


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.


In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

Wavelet-based density estimation for noise reduction in plasma simulations using particles

Romain Nguyen van yen^a, Diego del-Castillo-Negrete^{b,*}, Kai Schneider^c,
Marie Farge^a, Guangye Chen^b^aLaboratoire de Météorologie Dynamique-CNRS, École Normale Supérieure, Paris, France^bOak Ridge National Laboratory, Oak Ridge, TN, USA^cLaboratoire de Mécanique, Modélisation et Procédés Propres-CNRS, and Centre de Mathématiques et d'Informatique Université d'Aix-Marseille, France

ARTICLE INFO

Article history:

Received 1 September 2009

Received in revised form 7 December 2009

Accepted 8 December 2009

Available online 4 January 2010

Keywords:

Particle methods

Wavelets

Noise reduction

Plasma physics computations

ABSTRACT

For given computational resources, the accuracy of plasma simulations using particles is mainly limited by the noise due to limited statistical sampling in the reconstruction of the particle distribution function. A method based on wavelet analysis is proposed and tested to reduce this noise. The method, known as wavelet-based density estimation (WBDE), was previously introduced in the statistical literature to estimate probability densities given a finite number of independent measurements. Its novel application to plasma simulations can be viewed as a natural extension of the finite size particles (FSP) approach, with the advantage of estimating more accurately distribution functions that have localized sharp features. The proposed method preserves the moments of the particle distribution function to a good level of accuracy, has no constraints on the dimensionality of the system, does not require an a priori selection of a global smoothing scale, and its able to adapt locally to the smoothness of the density based on the given discrete particle data. Moreover, the computational cost of the denoising stage is of the same order as one time step of a FSP simulation. The method is compared with a recently proposed proper orthogonal decomposition based method, and it is tested with three particle data sets involving different levels of collisionality and interaction with external and self-consistent fields.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

Particle-based numerical methods are routinely used in plasma physics calculations [1,2]. In many cases these methods are more efficient and simpler to implement than the corresponding Eulerian methods. However, particle methods face the well-known statistical sampling limitation of attempting to simulate a physical system containing N particles using $N_p \ll N$ computational particles. Particle methods do not seek to reproduce the exact individual behavior of the particles, but rather to approximate statistical macroscopic quantities like density, current, and temperature. These quantities are determined from the particle distribution function. Therefore, a problem of relevance for the success of particle-based simulations is the reconstruction of the particle distribution function from discrete particle data.

The difference between the distribution function reconstructed from a simulation using N_p particles and the exact distribution function gives rise to a discretization error generically known as “particle noise” due to its random-like character. Understanding and reducing this error is a complex problem of importance in the validation and verification of particle

* Corresponding author. Tel.: +1 865 574 1127; fax: +1 865 576 7926.

E-mail address: delcastillod@ornl.gov (D. del-Castillo-Negrete).

codes, see for example Refs. [3–5] and references therein for a discussion in the context of gyrokinetic calculations. One obvious way to reduce particle noise is by increasing the number of computational particles. However, the unfavorable scaling of the error with the number of particles, $\sim 1/\sqrt{N_p}$ [6,7], puts a severe limitation on this straightforward approach. This has motivated the development of various noise reduction techniques including finite size particles (FSP) [8,9], Monte-Carlo methods [7], weight spreading [10], Fourier-filtering [11], coarse-graining [12], Krook operators [5], smooth interpolation [13], low noise collision operators [14], and Proper Orthogonal Decomposition (POD) methods [15] among others.

In the present paper we propose a wavelet-based method for noise reduction in the reconstruction of particle distribution functions from particle simulation data. The method, known as wavelet-based density estimation (WBDE), was originally introduced in Ref. [16] in the context of statistics to estimate probability densities given a finite number of independent measurements. However, to our knowledge, this method has not been applied before to particle-based computations. WBDE, as used here, is based on a truncation of the wavelet representation of the Dirac delta function associated with each particle. The method yields almost optimal results for functions with unknown local smoothness without compromising computational efficiency, assuming that the particles' coordinates are statistically independent. As a first step in the application of the WBDE method to plasma particle simulations, we limit our attention to “passive denoising”. That is the WBDE method is treated as a post-processing technique applied to independently generated particle data. The problem of “active denoising”, e.g. the application of WBDE methods in the evaluation of self-consistent fields in particle in cell simulations, will not be addressed. This simplification will allow us to assess the efficiency of the proposed noise reduction method in a simple setting. Another simplification pertains the dimensionality. Here, for the sake of simplicity, we limit attention to the reconstruction and denoising problem in two dimensions. However, the extension of the WBDE method to higher dimensions is in principle straightforward.

Collisions, or the absence of them, play an important role in plasma transport problems. Particle methods handle the collisional and non-collisional parts of the dynamics differently. Fokker–Planck-type collision operators are typically introduced in particle methods using Langevin-type stochastic differential equations. On the other hand, the non-collisional part of the dynamics is described using deterministic ordinary differential equations. Collisional dominated problems tend to wash out fine scale structures whereas collisionless problems typically develop fine scale filamentary structures in phase space. Therefore, it is important to test how the efficiency of denoising depends on the level of collisionality. Here we test the WBDE method in strongly collisional, weakly collisional and collisionless regimes. For the strongly collisional regime we consider particle data of force-free collisional relaxation involving energy and pinch-angle scattering. The weakly collisional regime is illustrated using guiding-center particle data of a magnetically confined plasma in toroidal geometry. The collisionless regime is studied using particle-in-cell (PIC) data corresponding to bump-on-tail and two streams instabilities in the Vlasov–Poisson system.

Beyond the role of collisions, the data sets that we are considering open the possibility of exploring the role of external and self-consistent fields in the reconstruction of the particle density. In the collisional relaxation problem, no forces act on the particles. In the guiding-center problem particles interact with an external magnetic field. In the Vlasov–Poisson problem particle interactions are incorporated through a self-consistent electrostatic mean field. One of the goals of this paper is to compare the WBDE method with the Proper Orthogonal Decomposition (POD) density reconstruction method proposed in Ref. [15].

