 1D convolution when (a) $m = 2c$ and (b) $m = 2c + 1$.

317 To multithread Procedure `pretransform`, we unroll two iterations of the
 318 loop from Procedure `build` in Ref. [4] to read, process, and write the entries
 319 for the elements indexed by k , $m - k$, $c + 1 - k$, and $m - c - 1 + k$ simultaneously,
 320 for $k = 1, \dots, \lceil c/2 \rceil$, as shown in Figure 4. The implicitly padded transformed
 321 data for remainders $r = 0$ and $r = 1$ is stored in the input data vector f ,
 322 whereas the data for remainder $r = -1$ is stored in the auxiliary vector u . In
 323 the frequently encountered case where $m = 2c$, the values at position $c - 1$

324 and c overlap for the remainders $r = 0$ and $r = 1$, whereas when $m = 2c + 1$
 325 there is only one overlapping value, at index c . In the noncompact case,
 326 any Nyquist inputs $f[m]$ and $g[m]$ are properly accounted for, but on output
 327 $f[m]$ is set to zero, for consistency with Hermitian symmetry. A similar
 328 loop unrolling is used in a revised implementation of the post-processing
 329 phase (see Procedure `posttransform`) to allow for an arbitrary number of
 330 inputs A and outputs B .

331 In the pseudocode, the portions of the arrays corresponding to remainders
 332 $r = 0$ and $r = 1$ are shown in Figure 4 and distinguished by the superscripts
 333 0 and 1. Since these portions overlap when written to the array f , additional
 334 code is required to save and restore the overlapping elements in the actual
 335 in-place implementation.

336 In our general Function `conv`, only $A + B - 2$ of the $3A + 3B$ FFTs can
 337 be performed out of place. When $A = B$ this is the best that one can do:
 338 for example, for an autoconvolution ($A = B = 1$), there is no free buffer
 339 available that would enable the use of out-of-place transforms. Nevertheless
 340 when $A > B$ or $B > A$ the optimized Function `convA` or `convB`, respectively,
 341 performs one FFT in place and the remaining $3A + 3B - 1$ FFTs out of place.
 342 Even with one thread, for $A = 2$ and $B = 1$, Functions `conv` and `convA` both
 343 have slightly better performance than Function `conv` in Ref. [4], primarily
 344 due to the removal of loop interdependence. Most importantly, `conv`, `convA`,
 345 and `convB` have fully parallelized pre- and post-processing phases and use
 346 FFTW's built-in parallel FFTs, which are typically much more efficient when
 347 the transforms are out of place. The new routines accept either compact
 348 or noncompact inputs and can therefore also benefit from the performance
 349 advantage of the noncompact data format discussed in Section 3.2.

350 In the case of a binary convolution with two input vectors and one output
 351 vector, a fully in-place convolution requires a total of nine Hermitian Four-
 352 tier transforms of size m , for an overall computational scaling of $\frac{9}{2}Km \log_2 m$
 353 operations, where $K = 34/9$ [4], in agreement with the leading-order scaling
 354 of an explicitly padded centered Hermitian convolution. In our new imple-
 355 mentation, eight of the nine Fourier transforms can now be performed out
 356 of place, using the same amount of memory ($6c + 2$ words in the compact
 357 case) as required to compute a centered Hermitian convolution with explicit
 358 padding.

