



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

A Fast Non-Fourier Method for Landau-Fluid Operators

A. M. Dimits, I. Joseph, M. V. Umansky

November 22, 2013

The Physics of Plasmas

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

A Fast Non-fourier Method For Landau-fluid Operators

A.M. Dimits,* Ilon Joseph and Maxim Umansky

Lawrence Livermore National Laboratory

L-637, P.O. Box 808, Livermore, CA 94511-0808

4/7/2014

*phone: 925 422 0211

fax: 925 423 3484

dimits1@llnl.gov

Abstract

An efficient and versatile non-Fourier method for the computation of Landau-fluid (LF) closure operators [G. W. Hammett and F. W. Perkins, *Physical Review Letters* **64**, 3019 (1990)] is presented, based on an approximation by a sum of modified-Helmholtz-equation solves (SMHS) in configuration space. This method can yield fast-Fourier-like scaling of the computational time requirements and also provides a very compact data representation of these operators, even for plasmas with large spatial nonuniformity. As a result, the method can give significant savings compared with direct application of “delocalization kernels” [e.g., G. P. Schurtz, et. al., *Phys. Plasmas* **7**, 4238 (2000)], both in terms of computational cost and memory requirements. The method is of interest for the implementation of Landau-fluid models in situations where the spatial nonuniformity, or particular geometry or boundary conditions render a Fourier implementation difficult or impossible. Systematic procedures have been developed to optimize the resulting operators for accuracy and computational cost. The four-moment Landau-fluid model of Hammett and Perkins has been implemented in the BOUT++ code using the SMHS method for LF closure. Excellent agreement has been obtained for the one-dimensional plasma density response function between driven initial-value calculations using this BOUT++ implementation and matrix eigenvalue calculations using both Fourier and SMHS non-Fourier implementations of the LF closures. The SMHS method also forms the basis for the implementation, which has been carried out in the BOUT++ code, of the parallel and toroidal drift-resonance LF closures. The method is a key enabling tool for the extension of gyro-Landau-fluid (GLF) models [e.g., M. A. Beer and G. W. Hammett, *Phys. Plasmas* **3**, 4046 (1996)] to codes that treat regions with strong profile variation, such as the tokamak edge and scrapeoff-layer.

I. INTRODUCTION

A. Motivation

Fluid-based models, i.e., models in which the dependent variables are fields that depend only on position and time rather than on additional velocity variables [1–5], have contributed much to the understanding of plasma systems, and can be valid under several circumstances. These can be thought of as “adiabatic regimes” where fluid closure is possible due to a wide separation of two dominant time scales. These include the “hydrodynamic” regimes in which the spatial scale length of structures of interest is much greater than the distance by which charged plasma particles (ions or electrons) separate in either a mean collision time (mean free path; i.e., the collisional hydrodynamic regime) [1, 2] or else in the characteristic time of the structures (cold-fluid “adiabatic” regime) [3, 4]. A fluid description can also be valid in the opposite (also sometimes called “adiabatic”) isothermal regime where the particles undergo oscillations (e.g., due to passage through spatially or temporally varying forces or in strong magnetic fields) which are rapid compared with the characteristic times of structures of interest. The fluid descriptions are very useful in that they have a lower inherent dimensionality than kinetic descriptions. This results in greater tractability both in some analytical calculations and more generally in computer simulations, where computational cost is often a limiting factor.

Our particular interest is in the edge and scrapeoff-layer region in magnetic fusion plasma devices such as tokamaks. Fluid simulations with computer codes such as the BOUT++ edge plasma fluid code [6] have been a useful tool for the understanding of transport and turbulence in such plasmas. In the past, edge simulation models have been based largely on collisional fluid equations. However, kinetic effects are known to be important in these regions, and it is highly desirable to extend the models to include kinetic effects. The gyro-Landau-fluid (GLF) approach [8, 9] is a way to incorporate some kinetic effects into fluid simulation codes such as BOUT++, while avoiding the increased dimensionality and associated computational cost of a full nonlinear gyrokinetic model. The “Landau-fluid” (LF) component [7, 10] of the GLF approach is a treatment of kinetic effects, in which the fluid equations are closed with specialized kinetic-fluid closure terms. These closure terms differ from those valid in collisional limits and typically involve nonlocal (integral) operators, but

are nevertheless expressed entirely in terms of the fluid fields that are the dependent variables in the fluid equation set. There is also a separate body of work on nonlocal kinetic closures using “delocalization kernels” [11, 12]. Indeed, the LF approach leads [9], in the three-moment (density, momentum and pressure moment) approximation, to a nonlocal kernel that is of the same form as that given in Ref. [11]. The LF and nonlocality-kernel closures can be thought of as specialized forms of “flux-limiters” in that they limit the spatially dependent component of the fluxes in cases where the collisionality is small or completely negligible. It is the spatially dependent or, more generally, the convergent or divergent parts of the fluxes that lead to changes in the values of the respective moments. The closures have the additional feature of having quantitative accuracy in the collisionless limit for both time-dependent and time-independent phenomena [7, 9, 11]. The LF approach has the advantage that it provides a clear path to closures in additional higher moment equations than those for the pressure or temperature, e.g., in the dynamical heat-flux equation, and this has been found to be essential for improved accuracy in fits of the temporal plasma responses [7–10].

In the simplest case of perturbations from a uniform one-dimensional Maxwellian plasma equilibrium with negligible collisionality, the Landau-fluid closure term can be represented in Fourier (wavenumber) space as a damping term with an effective phase-mixing damping [7, 10]

$$\gamma \propto -|k| v_{\text{th}}, \quad (1)$$

where k is the wavenumber and v_{th} is the mean thermal speed. This form of the damping rate can be arrived at in various ways. Hammett and coworkers introduce such a damping rate based on matching the Landau-damping rate given by a fluid system with a closure term involving this rate with the corresponding kinetic result. This was also derived in Ref. [13] via the adiabatic (low-frequency) limit of exact (but frequency-dependent) closed-form linear kinetic results for the closure terms. One can also arrive at such a rate from general considerations for the linear zero-collision-limit problem. If one assumes that spatial Fourier modes with wavenumber k in a periodic or infinite domain have a temporal linear damping rate that is independent of any collision rate and frequency (i.e., we do not allow the closure term to depend on time derivatives) then from purely dimensional considerations, the only available spatial scale is $1/k$ and the only available velocity scale is v_{th} . From symmetry

considerations, lack of dependence of the damping rate on the frequency dictates that the damping rate is independent of the sign of k . It follows that the damping coefficient of the form in Eq. (1) is the only possible one.

Collisional effects have been added in such treatments through a generalization of the above damping rate. At the level of a 4th-moment closure (closure of the dynamic heat flux equation by expressing the 4th velocity moment in terms of the heat flux) [9], Eq. (1) generalizes to

$$\gamma \propto -(|k|v_{\text{th}} + \nu_s).$$

Here, ν_s is an effective collision rate, so that ν_s/v_{th} is an effective inverse mean free path. By taking the adiabatic limit of the resulting dynamic heat flux equation, a closure for the heat flux itself in terms of the perturbed temperature can be obtained, which involves a damping rate of the form

$$\gamma \propto -\frac{k^2 v_{\text{th}}}{|k| + \nu_s/v_{\text{th}}}$$

at the level of a heat-flux closure [9, 11]. We note that the four-moment LF models, i.e., including a dynamical heat-flux equation, are of value both for improved accuracy in fits of the temporal plasma responses [7, 9], and because they provide transparency concerning the inclusion of additional physics such as the effects of a fluid flow velocity.