The rest of the paper is organized as follows. In Section 2 we review the main properties of kernel density estimation (KDE) and show its relationship with finite size particles (FSP). We then review basic notions on orthogonal wavelet and multiresolution analysis and outline a step by step algorithm for WBDE. For completeness, we also include in this section a brief description of the POD reconstruction method proposed in Ref. [15]. Section 3 discusses applications of the WBDE method and the comparison with the POD method. We start by post-processing a simulation of plasma relaxation by random collisions against a background thermostat. We then turn to a δf Monte-Carlo simulation in toroidal geometry, whose phase space has been reduced to two dimensions. Finally, we analyze the results of particle-in-cell (PIC) simulations of a 1D Vlasov–Poisson plasma. The conclusions are presented in Section 4.

2. Methods

This section presents the wavelet-based density estimation (WBDE) algorithm. We start by reviewing basic ideas on kernel density estimation (KDE) which is closely related to the use of finite size particles (FSP) in PIC simulations. Following this, we give a brief introduction to wavelet analysis and discuss the WBDE algorithm. For completeness, we also include a brief summary of the POD approach.

2.1. Kernel density estimation

Given a sequence of independent and identically distributed measurements, the non-parametric density estimation problem consists in finding the underlying probability density function (PDF), with no a priori assumptions on its functional form. Here we discuss general ideas on this difficult problem for which a variety of statistical methods have been developed. Further details can be found in the statistics literature, e.g. Ref. [17].

Consider a number N_p of statistically independent particles with phase space coordinates $(\mathbf{X}_n)_{1 \leq n \leq N_p}$ distributed in \mathbb{R}^d according to a PDF f . These data can come from a PIC or a Monte-Carlo, full f or δf simulation. Formally, the sample PDF can be written as

$$f^\delta(\mathbf{x}) = \frac{1}{N_p} \sum_{n=1}^{N_p} \delta(\mathbf{x} - \mathbf{X}_n) \tag{1}$$

where δ is the Dirac distribution. Because of its lack of smoothness, Eq. (1) is far from the actual distribution f according to most reasonable definitions of the error. The dependence of f^δ on the statistical fluctuations in (\mathbf{X}_n) can lead to an artificial increase of the collisionality, which could be problematic in the modeling of near collisionless plasmas of interest to controlled fusion. Beyond introducing dissipation, noise can lead to other problems including self-heating and momentum spread which, for example, is known to be an issue in laser–plasma interaction computations. Also, computations involving derivatives of f , like for example quasilinear fluxes in wave–particle interaction calculations, can be compromised by poor reconstruction techniques.

The simplest method to introduce some smoothness in f^δ is to use a histogram. Consider a tiling of the phase space by a Cartesian grid with N_g^d cells. Let $\{B_\lambda\}_{\lambda \in \Lambda}$ denote the set of all cells with characteristic function χ_λ defined as $\chi_\lambda = 1$ if $\mathbf{x} \in B_\lambda$ and $\chi_\lambda = 0$ otherwise. Then the histogram corresponding to the tiling is

$$f^H(\mathbf{x}) = \sum_{\lambda \in \Lambda} \left(\frac{1}{N_p} \sum_{n=1}^{N_p} \chi_\lambda(\mathbf{X}_n) \right) \chi_\lambda(\mathbf{x}) \tag{2}$$

which can also be viewed as the orthogonal projection of f^δ on the space spanned by the χ_λ . The main difference between f^δ and f^H is that the latter cannot vary at scales finer than the grid scale which is of order N_g^{-1} . By choosing N_g small enough, it is therefore possible to reduce the variance of f^H to very low levels, but the estimate then becomes more and more biased towards a piecewise continuous function, which is not smooth enough to be the true density. Histograms correspond to the nearest grid point (NGP) charge assignment scheme used in the early days of plasma physics computations [8].

One of the most popular methods to achieve higher level of smoothness is kernel density estimation (KDE) [18]. Given $(\mathbf{X}_n)_{1 \leq n \leq N_p}$, the kernel estimate of f is defined as

$$f^K(\mathbf{x}) = \frac{1}{N_p} \sum_{n=1}^{N_p} K(\mathbf{x} - \mathbf{X}_n) \tag{3}$$

where the smoothing kernel K is a positive function, normalized such that $\int K = 1$. Eq. (3) corresponds to the convolution of K with the Dirac delta measure corresponding to each particle. A typical example is the Gaussian kernel

$$K_h(\mathbf{x}) = \frac{1}{(\sqrt{2\pi}h)^d} e^{-\frac{\|\mathbf{x}\|^2}{2h^2}} \tag{4}$$

where the so-called “bandwidth”, or smoothing scale, h , is a free parameter. The optimal smoothing scale depends on how the error is measured. For example, in the one-dimensional case, to minimize the mean L^2 -error between the estimate and the true density, the smoothing volume h^d should scale like $N_p^{-\frac{1}{3}}$, and the resulting error scales like $N_p^{-\frac{2}{3}}$ [17]. As in the case of histograms, the choice of h relies on a trade-off between variance and bias. In the context of plasma physics simulations the kernel K corresponds to the charge assignment function [2].

A significant effort has been devoted to the choice of the function K since it has a strong impact on computational efficiency and on the conservation of global quantities. Concerning h , it has been shown that it should not be much larger than the Debye length λ_D of the plasma to obtain a realistic and stable simulation [1]. Given a certain amount of computational resources, the general tendency has thus been to reduce h as far as possible in order to fit more Debye lengths inside the simulation domain, which means that the effort has been concentrated on reducing the bias term in the error. Since the force fields depend on f through integral equations, like the Poisson equation, that tend to reduce the high wavenumber noise, we do not expect the disastrous scaling $h \propto N_p^{-\frac{1}{3}}$, which would mean $N_p \propto \lambda_D^{5d}$ in d dimensions, to hold. Nevertheless, the problem remains that if we want to preserve high resolution features of f or of the electromagnetic fields, we need to reduce h , and therefore greatly increase the number of particles to prevent the simulation from drowning into noise. Bandwidth selection has long been recognized as the central issue in kernel density estimation [19]. We are not aware of a theoretical or numerical prediction of the optimal value of h taking into account the noise term. To bypass this difficulty, it is possible to use new statistical methods which do not force us to choose a global smoothing parameter. Instead, they adapt locally to the behavior of the density f based on the available data. Wavelet based-density estimation, which we will introduce in the next two sections, is one of these methods.