359 As seen in Fig. 5, the efficiency of the resulting implicitly dealiased
 360 centered Hermitian convolution is comparable to an explicit implementa-
 361 tion. For each algorithm, we benchmark only those vector lengths that yield

```

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 ← ζ3m-k [Re f[k] + (-½, √3/2) Re f[m - k]]
    b ← -iζ3m-k [Im f[k] + (½, -√3/2) Im f[m - k]]
    A ← ζ3mk-c-1 [Re f[c + 1 - k] + (-½, √3/2) Re f[m - c - 1 + k]]
    B ← -iζ3mk-c-1 [Im f[c + 1 - k] + (½, -√3/2) Im f[m - c - 1 + k]]
    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 + √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 f , real w , vector u

Output: vector f

if noncompact then $f[m] \leftarrow 0$

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

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

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

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

$$f[1] \leftarrow a$$

$$f[c] \leftarrow A$$

$$f[m-1] \leftarrow b$$

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

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

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

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

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

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

$$f[k] \leftarrow a$$

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

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

$$f[m-k] \leftarrow b$$

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

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

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

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

$$f[d] \leftarrow a$$

$$f[m-d] \leftarrow b$$

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 f has length m ($m+1$) in the compact (noncompact) format, and the auxiliary vector 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: 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 ← fft-1(u_{A-1})
{u_b}_{b=0}^{B-1} ← mult({u_a}_{a=0}^{A-1})

u_{A-1} ← fft-1(f_{A-1}^0)
for a = A - 2 to 0 do
  | f_{a+1}^0 ← fft-1(f_a^0)
{f_b^0}_{b=1}^B ←
mult({f_a^0}_{a=1}^{A-1} ∪ {u_{A-1}})

u_{A-1} ← fft-1(f_{A-1}^1)
for a = A - 2 to 0 do
  | f_{a+1}^1 ← fft-1(f_a^1)
{f_b^1}_{b=1}^B ←
mult({f_a^1}_{a=1}^{A-1} ∪ {u_{A-1}})

for b = 0 to B - 1 do
  | u_{A-1} ← fft(u_b)
  | f_b^0 ← fft(f_{b+1}^0)
  | f_b^1 ← fft(f_{b+1}^1)
  | posttransform({f_b^0}_{k=0}^c ∪
  | {f_b^1}_{k=2c+2-m}, 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} ← fft-1(u_a)
  | f_{a+1}^0 ← fft-1(f_a^0)
  | f_{a+1}^1 ← fft-1(f_a^1)

{f_b^0}_{b=1}^{B-1} ∪ {u_0} ← mult({f_a^0}_{a=1}^A)
for b = 0 to B - 2 do
  | f_b^0 ← fft(f_{b+1}^0)
f_{B-1}^0 ← fft(u_0)

{f_b^1}_{b=1}^{B-1} ∪ {u_0} ← mult({f_a^1}_{a=1}^A)
for b = 0 to B - 2 do
  | f_b^1 ← fft(f_{b+1}^1)
f_{B-1}^1 ← fft(u_0)

{u_b}_{b=1}^{B-1} ∪ {u_0} ← mult({u_a}_{a=1}^A)
u_0 ← fft(u_0)

posttransform({f_{B-1}^0}_{k=0}^c ∪
{f_{B-1}^1}_{k=2c+2-m}, f_{B-1}^1[1], u_0)
for b = 0 to B - 2 do
  | u_0 ← fft(u_{b+1})
  | posttransform({f_b^0}_{k=0}^c ∪
  | {f_b^1}_{k=2c+2-m}, 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.

362 optimal performance. The optimal values of m for the explicit version are
 363 $\lfloor (2^p + 2)/3 \rfloor$ for natural numbers p , whereas for the implicit version the op-
 364 timal values are powers of two, so direct comparison of the methods using
 365 optimal problem sizes is not possible. Instead, we compare the two methods
 366 using a linear interpolation (with respect to $\log m$) of the execution time res-
 367 scaled by the computational complexity of the algorithm. With one thread,
 368 the implicit version runs between 5% and 23% faster for $m \geq 2048$; with four
 369 threads, the implicit version is between 12% and 105% faster for $m \geq 65536$,
 370 as shown in Fig. 5(a). We demonstrate the parallel efficiency of the implicit
 371 routine in Fig. 5(b) using one, two, and four threads. For $m \geq 8192$, the
 372 parallel efficiency of the implicit method with four threads is between 45%
 373 and 68%, giving a speedup of a factor of 1.8 to 2.7.

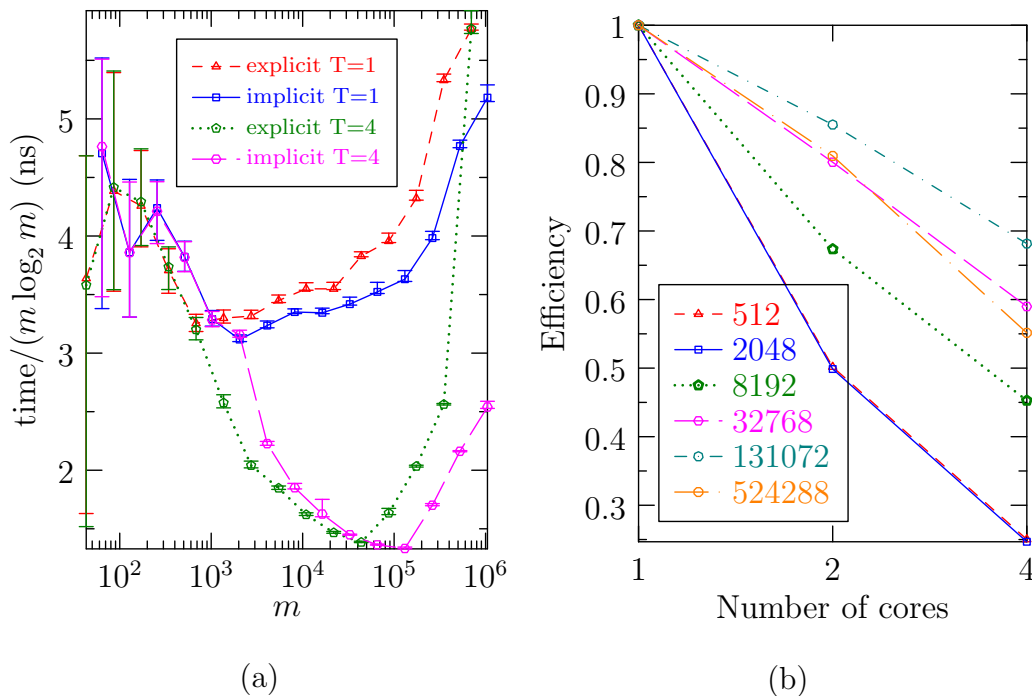


Figure 5: In-place 1D Hermitian convolutions of length m : (a) comparison of computation times for explicit and implicit dealiasing using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency of implicit dealiasing versus number of threads. For implicit convolutions, m is chosen to be a power of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