Specific results include the Hammet-Perkins expression for the divergence of the parallel conductive heat flux (parallel conduction of parallel internal energy)

$$\frac{\partial}{\partial z} Q_z(z) \approx -\sqrt{8/\pi} n_0 v_{\text{th}} \int dk e^{ikz} |k| \tilde{T}(k), \quad (2)$$

where n_0 is the number density. This can be expressed as a configuration-space operation on the temperature field

$$Q_z(z) = \frac{\partial}{\partial z} Q_z(z) \approx \text{P} \int_{-\infty}^{\infty} dz' G(z-z') \tilde{T}(z'),$$

$$G(z) = \left(\frac{2}{\pi}\right)^{3/2} n_0 v_{\text{th}} \frac{1}{z}.$$

A specific form of the closure of the heat flux in the energy equation was given [9] as

$$\frac{\partial}{\partial z} Q_z(z) \approx -8n_0 v_{\text{th}}^2 \int dk e^{ikz} \frac{k^2 \tilde{T}(k)}{\sqrt{8\pi} |k| v_{\text{th}} + (3\pi - 8) \nu_s}. \quad (3)$$

A very similar model was obtained empirically via a fit to Fourier-mode damping rates obtained from numerical Fokker-Planck-equation calculations by Epperlein and Short [11].

The above expressions make sense in the case where the background is uniform in the direction of the wavenumber. They then lead to straightforward and efficient evaluation in Fourier space of the application of the operators to a perturbation field. However, when strong spatial inhomogeneity is present, the locality of the operators in the Fourier space breaks down, and the Fourier treatment becomes much less useful. Such is the case in magnetic fusion edge plasmas, for which the BOUT++ code was developed, and also for the electrons in inertial confinement fusion capsules [14]. These plasmas can have large static or quasistatic spatial inhomogeneities in the moment quantities such as the density, temperature and flow velocity, and hence in the collisionality regime, as well as large relative fluctuation amplitudes. The particular discretizations, geometry and boundary conditions in simulation codes such as BOUT++ can themselves also make Fourier implementations difficult or impossible. In such cases a non-Fourier implementation of the LF closure operators is needed.

With spatial mesh-based schemes (finite difference, volume, element, etc.), rapid evaluation of operators of the form $k^n \leftrightarrow (-i\partial/\partial z)^n$ is possible because spatial discretizations (including high order accurate discretizations) of integer powers of the derivative with quite local kernels are available. However, for non-integer powers and non-analytic functions of the wavenumber, for example the operators $|k| \times k^n = k^{n+1} \text{sgn}(k)$ of interest here, the corresponding spatial kernels are intrinsically nonlocal, so that the straightforward evaluation of the application of these operators (through convolution or matrix multiplication involving the discretized kernel) results in an unfavorable scaling of the computational effort. Specifically, for a one-dimensional finite domain discretized with (a large integer) N_g mesh points, the computational effort for a direct convolution or matrix multiplication is expected to scale approximately as N_g^2 , versus $N_g \log(N_g)$ for the combination of discrete Fourier transformation and application of the operator in Fourier space. Furthermore, the discretized kernels themselves can be unwieldy in cases of interest. For a spatially inhomogeneous situation in

one dimension, the number of elements in the kernel matrix itself scales as N_g^2 . In a 2- or 3-dimensional magnetized plasma calculation, the LF operators themselves are still basically one dimensional. The parallel LF closure operator operates in the direction along the (possibly perturbed) magnetic field and the toroidal drift-resonance closures involve operations along streamlines of the combined ∇B and curvature drifts. In these cases the scaling of the number of matrix elements that must be calculated and/or stored is as $N_{\text{gtot}} \times N_{\text{gop}}$, where N_{gtot} is total number of grid cells in the system and N_{gop} is the number of grid cells in the direction in which the operator is applied, i.e., the parallel direction for the parallel LF operator and the minor-radius direction in the case of the toroidal drift-resonance operator. In a general (e.g., unmagnetized) spatially inhomogeneous multi-dimensional situation, the number of matrix elements can scale as N_{gtot}^2 or, if the kernel can be limited to some spatial scale such as a collisional mean free path, as $N_{\text{gtot}} \times N_{\text{gkernel}}$ where N_{gkernel} is a typical number of grid cells in the support of the kernel matrix. It is therefore of interest to look for configuration-space based methods that are accurate, yet have reduced computational (time and storage) costs relative to direct matrix multiplication.

We have developed an efficient and versatile non-Fourier method for the computation of the LF closure operators. This method is based on an approximation by a sum of modified-Helmholtz-equation solves in configuration space, and we will therefore refer to it here as “the SMHS method.” An brief preliminary report on this method, including some early results, was given in Ref. [15]. Here we present the method in greater detail, including systematic procedures for its construction and for the quantification and optimization of its accuracy and computational efficiency. The SMHS method uses the nonlocality of the kernel resulting from the inversion of (i.e., the Greens function for) a modified Helmholtz operator to efficiently capture the nonlocality in space of the $k^{n+1} \text{sgn}(k)$ operator. This can result in fast-Fourier-like scaling of the computational time requirements and also provides a very compact data representation of such operators, even for plasmas with significant spatial nonuniformity. As a result, the SMHS method can give significant savings compared with direct application of delocalization kernels, both in terms of computational cost and memory requirements.

The one-dimensional LF closure operator has been implemented in the BOUT++ code using the SMHS method, and in separate stand-alone verification codes using the SMHS method, the Fourier method as well as direct convolution and matrix multiplication. Em-

pirical comparisons of computational scaling between these methods are given. Excellent agreement has been obtained for the one-dimensional plasma density response function in the four-moment LF model of Hammett and Perkins between driven initial-value calculations and matrix-eigenvalue calculations using the SMHS and Fourier implementations of the closure terms.

The SMHS method also forms the basis for the implementations, which have been carried out in the BOUT++ code, of the parallel LF closure and the toroidal drift-resonance LF closure. The method is a key enabling tool for the extension of gyro-Landau-fluid (GLF) models [9] to codes that treat regions with strong profile variation, such as the tokamak edge and scrapeoff-layer.

The organization of the remainder of the paper is as follows. In subsection II A, the basic method and approximations are introduced. The computational performance of the method is compared for the simple, one-dimensional, case in subsection II B. A systematic procedure for optimizing the accuracy of the approximation of the LF operators by a sum of Lorentzians or modified Helmholtz solves is given in subsection II C. Subsection II D discusses optimizing the SMHS method for truncation error (finite-grid effects) and system-size effects. Subsection II E reports verification tests for the SMHS implementation using the one-dimensional plasma response function are discussed in Subsection II E. Subsection II F briefly summarizes the approaches that generalize the one-dimensional case to the closure terms that arise in the GLF models. Finally, a summary of the work and discussion of some further generalizations of interest and future directions are given in Section III.

II. APPROXIMATION BY A SUM OF MODIFIED-HELMHOLTZ-EQUATION SOLVES

A. The basic approximation method

Our method is based on a sum of Lorentzians in wavenumber (k) space and their representation as the solution of an inhomogeneous modified Helmholtz equation. Consider a function $\psi(z)$ (where z is for now a single real spatial variable) which satisfies the

inhomogeneous modified Helmholtz equation with source function $S(z)$:

$$\left(1 - \frac{\partial^2}{\partial z^2}\right) \psi(z) = S(z). \quad (4)$$

Eq. (4) can be discretized and solved by an appropriate elliptic-equation solver. In one dimension, the simplest discretization of the second derivative for a uniform mesh in one dimension is a 3-point second difference

$$\frac{\partial^2 \psi}{\partial z^2} \rightarrow \frac{1}{\Delta^2} (\psi_{i+1} + \psi_{i-1} - 2\psi_i), \quad (5)$$

which results in a tridiagonal matrix (or a generalization appropriate to the particular boundary conditions). For the one dimensional case and typical simulation mesh sizes, the resulting condition numbers of the matrices to be inverted are typically modest, and a direct banded solver can be used to solve the resulting matrix equation. The computational time for such a numerical solution scales directly in proportion to the number of mesh nodes N_g . This numerical solution can be thought of as an application of an integral operator to $S(z)$ which is nonlocal, but in Fourier space is approximately a multiplication by $1/(1+k^2)$ where k is the wavenumber. In other words, the operation of multiplying the Fourier coefficients of the field S by $1/(1+k^2)$ is the equivalent of applying the inverse of a modified Helmholtz operator, i.e.,

$$\left(1 - \frac{d^2}{dz^2}\right) \frac{1}{2\pi} \int_{-\infty}^{\infty} dk \frac{1}{1+k^2} \hat{S}_k e^{ikz} = \frac{1}{2\pi} \int_{-\infty}^{\infty} dk \hat{S}_k e^{ikz}$$

or, formally using the Fourier theorem,

$$\left(1 - \frac{d^2}{dz^2}\right)^{-1} S(z) = \frac{1}{2\pi} \int_{-\infty}^{\infty} dk \frac{1}{1+k^2} \hat{S}_k e^{ikz},$$

where

$$\hat{S}_k = \int_{-\infty}^{\infty} dz S(z) e^{-ikz}.$$

Thus, the numerical modified-Helmholtz-equation solution is a way to apply an operator which is highly nonlocal in z with a “fast” computational scaling. In contrast, the computational cost of applying the same operator directly in z configuration space via a direct

numerical convolution or matrix multiplication scales as N_g^2 .