2.2. Bases of orthogonal wavelets

Wavelets are a standard mathematical tool to analyze and compute non-stationary signals. Here we recall basic concepts and definitions. Further details can be found in Ref. [20] and references therein. The construction takes place in the Hilbert space $L^2(\mathbb{R})$ of square integrable functions. An orthonormal family $(\psi_{j,i}(\mathbf{x}))_{j \in \mathbb{N}, i \in \mathbb{Z}}$ is called a wavelet family when its members are dilations and translations of a fixed function ψ called the mother wavelet:

$$\psi_{j,i}(x) = 2^{j/2}\psi(2^jx - i) \tag{5}$$

where j indexes the scale of the wavelets and i their positions, and ψ satisfies $\int \psi = 0$. In the following we shall always assume that ψ has compact support of length S . The coefficients $\langle f|\psi_{j,i} \rangle = \int f\psi_{j,i}$ of a function f for this family are denoted by $(\tilde{f}_{j,i})$. These coefficients describe the fluctuations of f at scale 2^{-j} around position $\frac{i}{2^j}$. Large values of j correspond to fine scales, and small values to coarse scales. Some members of the commonly used Daubechies 6 wavelet family are shown in the left panel of Fig. 1.

It can be shown that the orthogonal complement in $L^2(\mathbb{R})$ of the linear space spanned by the wavelets is itself orthogonally spanned by the translates of a function φ , called the scaling function. Defining

$$\varphi_{L,i} = 2^{L/2}\varphi(2^Lx - i) \tag{6}$$

and the scaling coefficients $\tilde{f}_{L,i} = \langle f|\varphi_{L,i} \rangle$, one thus has the reconstruction formula:

$$f = \sum_{i=-\infty}^{\infty} \tilde{f}_{L,i}\varphi_{L,i} + \sum_{j=L}^{\infty} \sum_{i=-\infty}^{\infty} \tilde{f}_{j,i}\psi_{j,i} \tag{7}$$

The first sum on the right-hand side of Eq. (7) is a smooth approximation of f at the coarse scale, 2^{-L} , and the second sum corresponds to the addition of details at successively finer scales.

If the wavelet ψ has M vanishing moments:

$$\int x^m\psi(x)dx = 0 \tag{8}$$

for $0 \leq m < M$, and if f is locally m times continuously differentiable around some point x_0 , then a key property of the wavelet expansion is that the coefficients located near x_0 decay like $2^{-j(m+\frac{1}{2})}$ when $j \rightarrow \infty$ [21]. Hence, localized singularities or sharp features in f affect only a finite number of wavelet coefficients within each scale. Another important consequence of (8) of special relevance to particle methods is that, for $0 \leq m < M$, the moments $\int x^m f(x)dx$ of the particle distribution function depend only on its scaling coefficients, and not on its wavelet coefficients.

If the scaling coefficients $\tilde{f}_{j,i}$ at a certain scale J are known, all the wavelet coefficients at coarser scales ($j \leq J$) can be computed using the fast wavelet transform (FWT) algorithm [22]. We shall address the issue of computing the scaling coefficients themselves in Section 2.4.

The generalization to d dimensions involves tensor products of wavelets and scaling functions at the same scale. For example, given a wavelet basis on \mathbb{R} , a wavelet basis on \mathbb{R}^2 can be constructed in the following way:

$$\psi_{j,i_1,i_2}^1(x_1, x_2) = 2^j\psi(2^jx_1 - i_1)\varphi(2^jx_2 - i_2) \tag{9}$$

$$\psi_{j,i_1,i_2}^2(x_1, x_2) = 2^j\varphi(2^jx_1 - i_1)\psi(2^jx_2 - i_2) \tag{10}$$

$$\psi_{j,i_1,i_2}^3(x_1, x_2) = 2^j\psi(2^jx_1 - i_1)\psi(2^jx_2 - i_2) \tag{11}$$

where we refer to the exponent $\mu = 1, 2, 3$ as the direction of the wavelets. This name is easily understood by looking at different wavelets shown in Fig. 1 (right). The corresponding scaling functions are simply given by $2^j\varphi(2^jx_1 - i_1)\varphi(2^jx_2 - i_2)$.