374 4. Higher-dimensional convolutions

375 A d -dimensional convolution can be computed by performing an inverse
376 FFT of size $m_1 \times \dots \times m_d$, applying the appropriate multiplication on the
377 transformed data, followed by an FFT back to the original space. Equival-
378 ently, one can perform $\prod_{i=2}^d m_i$ inverse FFTs in the first dimension, followed
379 by m_1 convolutions of dimension $d - 1$, and finally $\prod_{i=2}^d m_i$ FFTs in the first
380 dimension. The innermost operation of a recursive multidimensional convo-
381 lution thus reduces to a 1D convolution. Using this decomposition, one can
382 reuse the work buffer for each implicitly dealiased subconvolution, thereby re-
383 ducing the total memory demand relative to the explicit d -dimensional deali-
384 asing requirement. For multithreaded implicitly dealiased convolutions, the
385 initial inverse FFT can be parallelized by dividing the $\prod_{i=2}^d m_i$ 1D FFTs and
386 m_1 subconvolutions between the T threads. Since each implicitly dealiased
387 subconvolution requires a work buffer, the total memory requirement grows
388 with the number of threads, but is still much lower than that required for
389 explicit multidimensional convolutions when $T \ll m_1$. When $T \leq m_1$, we
390 compute T subconvolutions at a time, using one inner thread per subconvo-
391 lution to avoid over-subscription. Otherwise, if $T > m_1$, we parallelize only
392 the inner subconvolution (over all T threads).

393 A single-threaded 2D implicitly 1/2-padded complex convolution is shown
394 in Figure 6. Each input buffer is implicitly padded and inverse Fourier trans-
395 formed in the x direction to produce the data shown in the square boxes. An
396 implicitly padded inverse FFT is then performed in the y direction, column-
397 by-column, using a one-dimensional work buffer, to produce a single column
398 of the Fourier transformed image, depicted in yellow. The Fourier trans-
399 formed columns of two inputs F and G are then multiplied pointwise and
400 stored back into the F column. At this point, the y forward-padded FFT can
401 then be performed, with the result stored in the lower-half of the column,
402 next to the previously processed data shown in red. The process is repeated
403 on the remaining columns, shifting and reusing the work buffer. Once all
404 the columns have been processed, a forward-padded FFT in the x direction
405 produces the final convolution in the left-hand half of the F buffer.

406 The reuse of subconvolution work memory allows the convolution to be
407 computed using less total memory: for 1/2 padded convolutions, the memory
408 requirement per input is only about twice what would be required if deali-
409 asing was outright ignored. This represents a memory savings of a factor
410 of 2^{d-1} as compared to explicit padding; for 2/3 padded convolutions, the

411 memory savings factor is $(3/2)^{d-1}$. In addition to having reduced memory
 412 requirements, implicitly dealiased multidimensional convolutions are significantly
 413 faster than their explicit counterparts, due to better data locality and
 414 cache management, along with the fact that transforms of data known to be
 415 zero are automatically avoided.

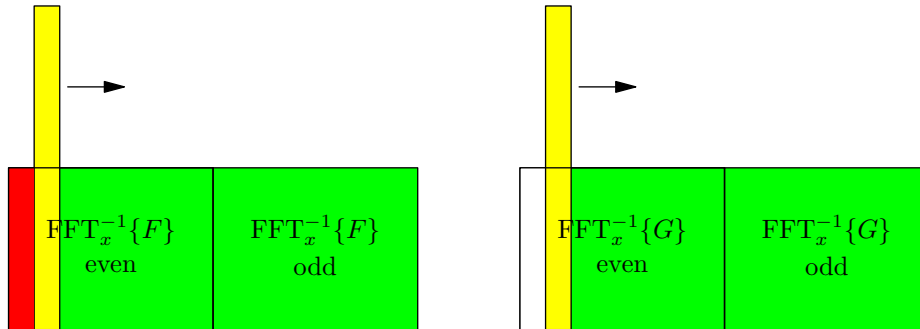


Figure 6: The reuse of memory in a 2D complex implicitly dealiased convolution: after applying a 1D y convolution to the yellow column, the upper half is reused for the next column.

416 In the following subsections, we show that the algorithms developed in
 417 Section 3 can be used as building blocks to construct efficient implicitly
 418 padded higher-dimensional convolutions.

419 4.1. Complex 2D convolution

420 Pseudocode for the implicitly padded transforms described by Eqs. (1)–
 421 (2) was given in Ref. [4] as Procedures `fftpadBackward` and `fftpadForward`.
 422 In order to compute a 2D convolution in parallel, the loops in these proced-
 423 ures were parallelized, and parallel FFTs were used. Since the input and
 424 output of these routines are multidimensional and the required FFT is one-
 425 dimensional, we use the FFTW multiple 1D FFT routine. A multithreaded