For higher-order compact finite-difference or finite-element discretizations, the matrix associated with the discretization generalizes to a narrowly banded matrix (or a generalization appropriate to the particular boundary conditions), for which the time taken for the solution of the resulting equation by an efficient matrix equation solver still scales directly in proportion to the number of mesh nodes. In the two-dimensional case, of interest for the toroidal drift-resonance Landau-fluid operators and the gyrofluid operators [9] that act perpendicular to the magnetic field, the associated matrix becomes a more general but still highly sparse matrix. In some such cases (e.g., for the toroidal drift-resonance closures when there is significant profile variation, but the perturbations have small amplitude), the matrix problem can be reduced to one dimension via a Fourier transform in the toroidal (binormal) direction. (See sec. IIF.) In more general cases, the representation through modified Helmholtz equations will still have the very significant advantage of sparseness. A more general but very efficient linear solution method, such as the algebraic multigrid method, will be required to preserve the fast-Fourier-like scaling of the computational time and overall usefulness of the SMHS method.

The second ingredient of our method is the approximation of $1/|k|$ as a sum of suitably scaled Lorentzians

$$\frac{1}{|k|} \approx \sum_n \frac{\beta_n}{k^2 + \alpha^{2n}},$$

where for now the range of n in the sum is not specified. We can compare the Lorentzian function $\beta/(k^2 + \alpha^2)$ and $1/|k|$. The Lorentzian is even in k , as is $1/|k|$. The values and gradients of these two functions agree at $k = \alpha$ if we set $\beta = 2\alpha$. This comparison is shown for $\alpha = 1.0$, i.e.,

$$\psi_1(k) = \frac{2}{1 + k^2},$$

in Fig. (1a). By adding a second Lorentzian with a different width and suitably scaled coefficients, the spectral range of good agreement (i.e., agreement to within some value of the relative or absolute error) can be greatly increased. This is seen in Fig. (1b), which shows a curve for the function

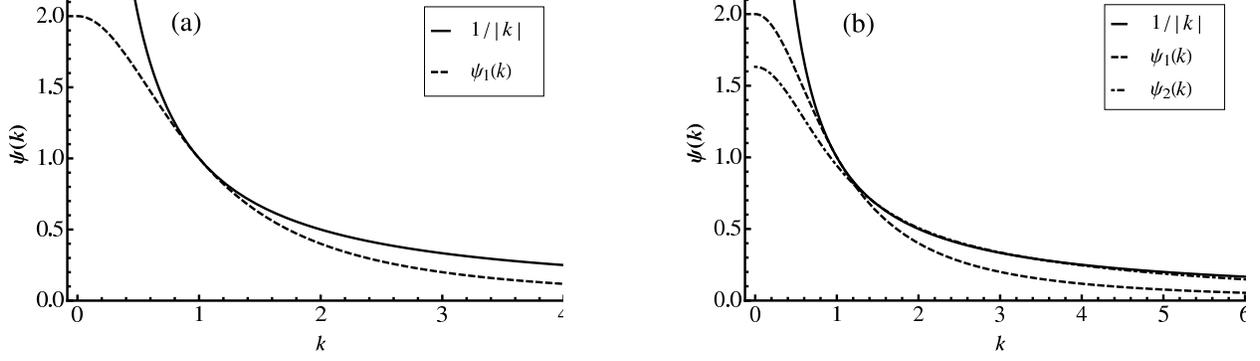


Figure 1: Plot of $1/|k|$ compared (a) with a single Lorentzian and (b) with a single Lorentzian and a sum of two Lorentzians.

$$\psi_2(k) = 0.68 \left(\frac{2}{1+k^2} + \frac{10}{25+k^2} \right),$$

in addition to the same curves as in Fig. (1a). Additional Lorentzians can be added to obtain good agreement over successively wider ranges of the wavenumber k .

Now consider the (bidirectionally infinite) series

$$\psi_\infty(k, \alpha) \approx \sum_{n=-\infty}^{\infty} \frac{\alpha^n}{k^2 + \alpha^{2n}},$$

where α is any real number greater than 1. This series converges pointwise for any real $k \neq 0$. This can easily be seen from the fact that for fixed k , the terms in the sum can be bounded by

$$\frac{\alpha^n}{k^2 + \alpha^{2n}} < \begin{cases} \alpha^{-n} & \text{as } n \rightarrow \infty \\ \alpha^n/k^2 & \text{as } n \rightarrow -\infty. \end{cases}$$

The sum $\psi_\infty(k, \alpha)$ satisfies

$$\begin{aligned} \psi_\infty(-k, \alpha) &= \psi_\infty(k, \alpha), \\ \psi_\infty(\alpha k, \alpha) &= \psi_\infty(k, \alpha)/\alpha. \end{aligned}$$

Thus a reasonable starting ansatz for fits to $1/|k|$ with sums of Lorentzians is a finite

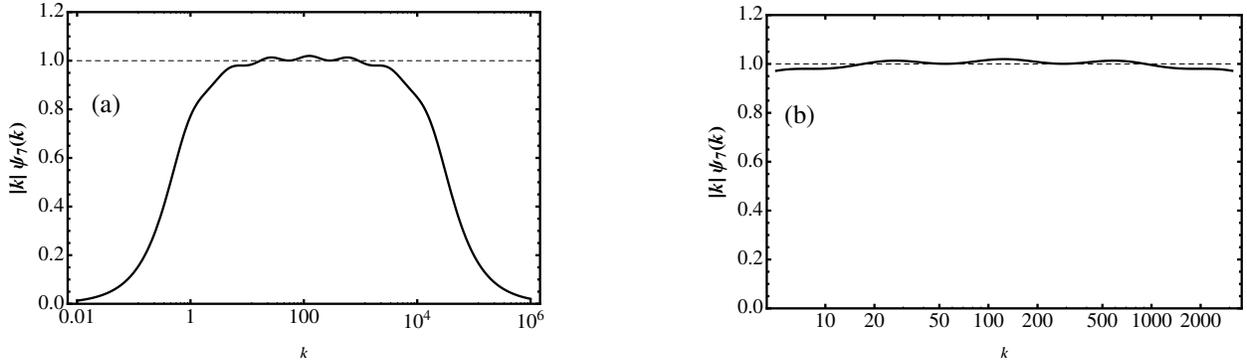


Figure 2: Plots of $|k|\psi_+(k; \alpha, \beta, N)$ for $\alpha = 5, \beta = 1.04, N = 7$. Frame (b) zooms in on the region of good fit of $\psi_+(k; \alpha, \beta, N)$ to $1/|k|$.

truncation of $\psi_\infty(k, \alpha)$, e.g.,

$$\psi_{N+}(k; \alpha, \beta, N) = \beta \sum_{n=0}^{N-1} \frac{\alpha^n}{k^2 + \alpha^{2n}}, \quad (6)$$

where N is a positive integer (which hopefully need not be large) and β is an overall multiplier, which for fitting purposes may be taken to be a function of α and N . Heuristically, $\psi_+(k; \alpha, \beta, N)$ can be expected to be a good fit over a wavenumber range delineated roughly by $\gamma \lesssim |k| \lesssim \alpha^{N-1}/\gamma$ where γ is a positive real number of order 1. This is indeed the case, as can be seen from plots of the ratio of ψ_+ to $1/|k|$, i.e., of

$$|k|\psi_+(k; \alpha, \beta, N) = \beta |k| \sum_{n=0}^{N-1} \frac{\alpha^n}{k^2 + \alpha^{2n}},$$

which are shown in Fig. 2.