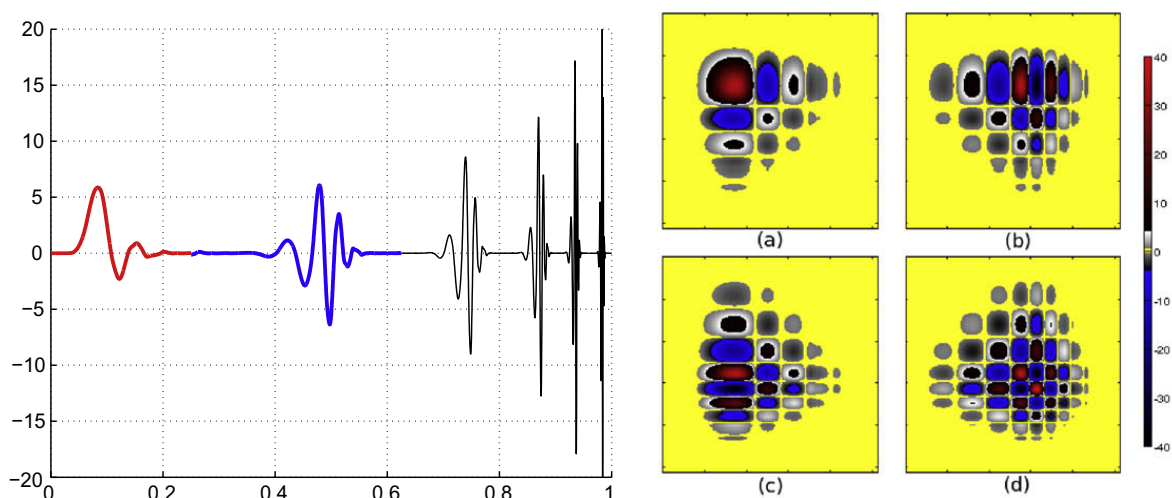


Fig. 1. Daubechies 6 wavelet family. Left, bold red: scaling function φ at scale $j = 5$. Left, bold blue: wavelet ψ at scale $j = 5$. Left, thin black, from left to right: wavelets at scales 6, 7, 8 and 9. Right: (a) 2D scaling function $\varphi(x_1)\varphi(x_2)$. (b) first 2D wavelet $\psi(x_1)\varphi(x_2)$. (c) Second 2D wavelet $\varphi(x_1)\psi(x_2)$. (d) Third 2D wavelet $\psi(x_1)\psi(x_2)$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Wavelets on \mathbb{R}^d are constructed exactly in the same way, but this time using $2^d - 1$ directions. To lighten the notation we write the d -dimensional analog of Eq. (7) as

$$f = \sum_{\lambda \in A_{\phi,L}} \bar{f}_{\lambda} \phi_{\lambda} + \sum_{\lambda \in A_{\psi,L}} \tilde{f}_{\lambda} \psi_{\lambda} \tag{12}$$

where $\lambda = (j, \mathbf{i}, \mu)$ is a multi-index, with the integer j denoting the scale and the integer vector $\mathbf{i} = (i_1, i_2, \dots)$ denoting the position of the wavelet and the scaling function and where $A_{\phi,L} = \{(j, \mathbf{i}, \mu) | j = L, 1 \leq i_1, \dots, i_d \leq 2^j, \mu = 1, 2^d - 1\}$ and $A_{\psi,L} = \{(j, \mathbf{i}, \mu) | j = 1, \dots, L, 1 \leq i_1, \dots, i_d \leq 2^j, \mu = 1, 2^d - 1\}$ are the corresponding index sets.

The wavelet multiresolution reconstruction formula in Eq. (7) involves an infinite sum over the position index i . One way of dealing with this sum is to determine a priori the non-zero coefficients in Eq. (7), and work only with these coefficients, but still retaining the full wavelet basis on \mathbb{R}^d as presented above. An alternative, which we have chosen because it is easier to implement, is to periodize the wavelet transform on a bounded domain [22]. Assuming that the coordinates have been rescaled so that all the particles lie in $[0, 1]^d$, we replace the wavelets and scaling functions by their periodized counterparts:

$$\psi_{j,i}(x) \rightarrow \sum_{l=-\infty}^{\infty} \psi_{j,i}(x+l) \tag{13}$$

$$\varphi_{j,i}(x) \rightarrow \sum_{l=-\infty}^{\infty} \varphi_{j,i}(x+l) \tag{14}$$

Throughout this paper we will consider only periodic wavelets. For the sake of completeness we mention a third alternative which is technically more complicated. It consists in constructing a wavelet basis on a bounded interval [23]. The advantage of this approach is that it does not introduce artificially large wavelet coefficients at the boundaries for functions f that are not periodic.

2.3. Wavelet-based density estimation

The multiscale nature of wavelets allows them to adapt locally to the smoothness of the analyzed function [22]. This fundamental property has triggered their use in a variety of problems. One of their most fruitful applications has been the denoising of intermittent signals [24]. The practical success of wavelet thresholding to reduce noise relies on the fact that the expansion of signals in a wavelet basis is typically sparse. Sparsity means that the interesting features of the signal are well summarized by a small fraction of large wavelet coefficients. On the contrary, the variance of the noise is spread over all the coefficients appearing in Eq. (12). Although the few large coefficients are of course also affected by noise, curing the noise in the small coefficients is already a very good improvement. The original setting of this technique, hereafter referred to as global wavelet shrinkage, requires the noise to be additive, stationary, Gaussian and white. It found a first application in plasma physics in Ref. [25], where coherent bursts were extracted out of plasma density signals. Since Ref. [24], wavelet denoising has been extended to a number of more general situations, like non-Gaussian or correlated additive noise, or to denoise the spectra of locally stationary time series [26]. In particular, the same ideas were developed in Ref. [27,16] to propose a wavelet-based density estimation (WBDE) method based on independent observations. At this point we would like to stress that WBDE assumes nothing about the Gaussianity of the noise, nor on its stationarity. In fact, under the independence hypothesis – which is admittedly quite strong – the statistical properties of the noise are entirely determined by standard probability theory. We refer to Ref. [28] for a review on the applications of wavelets in statistics. In Ref. [29], global wavelet shrinkage was applied directly to the charge density of a 2D PIC code, in a case were the statistical fluctuations were quasi Gaussian and stationary. In particular, an iterative algorithm [30], which crucially relies on the stationarity hypothesis, was used to determine the level of fluctuations. However, in the next section we will show an example where the noise is clearly non-stationary, and this procedure fails.