 426 version of this routine is available in the FFTW library, but we found that
 427 its parallel performance was sometimes lacking. This was somewhat surpris-
 428 ing, as there exists a simple algorithm to parallelize such problems: if one
 429 wishes to perform M FFTs using T threads, one can simply divide the M
 430 FFTs among the T threads, with any remaining r FFTs distributed among
 431 the first r threads. At run time, we automatically test for the possibility
 432 that this decomposition is faster than FFTW's parallel multiple FFT and use

433 whichever algorithm runs faster. This yielded a significant improvement in
 434 the parallel performance of our convolutions.

435 As shown in Fig. 7(a), the resulting implicit 2D algorithm dramatically
 436 outperforms the explicit version: using one thread, the mean speedup is
 437 a factor of 1.5, with a maximum speedup of 1.8. Using four threads, the
 438 mean speedup over the parallel explicit version is approximately 2.6, with
 439 a maximum speedup factor of 4.5. Fig. 7(b) shows the parallel efficiency
 440 of the 2D implicitly dealiased complex convolution for a variety of problem
 441 sizes. The parallel efficiency for the implicit routine ranges from 58% to 92%
 442 with four threads, for a speedup of 2.3 to 3.7 relative to one thread. The
 443 explicit routine has a parallel efficiency between 25% and 90%. Notably, the
 444 2D explicit version has poor parallel performance for problem sizes of 512^2
 445 and above using FFTW’s built-in multithreading.

446 Because the same temporary arrays u and v are used for each column
 447 of the convolution, the memory requirement is $2Cm_xm_y + Tm_y$ complex
 448 words using T threads, where $C = \max\{A, B\}$. Assuming that $T < 2m_x$,
 449 this is far less than the $4Cm_xm_y$ complex words needed for an explicitly
 450 padded convolution.

451 4.2. Centered Hermitian 2D convolution

452 In two dimensions, the Fourier-centered Hermitian symmetry appears as
 453 $F_{-k,-\ell} = \overline{F_{k,\ell}}$. This symmetry is exploited in the centered Hermitian convo-
 454 lution algorithm shown for the noncompact case in Function `conv2`. As with
 455 the 1D Hermitian convolution, one has the option to use a compact or non-
 456 compact data format. For the compact data format, the array has dimensions
 457 $\{-m_x + 1, \dots, m_x - 1\} \times \{0, \dots, m_y - 1\}$, whereas the noncompact version
 458 has dimensions $\{-m_x, \dots, m_x - 1\} \times \{0, \dots, m_y\}$. One can also perform convo-
 459 lutions on data that is compact in one direction and noncompact in the
 460 other. For serial computations, the best performance typically is achieved
 461 when the x direction is compact and the y direction is noncompact, so that
 462 each dimension is odd, to reduce cache associativity issues. However, when
 463 running on more than one thread, the noncompact format should be used
 464 in the x direction since Procedures `fft1padBackward` and `fft1padForward` are
 465 fully multithreaded.

466 While the noncompact case requires slightly more memory than the com-
 467 pact case, one advantage of the noncompact version is that the output of
 468 the Fourier transform of $(2m_x - 2) \times (2m_y - 2)$ real values corresponds to
 469 the modes $\{-m_x, \dots, m_x - 1\} \times \{0, \dots, m_y\}$, where the Fourier origin has

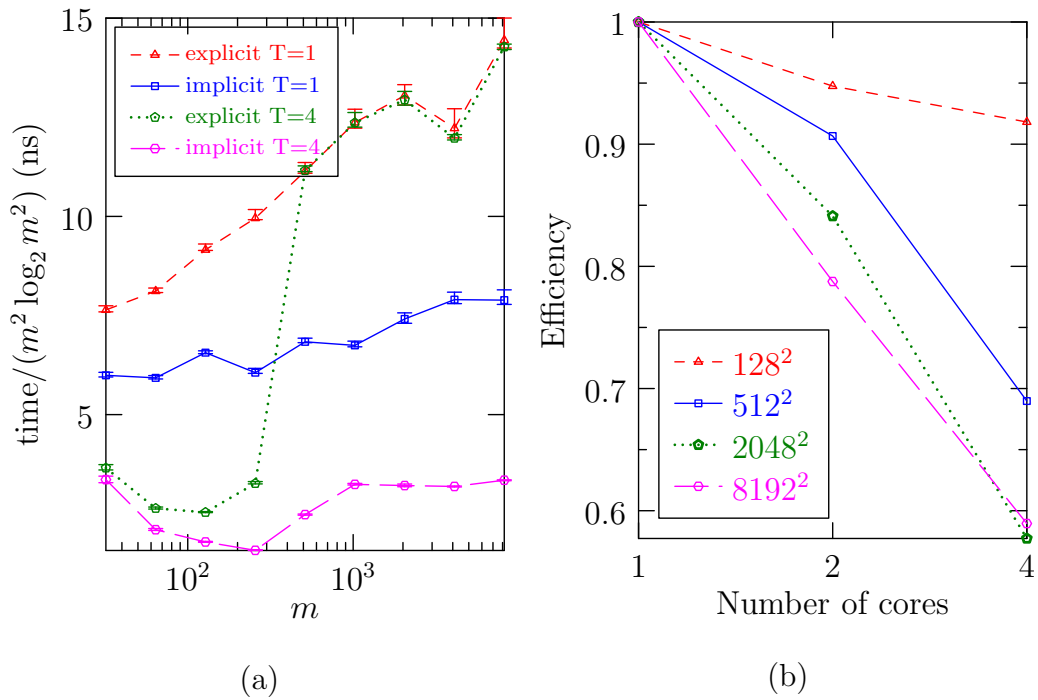


Figure 7: In-place 2D complex convolutions of size $m \times m$: (a) comparison of computation times for explicit and implicit dealiasing using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency of implicit dealiasing versus number of threads. Here m is chosen to be a power of two.

470 been shifted in the x direction to the middle of the array. Moreover, one
 471 is able to use the extra memory in the x direction for temporary storage,
 472 and having $m_y + 1$ variables in the y direction avoids latency issues with
 473 cache associativity when m_y is a power of two. These factors combine to
 474 give the noncompact format a performance advantage over the compact one:
 475 the noncompact case is typically slightly faster than the compact case when
 476 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
  | | fftpadBackward( $f_a^T[j], U_a^T[j]$ )
  | 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
  | | | fftpadForward( $f_b^T[j], U_b^T[j]$ )
  | return  $f$ 