This figure shows that $\psi_{N+}(k; \alpha, \beta, N)$ with 7 terms can approximate $1/|k|$ to within a relative error of 2% over a spectral range of approximately 10^3 .

Here we briefly comment on some favorable aspects of the SMHS approach. The approximation by a sum of Lorentzians or modified Helmholtz solves is applicable to cases where there is spatial inhomogeneity in the operator, for example through spatial variation in the thermal velocity or thermally averaged drift velocities in a toroidal implementation. This is not true of the Fourier or convolution methods. The matrix multiplication method is applicable, but requires calculation and storage of the kernel matrix, the size of which can be much larger than the computational mesh. The closure operator may in general be most

easily represented as a combination of local (e.g. differential) operators, quantities that are functions of position, and nonlocal operators. The SMHS method provides the flexibility to implement some of the spatial dependence through spatially dependent coefficients in the modified Helmholtz equation. Another consideration that favors the SMHS method over direct methods is the fact that in a toroidal application, e.g., a typical flux-tube based simulation geometry with offset-periodic boundary conditions in the direction parallel to the magnetic field, the perturbations may extend several poloidal circuits in the parallel direction. For such a perturbation the nonlocal operator should have good fidelity for these very long spatial scales (or low wavenumbers) and the effective value of N_g may be considerably larger than the number of cells in the parallel direction in the simulation domain.

B. Computational performance

We have implemented the SMHS method in a one-dimensional test code on a periodic domain. A comparison was made of computational time for the application of the operator by various methods, including the SMHS method, using a single process on a compute node of the NERSC Edison computer [16]. The timings are for 1000 applications of the operator in a loop of Fortran code. The other methods compared include the Fourier method, coded in Fortran using the FFTW3 library from within the Cray LibSci library, and matrix multiplication by the discretized kernel using the BLAS DGEMV routine from LibSci and the gfortran intrinsic MATMUL routine. The matrix multiplication is a realization of the direct application of a delocalization kernel [12]. The results are shown in Fig. 3. It is seen that the SMHS method has similar computational scaling to the Fourier method, approximately as $N_g \log(N_g)$, where N_g is the number of grid cells, and is uniformly a factor of 2 to 3 more expensive than the Fourier method over most of the range of N_g . In contrast, the timing for matrix multiplication scales as approximately as N_g^2 for the MATMUL implementation and approximately as $N_g^{2.5}$ (in the range $128 \leq N_g \leq 4096$) for DGEMV. Note that DGEMV is faster than MATMUL across the whole range of N_g values shown in Fig. 3, despite its apparent scaling as $N_g^{2.5}$ being less favorable than the N_g^2 scaling of MATMUL. The crossover point, at which the timing of the SMHS and direct methods is comparable is at $N_g \approx 128 - 512$. The SMHS method has an advantage in terms of computational-time cost for values of N_g larger than this.

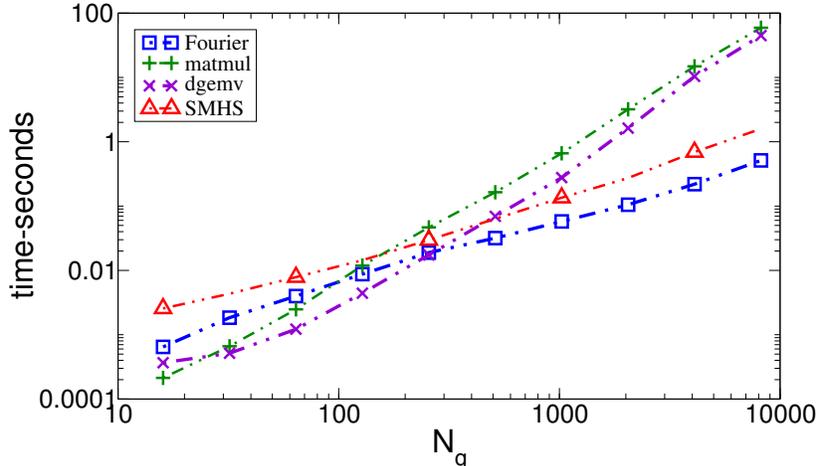


Figure 3: Timings for the application of the $|k|$ operator by various methods, as indicated in the legend.

Note, again, that the computational cost of the SMHS method in a particular setting will depend on the efficiency of the underlying modified Helmholtz solver used and on the boundary conditions. In multidimensional cases, we expect that this method can still be efficient and give the fast-Fourier-like computational time scaling provided that the associated modified Helmholtz equations are solved via a very efficient linear solution method, such as the algebraic multigrid method. Similarly, we observe that there are variations in the efficiency of the direct methods, depending on the particular matrix multiplication algorithm used.

C. Optimization of accuracy and spectral range

Improved fits to $1/|k|$ can be obtained by generalizing the coefficients in the sum in Eq. (6) to give

$$\tilde{\psi}_+(k; \alpha, \vec{\beta}, N) = \sum_{n=0}^{N-1} \frac{\alpha^n \beta_n}{k^2 + \alpha^{2n}}. \quad (7)$$

Here $\vec{\beta}$ is a N -dimensional real vector with components β_n , which can be determined by a systematic fitting procedure. One such procedure is a well-posed collocation calculation. By “well-posed” is meant a calculation involving a set of optimization constraints whose number exactly matches the number of free parameters in the fit. The collocation constraints used are exact matches of $\tilde{\psi}_+(k; \alpha, \vec{\beta}, N)$ with $1/|k|$ at a set of $|k|$ values most of which are evenly

N	2	3	4	5	6	7
range	7	20	100	500	2.5×10^3	1.3×10^4

Table I: Range of k for good fit of $\tilde{\psi}_+(k; \alpha, \vec{\beta}, N)$ to $1/|k|$.

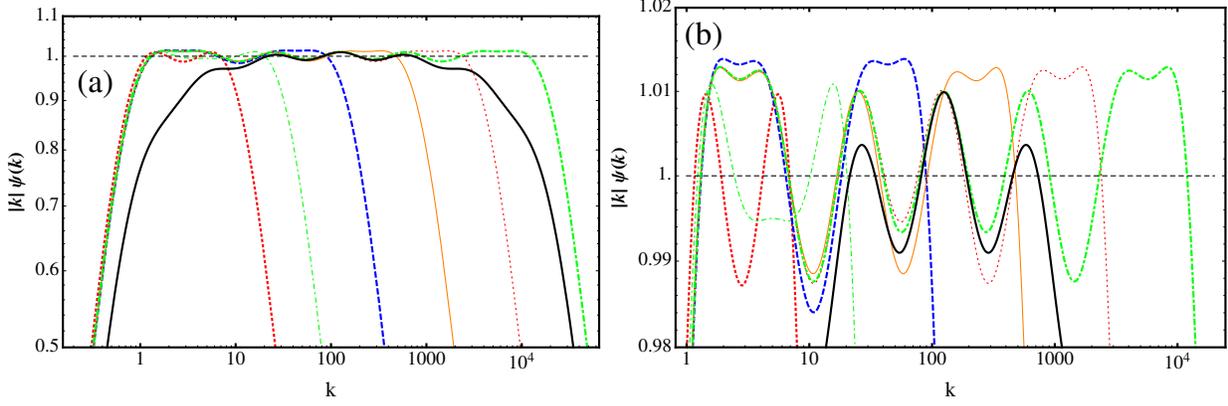


Figure 4: Results from the collocation-based optimized fits $\tilde{\psi}_+(k; \alpha, \vec{\beta}, N)$ to $1/|k|$ for different values of N . The plots show $|k| \tilde{\psi}_+(k; \alpha, \vec{\beta}, N)$ versus $|k|$ with different vertical scales for the two frames. Proceeding from narrower to wider range of fit, the curves are for $N = 2, 3, 4, 5, 6, 7$. The solid heavy black curve is the result from the truncation fit $|k| \psi_+(k; \alpha, \beta, N)$ for $N = 7$.