Let us now describe the WBDE method as we have generalized it to several dimensions. The first step is to expand the sample particle distribution function, f^{δ} , in Eq. (1) in a wavelet basis according to Eq. (12) with the wavelet coefficients

$$\bar{f}_{\lambda}^{\delta} = \langle f^{\delta} | \phi_{\lambda} \rangle = \frac{1}{N_p} \sum_{n=1}^{N_p} \varphi_{\lambda}(X_n) \tag{15}$$

$$\tilde{f}_{\lambda}^{\delta} = \langle f^{\delta} | \psi_{\lambda} \rangle = \frac{1}{N_p} \sum_{n=1}^{N_p} \psi_{\lambda}(X_n) \tag{16}$$

Since this reconstruction is exact, keeping all the wavelet coefficients does not improve the smoothness of f^{δ} . The simple and yet efficient remedy consists in keeping only a subset of the wavelet coefficients in Eq. (12). A straightforward prescription would be to discard all the wavelet coefficients at scales finer than a cut-off scale L . This approach corresponds to a generalization of the histogram method in Eq. (2) with $N_g = 2^L$. Because the characteristic functions χ_{λ} of the cells in a dyadic grid are the scaling functions associated with the Haar wavelet family, Eqs. (12) and (2) are in fact equivalent for this wavelet family. Accordingly, like in the histogram case, we would have to choose L quite low to obtain a stable estimate, at the risk of losing some sharp features of f . Better results can be obtained by keeping some wavelet coefficients down to a much finer scale $J > L$. However, to prevent that statistical fluctuations contaminate the estimate, only those coefficients whose modulus

are above a certain threshold should be kept. We are thus naturally led to a nonlinear thresholding procedure. In the one-dimensional case, values of J, L , and of the threshold within each scale that yield theoretically optimal results have been given in Ref. [16]. This reference discusses the precise smoothness requirements on f , which can accommodate well localized singularities, like shocks and filamentary structures known to arise in collisionless plasma simulations. There remains the question of how to compute the wavelet coefficients $\tilde{f}_{j,i}$ based on the positions of the particles. Although more accurate methods based on (16) may be developed in the future, our present approximation relies on the computation of a histogram, which creates errors of order N_g^{-1} . The complete procedure is described in the following *Wavelet-based density estimation* algorithm:

1. construct a histogram f^H of the particle data with $N_g = 2^{J_g}$ cells in each direction,
2. approximate the scaling coefficients at the finest scale J_g by:

$$\bar{f}_{J_g, \mathbf{i}} \simeq 2^{-J_g/2} f^H(2^{-J_g} \mathbf{i}) \tag{17}$$

3. compute all the needed wavelet coefficients using the FWT algorithm,
4. keep all the coefficients for scales coarser than L , defined by $2^{dL} \sim N_p^{1+2r_0}$ where r_0 is the order of regularity of the wavelet (1 in our case),
5. discard all the coefficients for scales strictly finer than J defined by $2^{dJ} \sim \frac{N_p}{\log_2 N_p}$,
6. for scales j in between L and J , keep only the wavelet coefficients \tilde{f}_λ such that $|\tilde{f}_\lambda| \geq T_j = C \sqrt{\frac{j}{N_p}}$ where C is a constant that must in principle depend on the smoothness of f and on the wavelet family [16].

The choice of parameters for the algorithm was justified rigourously in [16]. The dependence of the threshold on N_p and j can be intuitively understood as follows. Since the particles are assumed to be statistically independent, the standard deviation of each wavelet coefficient is proportional to $\frac{1}{\sqrt{N_p}}$. Because of the L^2 normalization of the wavelets (see (5)), the standard deviation is to a good approximation scale-independent. On the other hand, the central limit theorem implies that the fluctuations are almost Gaussian provided the number of particles is large enough. Therefore, to filter the fluctuations, the threshold should be larger than the typical value taken by a Gaussian random variable with standard deviation $\sigma \sim \frac{1}{\sqrt{N_p}}$.

A standard choice to ensure this level of denoising is to take the threshold proportional to $\sigma \sqrt{2 \ln(M)}$, where M is the number of samples [22]. Since there are 2^{dj} wavelet coefficients at scale j , we get a threshold proportional to $\sigma \sqrt{2 \ln 2^{dj}} = C \sqrt{1/N_p} \sqrt{j}$. In the following, except otherwise indicated, we will assume the proportionality constant to be, $C = \frac{1}{2}$. For the wavelet bases we used orthonormal Daubechies wavelets with 6 vanishing moments and thus support of size $S = 12$ [31]. In our case, $r_0 = 1$, which means that the wavelets have a first derivative but no second derivative, and the size of the wavelets at scale L for $d = 1$ is roughly $N_p^{\frac{1}{3}}$. Since $N_p \gg 1$, it follows from the definition at stage 5 of the algorithm that the size of the wavelets at scale J is orders of magnitude smaller than $N_p^{\frac{1}{3}}$. Using the adaptive properties of wavelets, we are thus able to detect fine scale structures of f without compromising the stability of the estimate. Note that the error at stage 2 could be reduced by using Coiflets [32] instead of Daubechies wavelets, but the gain would be negligible compared to the error made at stage 1. We will denote the WBDE estimate of f as f^W . In the one-dimensional case,

$$f^W = \sum_{i=1}^{2^L} \bar{f}_{L,i} \varphi_{L,i} + \sum_{j=L}^J \sum_{i=1}^{2^j} \rho_j(\tilde{f}_{j,i}) \psi_{j,i} \tag{18}$$

where ρ_j is the thresholding function as defined by stage 6 of the algorithm: $\rho_j(y) = 0$ if $|y| \leq T_j$ and $\rho_j(y) = y$ otherwise.

Finally, let us propose two methods for applying WBDE to postprocess δf simulations. Recall that the Lagrangian equations involved in the δf schemes are identical to their full f counterparts. The only difficulty introduced by the δf method lies in the evaluation of phase space integrals of the form $\delta I = \int A \cdot (f - f_0)$, where A is a function on phase space and f_0 is a known reference distribution function. In these integrals, $f - f_0$ should be replaced by δf , which is in turn written as a product wf , where w is a “weighting” function. Numerically, w is known via its values at particles positions, $w(X_n)$, and the usual expression for δI is thus $\delta I = \sum_{n=1}^{N_p} A(X_n) w(X_n)$. We cannot apply WBDE directly to δf , since this function is not a density function. An elegant approach would be to first apply WBDE to the unweighted distribution f^δ to determine the set of statistically significant wavelet coefficients, and to include the weights only in the final reconstruction (18) of f^W . A simpler approach, which we will illustrate in Section 3.2, consists in renormalizing δf , so that $\int |\delta f| = 1$, and treating it like a density.