```

477

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 $\{U_a\}_{a=0}^{A-1}$.

478 The explicit version requires storage for $9Cm_x(m_y + 1)/2$ complex words,
 479 where $C = \max\{A, B\}$. For the noncompact case, the implicit version using
 480 T threads requires storage for $3Cm_x(m_y + 1) + TC(\lfloor m_y/2 \rfloor + 1)$ complex
 481 words, which is much less than the explicit case when $m_x \geq T$. As shown
 482 in Fig. 8(a), implicit padding again yields a dramatic improvement in speed:
 483 the implicit version is on average 1.36 times faster than the explicit version
 484 when using one thread, and 2.93 times faster than the explicit version when
 485 using four threads. In Fig. 8(b) we show that the parallel efficiency of the
 486 implicit version is between 73% and 87% efficiency when using four threads,

487 giving a speedup of a factor of 2.9 to 3.5. As in the 2D complex case, the
 488 explicit version does not parallelize well for large problem sizes.

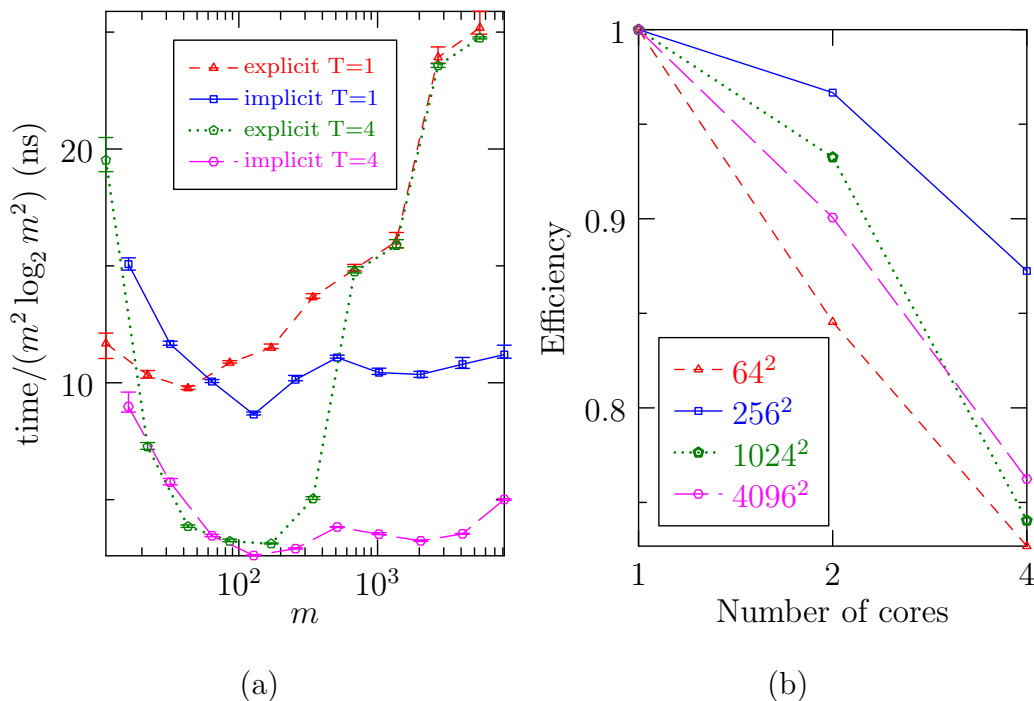


Figure 8: In-place 2D Hermitian convolutions of size $2m \times (m + 1)$: (a) comparison of computation times for explicit and implicit dealiasing using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency of implicit dealiasing versus number of threads. For implicit convolutions, m is chosen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