spaced in the logarithm of $|k|$. Specifically, we define $\kappa_n = \alpha^{n-1} \kappa_0$, $n = 0, 1, 2, \dots, N - 1$. One good choice for the collocation points is at $k = k_{\text{coll},n}$, with $k_{\text{coll},n} = \kappa_n$ for $n = 1, 3, \dots, N - 2$, and $k_{\text{coll},0} = \kappa_0/\eta$, $k_{\text{coll},N-1} = \eta \alpha^{N-1} \kappa_0$. With this choice, $\vec{\beta}$ is obtained as a solution of a (linear) matrix equation of the form $\mathbf{A} \vec{\beta} = \mathbf{R}$, where \mathbf{A} is a $N \times N$ matrix, and \mathbf{R} is a N -dimensional vector. Results from the collocation-based optimization are shown in Fig. (4). For a given N , the optimized fits give an increase in the spectral range of good fit of order 100 over the simple truncation of Eq. (6).

We note also that the curves in Figs. (2) and (4) display a symmetry about $\alpha^{(N-1)/2}$. This symmetry can be shown to hold provided $\beta_{N-1-n} = \beta_n$, which is the case for the values of β_n obtained by the above collocation procedure. Formally, this symmetry can be expressed as

$$k \tilde{\psi}_+(\alpha^{(N-1)/2} k; \alpha, \vec{\beta}, N) = \frac{1}{k} \tilde{\psi}_+\left(\frac{\alpha^{(N-1)/2}}{k}; \alpha, \vec{\beta}, N\right). \quad (8)$$

For the collisional case, e.g., for the operator in Eq. (3), if v_{th} and ν_s are uniform, then by a normalization of the spatial scale (to $\sqrt{8\pi} v_{\text{th}} / [(3\pi - 8) \nu_s]$), the basic nonlocal operator in can be represented as $1/(|k| + 1)$. Various choices are possible for the summation limits

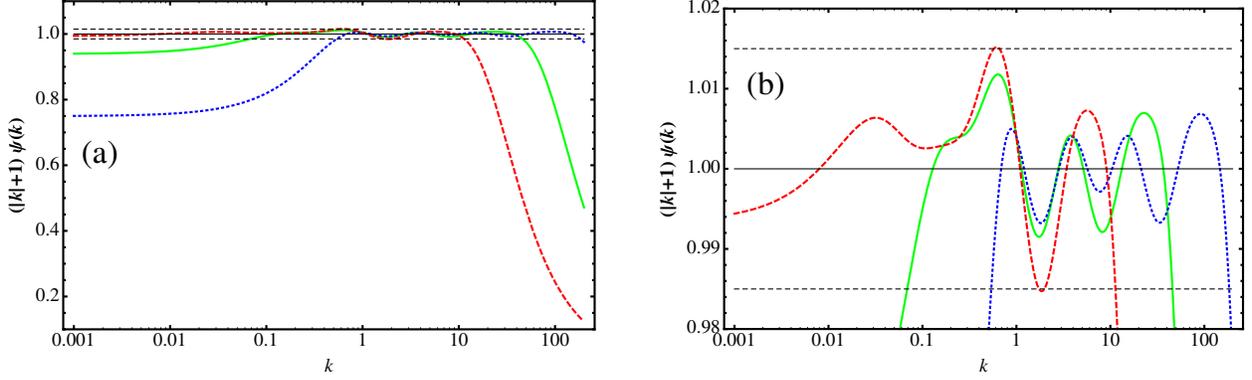


Figure 5: Results from the collocation-based optimized fits $\tilde{\psi}(k; \alpha, \vec{\beta}, M, N)$ to $1/(1+|k|)$ for $\alpha = 3$, $N + M + 1 = 5$. The plots show $(1+|k|)\tilde{\psi}(k; \alpha, \vec{\beta}, M, N)$ versus $1+|k|$ with different vertical scales for the two frames. Proceeding from narrower to wider range of fit, the curves are for $M = -2, -1, 0$. Frame (b) magnifies the vertical axis to show the relative accuracy of the fits more quantitatively.

and collocation points in a SMHS approximation of this operator. One example, which gives a good fit (to less than 1.5% relative error) over a range of normalized wavenumbers of approximately 300 irrespective of whether that wavenumber range is large or small compared with 1 (the inverse mean free path) or straddles the transition between the two regimes, is

$$\frac{1}{1+|k|} \approx \tilde{\psi}(k; \alpha, \vec{\beta}, M, N) = \sum_{n=-M}^N \frac{\alpha^n \beta_n}{k^2 + \alpha^{2n}}, \quad (9)$$

with $\alpha = 3$, $N + M + 1 = 5$. The collocation points chosen as $k_n = \alpha^n$, $n = -M, \dots, N-1$, $k_N = \eta\alpha^N$, and the collocation is now done between $\gamma/(1+|k|)$ and $\tilde{\psi}(k; \alpha, \vec{\beta}, M, N)$. Here γ is an additional free parameter (we choose $\gamma = 1.004$), which helps to center the relative error about zero over a wider spectral range. For different M values, with $N = 4 - M$, the range of fit can be chosen to encompass high collisionality ($|k| \ll 1$, using $M = 2$), low collisionality ($|k| \gg 1$, using $M \leq -1$), or to encompass a collisionality range extending from low collisionality to moderate collisionality (using $M = 0$ or 1).

The series approximations in Eqs. (6)-(9) are specific cases of the more general form

$$\psi_A(k; \alpha, \vec{\beta}, N) = \sum_{n=N_l}^{N_r} \frac{\alpha^n \beta_n}{k^2 + \alpha^{2n}}.$$

We have also implemented a least-squares based procedure, which will be described else-

where.

D. Matching between the spectral range of good approximation and that of the simulation mesh

In a practical implementation of the SMHS method, the effects of the spatial discretization and the system size can be taken into account in the optimization of the approximation.

Associated with the simulation mesh size Δ is the Nyquist wavenumber $k_{\text{Nyq}} = \pi/\Delta$, and the largest wavenumber that is represented with some acceptable level of accuracy by the discretization $k_{\text{dscr}} = \eta k_{\text{Nyq}}$, where η is a number less than 1. The negative of the eigenvalues of the centered 3- and 5-point difference discretizations of the second derivative on a uniform mesh with mesh size Δ , Eq.(5), are

$$\begin{aligned} K_2^{(2)}(k\Delta) &= [2 \sin(k\Delta/2)/\Delta]^2, \\ K_4^{(2)}(k\Delta) &= \frac{4}{3} [2 \sin(k\Delta/2)/\Delta]^2 - \frac{1}{3} [\sin(k\Delta)/\Delta]^2, \end{aligned}$$

where k is the wavenumber. These are compared with the exact value k^2 in Fig. (6). If, for example, one defines k_{dscr} as the wavenumber below which the relative error in a discretization is less than 2%, then the values of $k_{\text{dscr}}\Delta$ for these 3- and 5-point discretizations are respectively 0.45 and 1.2. The corresponding values of η are then respectively 0.14 and 0.38.

There is also generally a longest wavelength or spatial length scale L that needs to be represented, and an associated wavenumber $k_L = 2\pi/L$. L may be the system size or, in the case of the parallel operator on an offset-periodic mesh as arises when using magnetic-field-aligned coordinates, may be many times the actual domain size. Even in this case, the value of N required in Eq. (7) can still be quite modest.