2.4. Further issues related to practical implementation

In this section we discuss how the WBDE method handles two issues of direct relevance to plasma simulations: conservation of moments and computational efficiency. As mentioned before, due to the vanishing moments of the wavelets in Eq. (8), the moments up to order M of the particle distribution are solely determined by its scaling function coefficients. As a consequence, we expect the thresholding procedure to conserve these moments, in the sense that

$$\mathcal{M}_{m,k}^W = \int x_k^m f^W(\mathbf{x}) d\mathbf{x} \simeq \int x_k^m f^\delta(\mathbf{x}) d\mathbf{x} = \mathcal{M}_{m,k}^\delta \quad (19)$$

for $0 \leq m \leq M - 1$ and for all $i \in \{1, \dots, d\}$. This conservation holds up to round-off error if the wavelet coefficients can be computed exactly. Due to the type of wavelets that we have used, we were not able to achieve this in the results presented here. There remains a small error related to stages 1 and 2 of the algorithm, namely the construction of f^H and the approximation of the scaling function coefficients by Eq. (17). They are both of order N_g^{-1} . We will present numerical examples of the moments of f^W in the next section.

Conservation of moments is closely related to a peculiarity of the denoised distribution function resulting from the WBDE algorithm: it is not necessarily everywhere positive. Indeed, wavelets are oscillating functions by definition, and removing wavelet coefficients therefore cannot preserve positivity in general. Further studies are needed to assess if this creates numerical instabilities when f^W is used in the computation of self-consistent fields. The same issue was discussed in Ref. [33] where a kernel with two vanishing moments was used to linearly smooth the distribution function. The fact that this kernel is not everywhere positive was not considered harmful in this reference. We acknowledge that it may render the resampling of new particles from f^W , if it is needed in the future, more difficult. There are ways of forcing f^W to be positive, for example by applying the method to \sqrt{f} and then taking the square of the resulting estimate, but this implies the loss of moment conservation, and we have not pursued in this direction.

The number of multiplications required to perform forward and inverse wavelet transforms between scales J and L using the FWT algorithm in d dimensions is $4dS2^{d(J-L)}$, where S is the length of the wavelet filter (12 for the Daubechies filter that we are using). From the definitions of J and L it follows that $2^{d(J-L)} \sim N_p^2 / \log_2 N_p$. Thus, for a simulation in d dimensions involving $N_p = 10^m$ particles, the number of operations scales as $\sim (4Sd \log_2 2/m) \times 10^{2m/3}$. For example, for a 2-dimensional simulation with 10^7 particles, $J = 10$ and $L = 4$. In this case, the computation of the wavelet coefficients require $\sim 1.5 \times 10^6$ multiplications which takes about 0.3 s of CPU time on a desktop computer. For a more demanding 4-dimensional simulation with 10^8 particles, we find that $J = 6$ and $L = 3$, which increases the number of multiplications to 12.5×10^6 . The cost of the binning stage is of order N_p , so that the total number of operations for computing f^W is proportional to N_p , but overall the number of multiplications per particle remains small. On the other hand, advancing the particles in time is likely to require several multiplications per particle. If one wishes to use a finer grid to ensure high accuracy conservation of moments, the storage requirement grows like N_g^d . Thanks to optimized in-place algorithms, the amount of additional memory needed during the computation does not exceed $3S$. Also, an important issue that needs to be kept in mind is that the FWT algorithm requires N_g to be an integer multiple of 2^{J-L} . For comparison purposes, let us recall that most algorithms to compute the POD in 2 dimensions have a complexity proportional to N_g^3 , in addition to the order N_p cost of binning.

To conclude this subsection, Fig. 2 presents an example of the reconstruction of a 1D discontinuous density that illustrates the difference between the KDE and WBDE methods. The probability density function is uniform on the interval $[1/3, 2/3]$ and the estimates were computed on $[0, 1]$ to include the discontinuities. The sample size was 2^{14} , and the binning used $N_g = 2^{16}$ cells to compute the scaling function coefficients. For this 1D case the value $C = 2$ was used to determine the thresholds (step 6 of the algorithm). The KDE estimate is computed using a Gaussian kernel with smoothing scale $h = 0.0138$ [34]. The relative mean square errors associated with the KDE and WBDE estimates are, respectively, 19.6×10^{-3} and 6.97×10^{-3} . The error in the KDE estimate comes mostly from the smoothing of the discontinuities. The better performance of WBDE stems from the much sharper representation of these discontinuities. It is also observed that the WBDE estimate is not everywhere positive. The approximate conservation of moments is demonstrated on Table 1. Note that the error on all these moments for




Fig. 2. Estimation of the density of a sample of size 2^{14} drawn uniformly in $[1/3, 2/3]$, using Gaussian kernels (left) or wavelets (right). The discontinuous analytical density is plotted with a dashed line in the two cases.

Table 1

Relative absolute difference between the moments of f^s and those of f^K and f^W , for the distribution function corresponding to Fig. 2.

	$m = 0$	$m = 1$	$m = 2$	$m = 4$
f^K	$1.81 \cdot 10^{-5}$	$1.70 \cdot 10^{-5}$	$7.52 \cdot 10^{-4}$	$3.90 \cdot 10^{-3}$
f^W	$1.08 \cdot 10^{-11}$	$1.52 \cdot 10^{-5}$	$2.93 \cdot 10^{-5}$	$5.52 \cdot 10^{-5}$

f^W could be made arbitrary low by increasing N_g . The overshoots could also be mitigated by using nearly shift invariant wavelets [35].

2.5. Proper orthogonal decomposition method

For completeness, in this subsection we present a brief review of the POD density reconstruction method. For the sake of comparison with the WBDE method, we limit attention to the time independent case. Further details, including the reconstruction of time dependent densities using POD methods can be found in Ref. [15].