489 4.3. Complex 3D convolution

490 The decoupling of the 2D work arrays in Function `cconv2` facilitates the
 491 construction of an efficient 3D implicit complex convolution, as described
 492 in Function `cconv3`. For A inputs with dimensions $m_x \times m_y \times m_z$ and
 493 B outputs, the explicit version requires $8Cm_xm_y m_z$ complex words, where
 494 $C = \max\{A, B\}$. In contrast, the implicit version with T threads requires
 495 $2Cm_xm_y m_z + TCm_y m_z + TCm_z$ complex words, approximately one quarter
 496 the storage requirements for the explicit version when $m_x \gg T$. As shown
 497 in Fig. 9(a), implicitly dealiasied convolutions are consequently much faster
 498 than their explicit counterparts. For a single thread, the implicit version

499 is on average 1.9 times as fast as the explicit version and 4.3 times faster
500 on average when comparing execution times over four threads. Fig. 9(b)
501 shows the parallel efficiency of the implicit version, which is between 65%
502 and 86% efficient when using four threads, giving a speedup of a factor of 2.6
503 to 3.4 over one thread. The explicit version has reasonable parallel efficiency
504 for small problem sizes, but this drops to roughly 25% on four threads for
505 problem size $m \geq 64$.

506 4.4. Centered Hermitian 3D convolution

507 As with the 1D and 2D cases, we offer compact and noncompact ver-
508 sions of a 3D Hermitian convolution, and users can choose formats that are
509 compact/noncompact in each direction separately. For serial computations,
510 the best performance typically is achieved when the x and y directions are
511 compact and the z direction is noncompact, so that each dimension is odd,
512 in the interest of cache associativity. However, just as for 2D Hermitian con-
513 volutions, when running on more than one thread, the x direction should be
514 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**
 | $\text{fftpadBackward}(f_a^T[k][j], U_a^T[k][j])$
 parallel for $i = 0$ **to** $m_x - 1$ **do**
 | $\text{cconv2}(\{f_a[i]\}_{a=0}^{A-1})$
 | $\text{cconv2}(\{U_a[i]\}_{a=0}^{A-1})$
 for $b = 0$ **to** $B - 1$ **do**
 | **parallel foreach** $(j, k) \in R$ **do**
 | | $\text{fftpadForward}(f_b^T[k][j], U_b^T[k][j])$
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 $\{U_a\}_{a=0}^{A-1}$.

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**
 | $\text{fft1padBackward}(f_a^T[k][j], U_a^T[k][j])$
 parallel for $i = 0$ **to** $2m_x - 1$ **do**
 | $\text{conv2}(\{f_a[i]\}_{a=0}^{A-1})$
 | $\text{conv2}(\{U_a[i]\}_{a=0}^{A-1})$
 for $b = 0$ **to** $B - 1$ **do**
 | **parallel foreach** $(j, k) \in R$ **do**
 | | $\text{fft1padForward}(f_b^T[k][j], U_b^T[k][j])$
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 $\{U_a\}_{a=0}^{A-1}$.

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

522 In the noncompact format, the memory requirements for an explicit 3D
 523 Hermitian convolution with A inputs and B outputs is $\frac{27}{2}Cm_xm_y(m_z + 1)$
 524 complex words, whereas the implicit version requires only $6Cm_xm_y(m_z +$
 525 $1) + TCm_y(m_z + 1) + TC(\lfloor m_z/2 \rfloor + 1)$ complex words using T threads, where
 526 $C = \max\{A, B\}$. We did not implement a high-performance version of the
 527 explicit routine, so instead we show the execution time of the implicit routine
 528 using one and four threads in Fig. 10 (a). The parallel efficiency is shown in
 529 Fig. 10(b) and ranges between 65% and 92%, which translates to a speedup
 530 of a factor of 2.6 to 3.7 using four threads instead of one.

531 5. Concluding remarks

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

540 Implicitly dealiased convolutions require less memory, are faster, and have
 541 greater parallel efficiency than their explicitly dealiased counterparts. Spe-
 542 cifically, in d dimensions the memory savings for $1/2$ padding is a factor
 543 of 2^{d-1} ; for $2/3$ padding the savings factor is $(3/2)^{d-1}$. The decoupling of
 544 temporary storage and user data means that even in one dimension, users
 545 can save memory by not having to copy their data to a separate buffer. In
 546 higher dimensions, this decoupling allows one to reuse work memory. By
 547 avoiding the need to compute the entire Fourier image at once, one obtains
 548 a dramatic reduction in total memory use. Moreover, the resulting increased
 549 data locality significantly enhances performance, particularly under parallel-
 550 ization. For example, a 3D implicitly dealiased complex convolution runs
 551 about twice as fast as an explicitly dealiased convolution on one thread, and