Recall that $\psi_{N+}(k; \alpha, \beta, N)$ or $\tilde{\psi}_{N+}(k; \alpha, \vec{\beta}, N)$ can be expected to be a good fit over a wavenumber range delineated roughly by $\gamma \lesssim |k| \lesssim \alpha^{N-1}/\gamma$ where γ is a positive real number of order 1 [see also Eq.(8)]. In application to a particular simulation mesh, this region of good fit must be normalized and “positioned” in k space to encompass the range of wavenumbers (e.g., $k_L \lesssim k \lesssim k_{\text{dscr}}$) for which a good fit is needed. This can be accomplished via a simple scaling of the wavenumber and overall multiplier, as shown in Fig. 7.

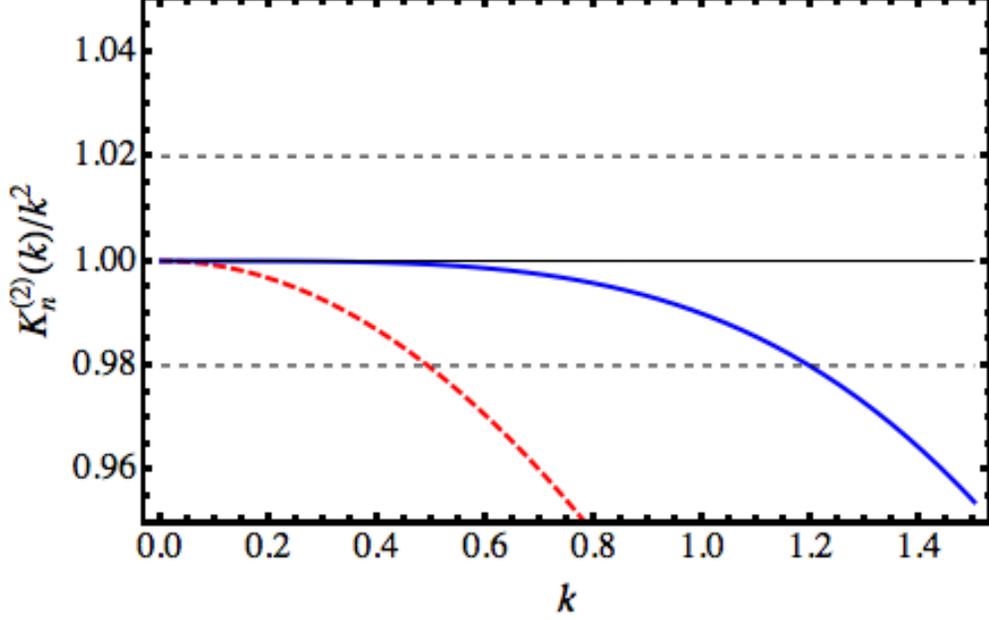


Figure 6: Curves show $K_2^{(2)}(k)/k^2$ (red) and $K_4^{(2)}(k)/k^2$ (blue) for comparison of 3-point and 5-point approximations of $-\partial^2/\partial x^2$ to k^2 , for $\Delta = 1$.

If $\tilde{\psi}_{N+} \left(k; \alpha, \vec{\beta}, N \right)$ approximates $1/|k|$ adequately for $\gamma \lesssim |k| \lesssim \alpha^{N-1}/\gamma$ according to some criterion (e.g., to within a specified relative error), then the function

$$\frac{1}{k_{\text{sh}}} \tilde{\psi}_{N+} \left(\frac{k}{k_{\text{sh}}}; \alpha, \vec{\beta}, N \right)$$

approximates $1/|k|$ adequately for $\gamma k_{\text{sh}} \lesssim |k| \lesssim \alpha^{N-1} k_{\text{sh}}/\gamma$. Choosing $k_{\text{sh}} = k_L \gamma$, which corresponds to moving the lower wavenumber boundary of the range of good fit to coincide with k_L , gives that

$$\frac{1}{\gamma k_L} \tilde{\psi}_{N+} \left(\frac{k}{\gamma k_L}; \alpha, \vec{\beta}, N \right)$$

approximates $1/|k|$ adequately for $k_L \lesssim |k| \lesssim \alpha^{N-1} k_L/\gamma^2$. If $\alpha^{N-1} k_L/\gamma^2 \geq k_{\text{dscr}}$, then the range of good fit encompasses the entire spectral range of the simulation. Alternatively, one can make the upper boundary of the range of good fit coincide with k_{dscr} by choosing $k_{\text{sh}} = \gamma k_{\text{dscr}}/\alpha^{N-1}$.

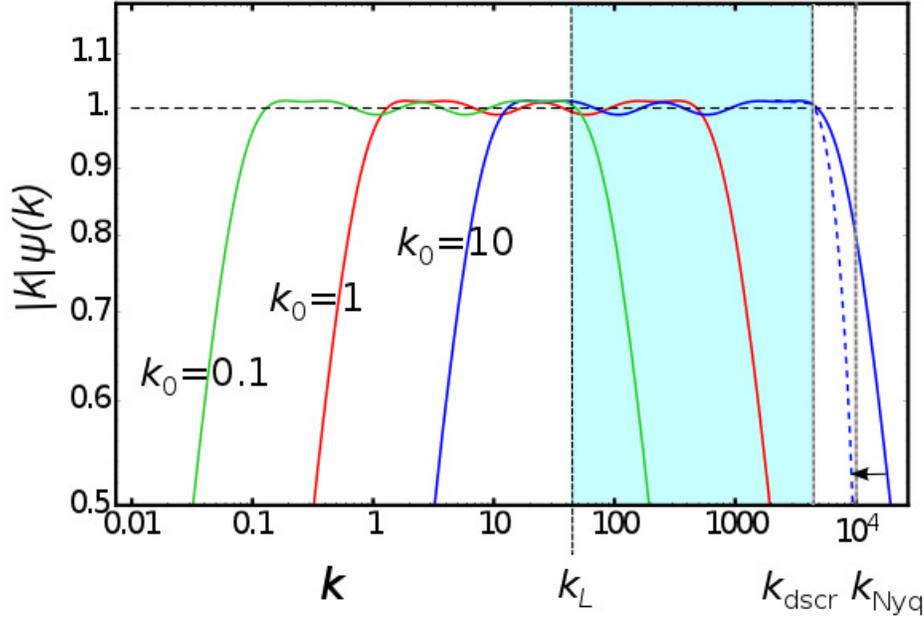


Figure 7: Shifting the spectral range of good fit of the sum-of-Lorentzians approximation to $1/|k|$. Here, the function $\tilde{\psi}_{N+}(k; 5, \vec{\beta}, 5)$ with a suitable choice of $\vec{\beta}$ is shown, which approximates $1/|k|$ to within a relative error of less than 1.5% for $1.2 \lesssim |k| \lesssim 5^4/1.2$. This can be made to fit over the wavenumber range $12 \lesssim |k| \lesssim 5^4 \times 10/1.2$ by setting $k_{\text{sh}} = 10$. Shown also is an example of a simulation wavenumber range with $k_L = 45$, $k_{\text{Nyq}} = 10^4$, and $k_{\text{dscr}} \approx 4 \times 10^3$ which is encompassed by the wavenumber range of good fit of the shifted approximating function $\frac{1}{10}\tilde{\psi}_{N+}(10k; 5, \vec{\beta}, 5)$. The two curve segments connected by the arrow show schematically the effect of truncation error on the fit.

E. Verification through computation of the plasma response function

The one-dimensional LF closure operator has been implemented in the BOUT++ code using the SMHS method, and in separate stand-alone initial-value and matrix-eigenvalue verification codes using the SMHS method and the Fourier method. We have used these implementations to verify the applicability and accuracy of the SMHS method for the linear one-dimensional collisionless Landau response in the four-moment Landau-fluid model of Hammett and Perkins [7]. The (density) response function $-R(k, \omega) = \tilde{n}(k, \omega)T_0/n_0q\tilde{\phi}(k, \omega)$ for the four-moment LF model, where \tilde{n} is the perturbed density, T_0 is the temperature, q is the charge, $\tilde{\phi}$ is the perturbed electrostatic potential, and ω is the frequency, has been calculated versus normalized phase velocity $\omega/(kv_{\text{th}})$ using all of the above implementations. For the initial-value codes, driven initial-value calculations were carried out, in which an elec-

trostatic potential with sinusoidal temporal and spatial dependence was imposed. Figure 8 shows the results of a subset of these calculations. All of calculations show excellent agreement with each other and with the results in Fig. 1 of Ref. [7], providing verification that the SMHS implementation of the LF closures can accurately capture the plasma response in this situation.