The first step in the POD method is to construct the histogram f^H from the particle data. This density is represented by an $N_x \times N_y$ matrix \hat{f}_{ij} containing the fraction of particles with coordinates (x, y) such that $X_i \leq x < X_{i+1}$ and $Y_i \leq y < Y_{i+1}$. In two dimensions, the POD method is based on the singular value decomposition (SVD) of the histogram. According to the SVD theorem [37], the matrix \hat{f} can always be factorized as $\hat{f} = UWV^t$, where U and V are $N_x \times N_x$ and $N_y \times N_y$ orthogonal matrices, $UU^t = VV^t = I$, and W is a diagonal matrix, $W = \text{diag}(w_1, w_2, \dots, w_N)$, such that $w_1 \geq w_2 \geq \dots \geq w_N \geq 0$ with $N = \min(N_x, N_y)$.

In vector form, the decomposition can be expressed as

$$\hat{f}_{ij} = \sum_{k=1}^N w_k u_i^{(k)} v_j^{(k)} \tag{20}$$

where the N_x -dimensional vectors, $u_i^{(k)}$, and the N_y -dimensional vectors, $v_j^{(k)}$, are the orthonormal POD modes and correspond to the columns of the matrices U and V , respectively. Given the decomposition in Eq. (20), we define the rank- r approximation of \hat{f} as

$$\hat{f}_{ij}^{(r)} = \sum_{k=1}^r w_k u_i^{(k)} v_j^{(k)} \tag{21}$$

where $1 \leq r < N$, and define the corresponding rank- r reconstruction error as

$$e(r) = \|\hat{f} - \hat{f}^{(r)}\|^2 = \sum_{i=r+1}^N w_i^2 \tag{22}$$

where $\|A\| = \sqrt{\sum_{ij} A_{ij}^2}$ is the Frobenius norm. Since $\hat{f}^{(r=N)} = \hat{f}$, we define $e(N) = 0$. The key property of the POD is that the approximation in Eq. (21) is optimal in the sense that

$$e(r) = \min \left\{ \|\hat{f} - g\|^2 \mid \text{rank}(g) = r \right\} \tag{23}$$

That is, of all the possible rank- r Cartesian product approximations of \hat{f} , $\hat{f}^{(r)}$ is the closest to \hat{f} in the Frobenius norm.

The SVD spectrum, $\{w_k\}$, of noise free coherent signals decays very rapidly after a few modes, but the spectrum of noise dominated signals is relatively flat and decays very slowly. When a coherent signal is contaminated with low level noise, the SVD spectrum exhibits an initial rapid decay followed by a weakly decaying spectrum known as the noisy plateau. In the POD method the denoised density is defined as the truncation $f^P = \hat{f}^{(r_c)}$, where r_c corresponds to the rank where the noisy plateau starts. In general it is difficult to provide a precise a priori estimate of r_c , and this is one of the potential limitations of the POD method. One possible quantitative criterion used in Ref. [15] is to consider the relative decay of the spectrum, $\Delta(k) = (w_{k+1} - w_k)/(w_2 - w_1)$, for $k > 1$, and define r_c by the condition $\Delta(r_c) \leq \Delta_c$ where Δ_c is a predetermined threshold.

3. Applications

In this section, we apply the WBDE method to reconstruct and denoise the particle distribution function starting from discrete particle data. The data corresponds to three different groups of simulations: collisional thermalization with a background plasma, guiding center transport in toroidal geometry, and Vlasov–Poisson electrostatic instabilities. We will compare the WBDE and POD methods in all three cases. Note that the first two groups of simulations were already analyzed using POD methods in Ref. [15]. The third group, which is new here, allows the testing of the reconstruction algorithms in a collisionless system that incorporates the self-consistent evaluation of the forces acting on the particles, as opposed to the collisional, test particle problems analyzed before. When comparing the two methods, it is important to keep in mind that POD has one free parameter, namely the number r of singular vectors that are retained to reconstruct the denoised

distribution function. In the cases studied here, we used a best guess for r based on the properties of the reconstruction. In Ref. [15] the POD method was developed and applied to time independent and time dependent data sets. However, in the comparison with the WBDE method, we limit attention to 2-dimensional time independent data sets.

The accuracy of the reconstruction of the density at a fixed time t will be monitored using the mean square error

$$e = \sum_{ij} |f^{est}(x_i, y_j; t) - f^{ref}(x_i, y_j; t)|^2 \quad (24)$$

where (x_i, y_j) are the coordinates of the nodes of a prescribed $N_g \times N_g$ grid in phase space, and f^{est} denotes the estimated density computed from a sample with N_p particles. For the WBDE method $f^{est} = f^W$, and for the POD method $f^{est} = f^P$. In principle, the reference density, f^{ref} , in Eq. (24) should be the density function obtained from the exact solution of the corresponding continuum model, e.g. the Fokker–Planck or the Vlasov–Poisson system. However, when no explicit solution is available, we will set $f^{ref} = f^H$, where f^H is the histogram corresponding to a simulation with a maximum number of particles available, which in the cases reported here correspond to $N_p = 10^6$. We will also use the normalized error

$$e_0 = \frac{e}{\sum_{i,j} |f^{ref}(x_i, y_j; t)|^2} \quad (25)$$

3.1. Collisional thermalization with a background plasma

This first example models the relaxation of a non-equilibrium plasma by collisional damping and pitch angle scattering on a thermal background. The plasma is spatially homogeneous and is represented by an ensemble of N_p particles in a three-dimensional velocity space. Assuming a strong magnetic field, the dynamics can be reduced to two degrees of freedom: the magnitude of the particle velocity, v , and the particle pitch, $\lambda = \cos \theta$, where θ is the angle between the particle velocity and the magnetic field. In the continuum limit the particle distribution function is governed by the Fokker–Planck equation, which in the particle description corresponds to the stochastic differential equations

$$d\lambda = -\lambda v_D dt - \sqrt{v_D(1-\lambda^2)} d\eta_\lambda \quad (26)$$

$$dv = -\left[\alpha v_s v - \frac{1}{2v^2} \frac{d}{dv} (v_\parallel v^4)\right] dt + \sqrt{v^2 v_\parallel} d\eta_v \quad (27)$$

describing the evolution of $v \in (0, \infty)$ and $\lambda \in [-1, 1]$ for each particle, where $d\eta_\lambda$ and $d\eta_v$ are independent Wiener stochastic processes and v_D, v_s and v_\parallel are functions of v . For further details on the model see Ref. [15] and references therein.