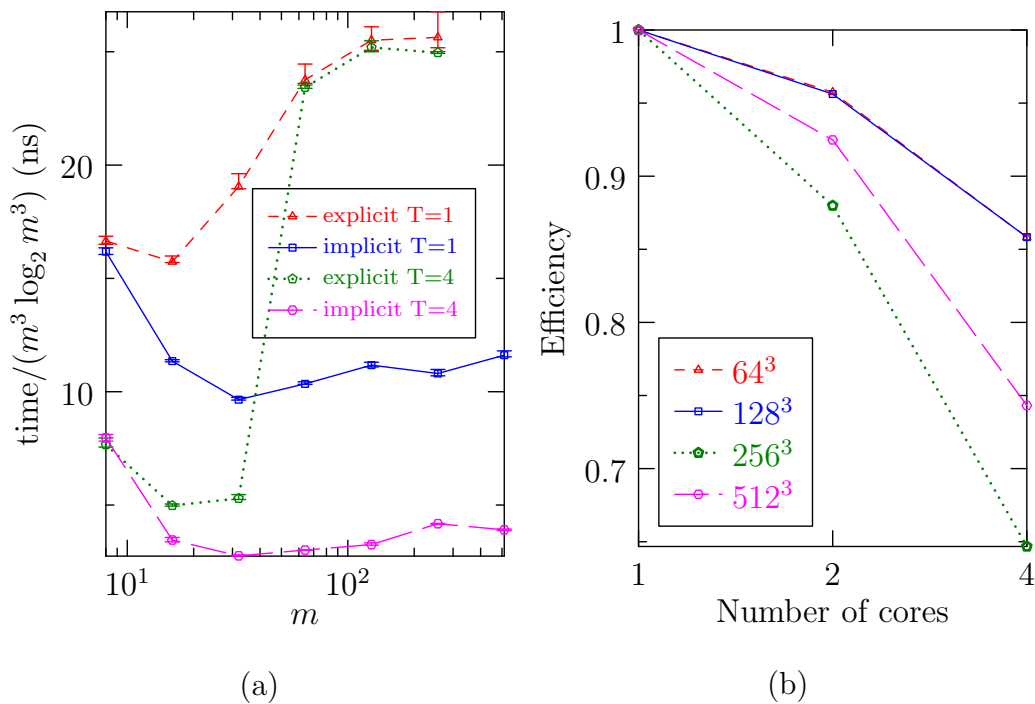


Figure 9: In-place 3D complex convolutions of size $m \times m \times m$: (a) comparison of computation times for explicit and implicit dealiasing using $T = 1$ thread and $T = 4$ threads; (b) parallel efficiency of implicit dealiasing versus number of threads. Here m is chosen to be a power of two.

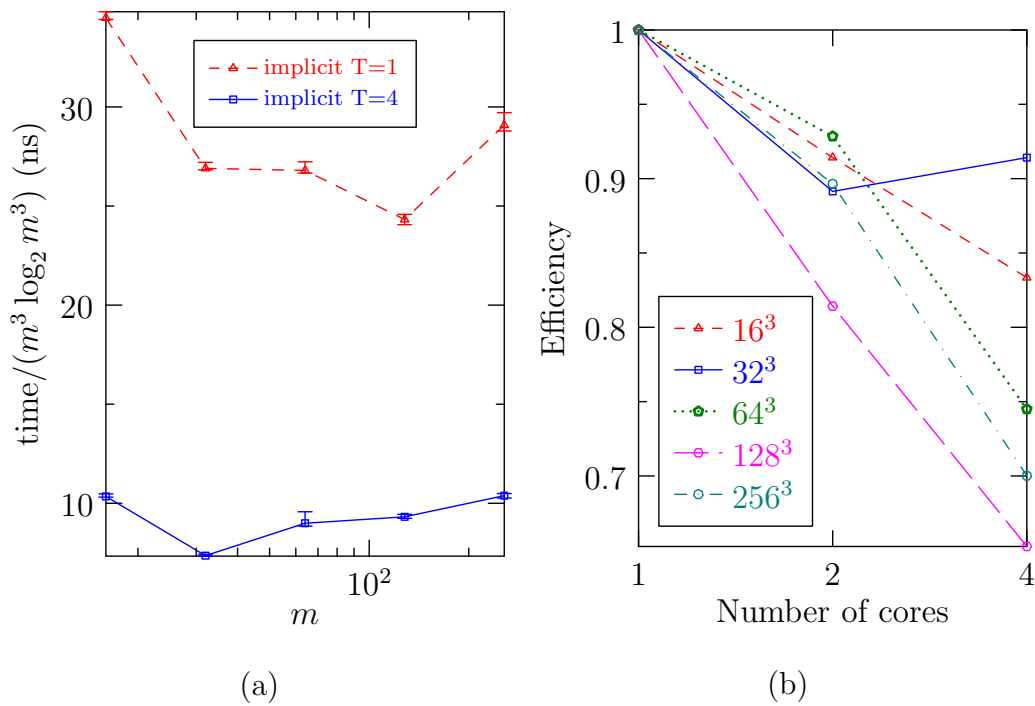


Figure 10: Implicitly dealiased in-place 3D Hermitian convolutions of size $(2m-1) \times (2m-1) \times (m+1)$ for $T=1$ and $2m \times (2m-1) \times (m+1)$ for $T=4$: (a) computation times using $T=1$ thread and $T=4$ threads; (b) parallel efficiency versus number of threads. Here m is chosen to be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

552 over four times faster than the explicit method when both are parallelized
 553 over four threads. For large problem sizes, an implicit complex convolution
 554 requires one-half of the memory needed for a zero-padded convolution in two
 555 dimensions and one-quarter in three dimensions. In the centered Hermitian
 556 case, the memory use in two dimensions is 2/3 of the amount used for an ex-
 557 plicit convolution and 4/9 of the corresponding storage requirement in three
 558 dimensions.