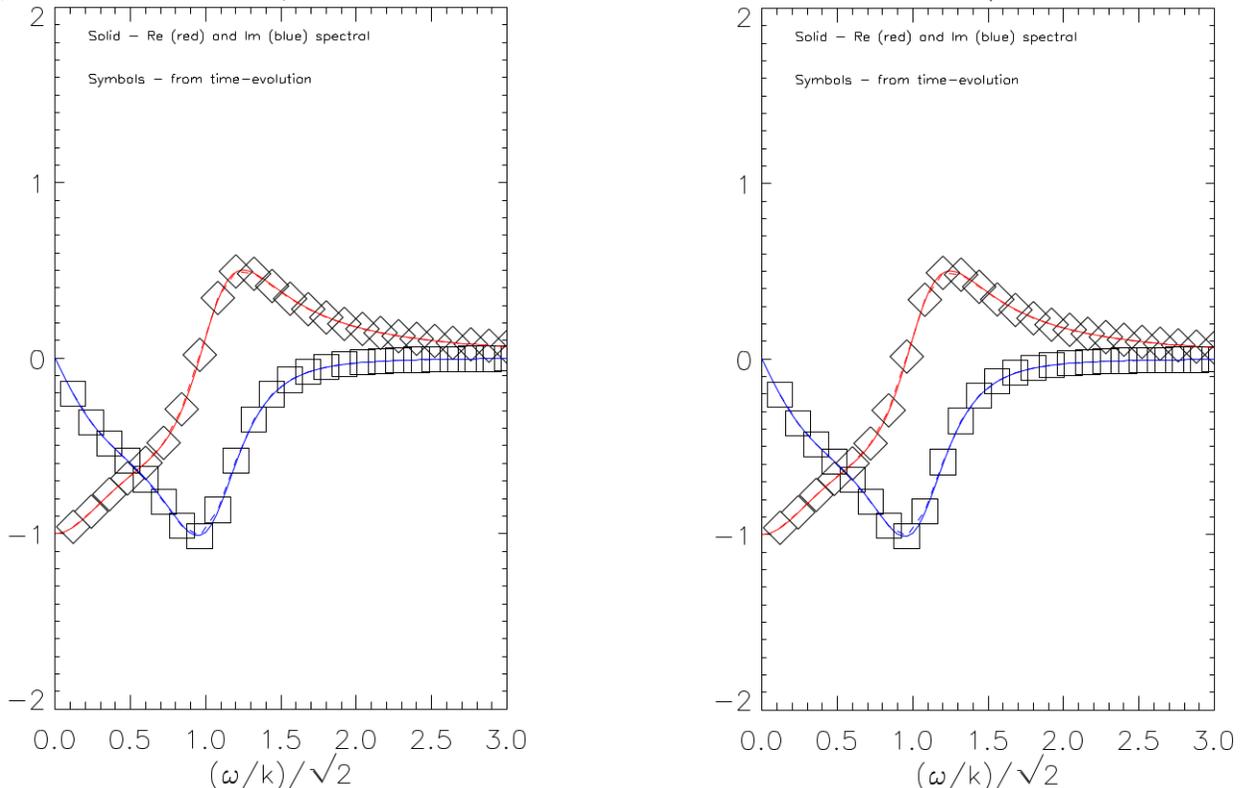


Figure 8: Real and imaginary parts of the negative of the response function $-R = \tilde{n}T_0/n_0q\tilde{\phi}$ versus normalized phase velocity for matrix-eigenvalue and driven initial-value calculations for the closure term in the four-moment Landau-fluid model of Ref. [7], compared with the results of matrix-eigenvalue calculations using. The left frame shows results from the SMHS method, and the right frame shows results from the Fourier method. The points show the real (diamonds) and imaginary (squares) parts from the initial-value calculations, while the solid curves show the results from the matrix-eigenvalue calculations.

F. Toroidal extensions

We have implemented Landau-fluid terms in the BOUT++ magnetic-fusion edge-physics computer code [6]. For the terms that represent Landau damping and phase mixing in the direction parallel to the magnetic field (i.e., $|k_{\parallel}|$ phase-mixing) we use the available

parallel-direction elliptic solver. This solver accounts for the correct offset-periodic boundary conditions in the parallel direction in the core region (see, e.g., Ref. [9]), and therefore allow for perturbations that may extend many parallel passes through the simulation domain (or, physically, poloidal circuits in the parallel direction). For the toroidal-drift-resonance (“ $|\omega_d|$ ”) terms, we have modified the existing perpendicular elliptic solve to solve the required equation, which is a modified Helmholtz equation along lines that represent the trajectories of the combined, thermally averaged grad-B and curvature drifts:

$$\left[\alpha^{2n} k_{\phi 0}^2 - \left(\hat{\mathbf{V}}_{\mathbf{d}} \cdot \nabla \right)^2 \right] \Psi = S.$$

Here, α^{2n} is the coefficient that arises in the SMHS approximation of Eq. (7), $k_{\phi 0}$ is a normalizing toroidal wavenumber, and $\hat{\mathbf{V}}_{\mathbf{d}}$ is the unit vector in the direction of the combined, thermally averaged grad-B and curvature drifts.

Preliminary verification studies of our implementation of the LF terms in a four-moment GLF model in the BOUT++ code have been undertaken for toroidal ion temperature gradient driven modes [17]. The linear growth rates of these modes are found to agree well between the GLF model with the above implementations of the LF terms and gyrokinetic computations using a four-moment toroidal GLF model [9]. These verification studies will be presented in detail elsewhere.

III. SUMMARY AND OUTLOOK

We have developed a new non-Fourier method for the calculation of Landau-fluid (LF) operators. This method is useful for situations with large (including background) spatial inhomogeneities, and for computer codes in which the discretizations are generally configuration-space (not Fourier) based. The method is based on approximations of the nonlocal LF operators by sums of Lorentzians in wavenumber space, which correspond to modified-Helmholtz-equation solves (SMHS) in configuration space. These approximations yield excellent accuracy. A relative error of $\lesssim 1.5\%$ over a wide spectral range is readily achievable with a small enough number of terms that the computational cost is highly competitive with the Fourier method, in the subset of cases where the latter is applicable.

We show for the one-dimensional case that the SMHS method has fast-Fourier-like compu-

tational time scaling, and becomes faster than direct methods such as matrix multiplication once the number of grid cells N_g exceeds approximately 128 – 512. While our timings are specifically for the one-dimensional case, this is already relevant for some implementations that are useful for microturbulence calculations in the tokamak plasma edge region. Examples include the parallel LF operator for electrostatic cases and the toroidal drift resonance ($|\omega_a|$) phase-mixing LF operator provided that the saturated fluctuation amplitudes are small. In truly multidimensional cases, we expect that this method can still be very efficient and give the fast-Fourier-like computational time scaling provided that the associated modified Helmholtz equations are solved via a very efficient linear solution method, such as the algebraic multigrid method.

The SMHS method also has the important advantage for spatially inhomogeneous plasma cases of giving a much more (data-volume) efficient representation of the LF operators than a direct calculation and storage of the nonlocal kernel matrix. A verification of this method using the one-dimensional four-moment linear plasma response has been carried out, both with semi-analytical calculations comparing the response using the SMHS method with the Fourier method, and with driven initial value calculations of this linear response. Excellent agreement is found for the linear response between Fourier and SMHS methods. We remark that this method is also likely to be useful for capturing correct asymptotic (large-wavenumber) form of the various perpendicular (gyrofluid) operators [8].