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

565 In future work, we wish to develop specialized implicit convolutions of
 566 real data for applications in signal processing, such as computing cross cor-
 567 relations and autocorrelations of time series. We are also exploring novel ap-
 568 plications of implicitly dealiased convolutions for computing sparse Fourier
 569 transforms [10], fractional phase Fourier (chirp-z) transforms [1], and partial
 570 Fourier transforms [14, 3].

571 **Acknowledgment**

572 Financial support for this work was provided by Discovery Grant RES0020458
 573 from the Natural Sciences and Engineering Research Council of Canada.

574 **AppendixA. Basdevant formulation**

575 *AppendixA.1. 3D incompressible Navier–Stokes equation*

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

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

578 where $D_{ij} = u_i u_j$, requires three backward FFTs to compute the velocity
 579 components from their spectral representations and six forward FFTs of the
 580 independent components of the symmetric tensor D_{ij} , for a total of nine FFTs
 581 per integration stage. However Basdevant [2] showed that this number can

582 be reduced to eight, by subtracting the divergence of the symmetric matrix
 583 $S_{ij} = \delta_{ij} \text{tr } D/3$ from both sides of Eq. (A.1):

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

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

589 *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 in terms of $\tilde{D}_{ij} \doteq D_{ij} - S_{ij}$:

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

on recalling that S is diagonal and $S_{11} = S_{22}$. The scalar vorticity ω thus evolves as

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

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

594 [1] Bailey, D. H., Swarztrauber, P. N., 1991. The fractional Fourier trans-
 595 form and applications. *SIAM review* 33 (3), 389–404.

596 [2] Basdevant, C., 1983. Technical improvements for direct numerical sim-
 597 ulation of homogeneous three-dimensional turbulence. *Journal of Com-
 598 putational Physics* 50 (2), 209–214.

- 599 [3] Bowman, J. C., Ghoggali, Z., 2017. The partial fast Fourier transform.
600 Submitted to J. Sci. Comput.
- 601 [4] Bowman, J. C., Roberts, M., 2011. Efficient dealiased convolutions
602 without padding. SIAM J. Sci. Comput. 33 (1), 386–406.
- 603 [5] Bowman, J. C., Roberts, M., May 6, 2010. FFTW++: A fast Four-
604 vier transform C++ header class for the FFTW3 library. [http://fftwpp.](http://fftwpp.sourceforge.net)
605 [sourceforge.net](http://fftwpp.sourceforge.net).
- 606 [6] Cooley, J. W., Tukey, J. W., April 1965. An algorithm for the ma-
607 chine calculation of complex Fourier series. Mathematics of Computa-
608 tion 19 (90), 297–301.
- 609 [7] Frigo, M., Johnson, S. G., 2005. The design and implementation of
610 FFTW3. Proceedings of the IEEE 93 (2), 216–231.
- 611 [8] Gauss, C. F., 1866. Nachlass: Theoria interpolationis methodo nova
612 tractata. In: Carl Friedrich Gauss Werke. Vol. 3. Königliche Gesellschaft
613 der Wissenschaften, Göttingen, pp. 265–327.
- 614 [9] Gottlieb, D., Orszag, S. A., 1977. Numerical Analysis of Spectral Meth-
615 ods: Theory and Applications. Society for Industrial and Applied Math-
616 ematics, Philadelphia.
- 617 [10] Hassanieh, H., Indyk, P., Katabi, D., Price, E., 2012. Simple and prac-
618 tical algorithm for sparse Fourier transform. In: Proceedings of the
619 Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms.
620 SIAM, pp. 1183–1194.
- 621 [11] Orszag, S. A., 1971. Elimination of aliasing in finite-difference schemes
622 by filtering high-wavenumber components. Journal of the Atmospheric
623 Sciences 28, 1074.
- 624 [12] Roberts, M., 2011. Multispectral reduction of two-dimensional turbu-
625 lence. Ph.D. thesis, University of Alberta, Edmonton, AB, Canada,
626 http://www.math.ualberta.ca/~bowman/group/roberts_phd.pdf.
- 627 [13] Roberts, M., Leroy, M., Morales, J., Bos, W., Schneider, K., 2014. Self-
628 organization of helically forced MHD flow in confined cylindrical geo-
629 metries. Fluid Dynamics Research 46 (6), 061422.
630 URL <http://stacks.iop.org/1873-7005/46/i=6/a=061422>

- ⁶³¹ [14] Ying, L., Fomel, S., 2009. Fast computation of partial Fourier trans-
⁶³² forms. *Multiscale Modeling and Simulation* 8 (1), 110–124.