Spatial inhomogeneities in the operators can arise, for example, through spatial dependences of the temperature or the collision rates. While the underlying theoretical justification for particular forms of the operators in spatially dependent settings is under active investigation, some results are available either for simple cases or based on general considerations of desirable properties of the closure terms. One result in a simplified case is that for uniform temperature (and hence thermal speed), and for a Krook collision operator with a single collisional damping rate (eigenvalue) that has a spatial dependence, the spatial dependence can be entirely folded into a spatial-variable transformation to an “optical depth” variable

$$S(z) = \int^z dz' / \lambda(z')$$

where $\lambda(z)$ is the spatial damping coefficient. In the presence of spatial inhomogeneities also in the temperature, the theory is more complicated because it involves multiple competing

spatially inhomogeneous quantities. However, when there are significant spatial variations, a reasonable generalization for the Landau-fluid operators is one that conserves field quantity upon which it operates, but dissipates fluctuations in that field. For the collisionless case with spatially dependent temperature, examples of such operators are

$$\begin{aligned}
-|k|v_{\text{th}}\tilde{\Psi} &\rightarrow \frac{\partial}{\partial z} \left\{ \sqrt{v_{\text{th}}(z)} \frac{1}{|k|} \left[\sqrt{v_{\text{th}}(z)} \frac{\partial \tilde{\Psi}}{\partial z} \right] \right\}, \\
-|\mathbf{k} \cdot \mathbf{V}_{\mathbf{d}}| \tilde{\Psi} &\rightarrow \nabla \cdot \left\{ \hat{\mathbf{V}}_{\mathbf{d}} \sqrt{V_{\mathbf{d}}(\mathbf{x})} \frac{1}{|\mathbf{k} \cdot \hat{\mathbf{V}}_{\mathbf{d}}|} \left[\sqrt{V_{\mathbf{d}}(\mathbf{x})} (\hat{\mathbf{V}}_{\mathbf{d}} \cdot \nabla \tilde{\Psi}) \right] \right\}.
\end{aligned}$$

Here, $\mathbf{V}_{\mathbf{d}}$ is the thermal average of the magnetic guiding-center (grad-B and curvature) drifts, $\hat{\mathbf{V}}_{\mathbf{d}}$ is the unit vector in its direction, and $V_{\mathbf{d}}$ is its magnitude. The latter form can also be used to implement the nonlinear phase mixing closure [8] that has been found to be important in the nonlinear saturation of toroidal ion temperature gradient driven modes [18]. For the latter implementation, the equilibrium magnetic drift $\mathbf{V}_{\mathbf{d}}$ is replaced by a vector field that represents the finite-gyroradius corrections to the $\mathbf{E} \times \mathbf{B}$ drift. It is also possible to generalize the operators to allow simultaneously for a spatially dependent collisional damping rate and spatially dependent thermal velocity, while preserving the desirable properties of conservation of the field quantity upon which it operates and dissipation of fluctuations in that field. Such a generalization is

$$L(z) \tilde{\Psi} = \frac{\partial}{\partial z} \left[N(z) \frac{\partial \tilde{\Psi}}{\partial z} \right],$$

where $N(z)$ is a positive-definite nonlocal operator such as

$$N(z) \Phi = \sum_n \alpha^n \sqrt{v_{\text{th}}(z) \beta_n(z)} \left[\alpha^{2n} - \frac{\partial^2}{\partial z^2} \right]^{-1} \left[\sqrt{v_{\text{th}}(z) \beta_n(z)} \Phi \right].$$

Further investigations of the theoretical basis for extensions of the closures to include significant spatial inhomogeneity, as well as comparisons with Fokker-Planck calculations are underway [19].

IV. ACKNOWLEDGMENTS

This work was performed for U.S. Department of Energy by LLNL under Contract DE-AC52-07NA27344 and LLNL LDRD project 12-ERD-022. The timing runs shown in Fig. 3 were performed on the Edison computer at NERSC, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. We wish to acknowledge a careful reading of the manuscript and thoughtful comments by the referee. It is also a pleasure to acknowledge fruitful discussions with S. S. Kim, P. W. Xi, X. Q. Xu, G. W. Hammett, D. B. Del-Castillo-Negrete, and E. Held.

References

- [1] Chapman and Cowling *The Mathematical Theory of Non-Uniform Gases*: S. Chapman, and T.G. Cowling (Cambridge University Press, Cambridge UK, 1953).
- [2] S.I. Braginskii, *Transport Processes in a Plasma*, in *Reviews of Plasma Physics* (Consultants Bureau, New York NY, 1965), Vol. 1, p. 205.
- [3] T. H. Stix, *Waves in Plasmas* (1992, Springer-Verlag, New York, Inc.)
- [4] G. F. Chew, M. L. Goldberger, and F. E. Low, *Proc. Roy. Soc., (London)* **A236**, 112 (1956).
- [5] F. L. Hinton and R. D. Hazeltine, *Rev. Mod. Phys.* **48**, 239–308 (1976).
- [6] B. D. Dudson, M. V. Umansky, X. Q. Xu, P. B. Snyder and H. R. Wilson "BOUT++: a framework for parallel plasma fluid simulations" *Computer Physics Communications* **180**, 1467 (2009).
- [7] Gregory W. Hammett and Francis W. Perkins, "Fluid Moment Models for Landau Damping with Application to the Ion-Temperature-Gradient Instability," *Physical Review Letters* **64**, 3019 (1990).
- [8] W. Dorland and G. W. Hammett, "Gyrofluid turbulence models with kinetic effects," *Physics of Fluids B: Plasma Physics* **5** , 812 (1993); doi: 10.1063/1.860934.
- [9] M. A. Beer and G. W. Hammett, "Toroidal Gyrofluid Equations for Simulations of Tokamak Turbulence," *Physics of Plasmas* **3** , 4046 (1996); doi: 10.1063/1.871538.
- [10] G. W. Hammett, W. Dorland, and F. W. Perkins, "Fluid Moment Models of Phase Mixing, Landau Damping, and Nonlinear Gyrokinetic Dynamics," *Phys. Fluids B* **4**, 2052 (1992).
- [11] E. M. Epperlein and R. W. Short, "A practical nonlocal model for electron heat transport in laser plasmas," *Phys. Fluids B* **3**, 3092 (1991); doi: 10.1063/1.859789.
- [12] G. P. Schurtz, Ph. D. Nicolai, and M. Busquet, *Phys. Plasmas* **7**, 4238 (2000).
- [13] Zuoyang Chang and J. D. Callen, "Unified fluid/kinetic description of plasma microinstabilities. Part I: Basic equations in a sheared slab geometry," *Phys. Fluids B* **4**, 1167 (1992).
- [14] John D. Lindl, Peter Amendt, Richard L. Berger, S. Gail Glendinning, Siegfried H. Glenzer, Steven W. Haan, Robert L. Kauffman, Otto L. Landen and Laurence J. Suter, *Phys. Plasmas* **11**, 339 (2004); doi: 10.1063/1.1578638.

- [15] X. Q. Xu, P. W. Xi, A. Dimits, I. Joseph, M. V. Umansky, T. Y. Xia, B. Gui, S. S. Kim, G. Y. Park, T. Rhee, H. Jhang, P. H. Diamond, B. Dudson, and P. B. Snyder, "Gyro-fluid and two-fluid theory and simulations of edge-localized-modes," *Phys. Plasmas* **20**, 056113 (2013).
- [16] The specifications of the compute nodes of the NERSC Edison computer can be found at <http://www.nersc.gov/users/computational-systems/edison/configuration/>.
- [17] S. S. Kim and P. W. Xi, 2013 private communication.
- [18] K. Despain, PhD. thesis, University of Maryland, 2011, drum.lib.umd.edu/bitstream/1903/11465/1/Despain_umd_0117E_11932.pdf
- [19] I. Joseph, M. V. Umansky, and A. M. Dimits, 2013, in preparation.