We considered simulations with $N_p = 10^3, 10^4, 10^5$ and 10^6 particles. The initial conditions of the ensemble of particles were obtained by sampling a distribution of the form

$$f(v, \lambda, t = 0) = Av^2 \exp \left\{ -\frac{1}{2} \left[\frac{(\lambda - \lambda_0)^2}{\sigma_\lambda^2} + \frac{(v - v_0)^2}{\sigma_v^2} \right] \right\} \quad (28)$$

where a v^2 factor has been included in the definition of the initial condition so that the volume element is simply $dv d\lambda$. A is a normalization constant, $\lambda_0 = 0.25$, $v_0 = 5$, $\sigma_\lambda = 0.25$ and $\sigma_v = 0.75$. This relatively simple problem is particularly well suited for the WBDE method because the simulated particles do not interact and therefore statistical correlations cannot build-up between them.

Before applying the WBDE method, we analyze the sparsity of the wavelet expansion of f^δ , and compare the number of modes kept and the reconstruction error for different thresholding rules. The plot in the upper left panel of Fig. 3 shows the absolute values of the wavelet coefficients in decreasing order at different fixed times. The wavelet coefficients exhibit a clear rapid decay beyond the few significant modes corresponding to the gross shape of the Maxwellian distribution. A similar trend is observed in the coefficients of the POD expansion shown in the upper right panel of Fig. 3. However, in the wavelet case the exponential decay starts after more than 100 modes, whereas in the POD case the exponential decay starts after only one mode.

The two panels at the bottom of Fig. 3 show the square root of the reconstruction error normalized by $N_g, \sqrt{e}/N_g^2$, in the WBDE and POD methods. Because in this case we do not have access to the exact solution of the corresponding Fokker–Planck equation at the prescribed time, we used f^H computed using $N_p = 10^6$ particles as the reference density f^{ref} in Eq. (24). The error observed when applying a global threshold to the wavelet coefficients (bottom left panel in Fig. 3) is minimal when around 100 modes are kept whereas in the POD case (bottom right panel in Fig. 3) the minimal error is reached with about two or three modes. Fig. 3 also shows the wavelet threshold obtained by applying the iterative algorithm based on the stationary Gaussian white noise hypothesis [30,25]. The error corresponding to this threshold is larger than the optimal error because the noise in this problem is very non-stationary due to the lack of statistical fluctuations in the regions where particles are absent. In contrast, the error corresponding to the WBDE procedure (dash-dotted line in Fig. 3) is typically smaller than the optimal error obtained by global thresholding. This is not a contradiction, because the WBDE procedure is not a global threshold, but a level dependent threshold.




Fig. 3. Wavelet and POD analyses of collisional relaxation particle data at different fixed times, with $N_p = 10^5$. Top left: absolute values of the wavelet coefficients sorted by decreasing order (full lines), and thresholds given by the global wavelet shrinkage algorithm (dashed lines). Top right: POD singular values of the histogram used to construct f^p . Bottom left: error estimate $\frac{e^{1/2}}{N_g^2}$ with respect to the run for $N_p = 10^6$ as a function of the number of retained wavelet coefficients (full lines), error obtained when using the global wavelet shrinkage threshold (dashed lines), and error obtained using the WBDE method (dash-dotted lines). Bottom right: error estimate $\frac{e^{1/2}}{N_g^2}$ for f^p as a function of the number l of retained POD singular values.

Fig. 4 compares at different times the densities estimated with the WBDE and the POD (retaining only three modes) methods using $N_p = 10^5$ particles with the histograms computed using $N_p = 10^5$ and 10^6 particles. The key feature to observe is that the level of smoothness of f^W and f^p corresponding to $N_p = 10^5$ is similar, if not greater, than the level of smoothness in f^H computed using ten times more particles, i.e. $N_p = 10^6$ particles. Table 2 summarizes the normalized reconstruction errors for $N_p = 10^5$ according to Eq. (24) using f^H with $N_p = 10^6$ as f^{ref} . The WBDE and POD denoising methods offer a significant improvement, approximately by a factor 2, over the raw histogram method.

A more detailed comparison of the estimates can be achieved by focusing on the Maxwellian final equilibrium state

$$f_M(v) = \frac{2}{\sqrt{\pi}} v^2 e^{-v^2} \tag{29}$$

where, as in Eq. (28), the v^2 metric factor has been included in the definition of the distribution. For these calculations we considered sets of particles sampled from Eq. (29) in the compact domain $[-1, 1] \times [0, 4]$. Since f_M is an exact equilibrium solution of the Fokker–Planck equation, the ensemble of particles is in statistical equilibrium but it exhibits fluctuations due to the finite number of particles. Fig. 5 shows the dependence of the square root of the reconstruction error, e (normalized by N_g^2) on the number of particles N_p and the grid resolution N_g for the WBDE and POD methods. The main advantage of this example is that the exact density f^M can be used as the reference density f^{ref} in the evaluation of the error.

3.2. Collisional guiding center transport in toroidal geometry

The previous example focused on collisional dynamics. However, in addition to collisions, plasma transport involves external and self-consistent electromagnetic fields and it is of interest to test the particle density reconstruction algorithms in these more complicated settings. As a first step for solving this challenging problem, we consider a plasma subject to collisions and an externally applied fixed magnetic field in toroidal geometry. The choice of the field geometry and structure was motivated by problems of interest to magnetically confined fusion plasmas. The data were presented and analyzed using

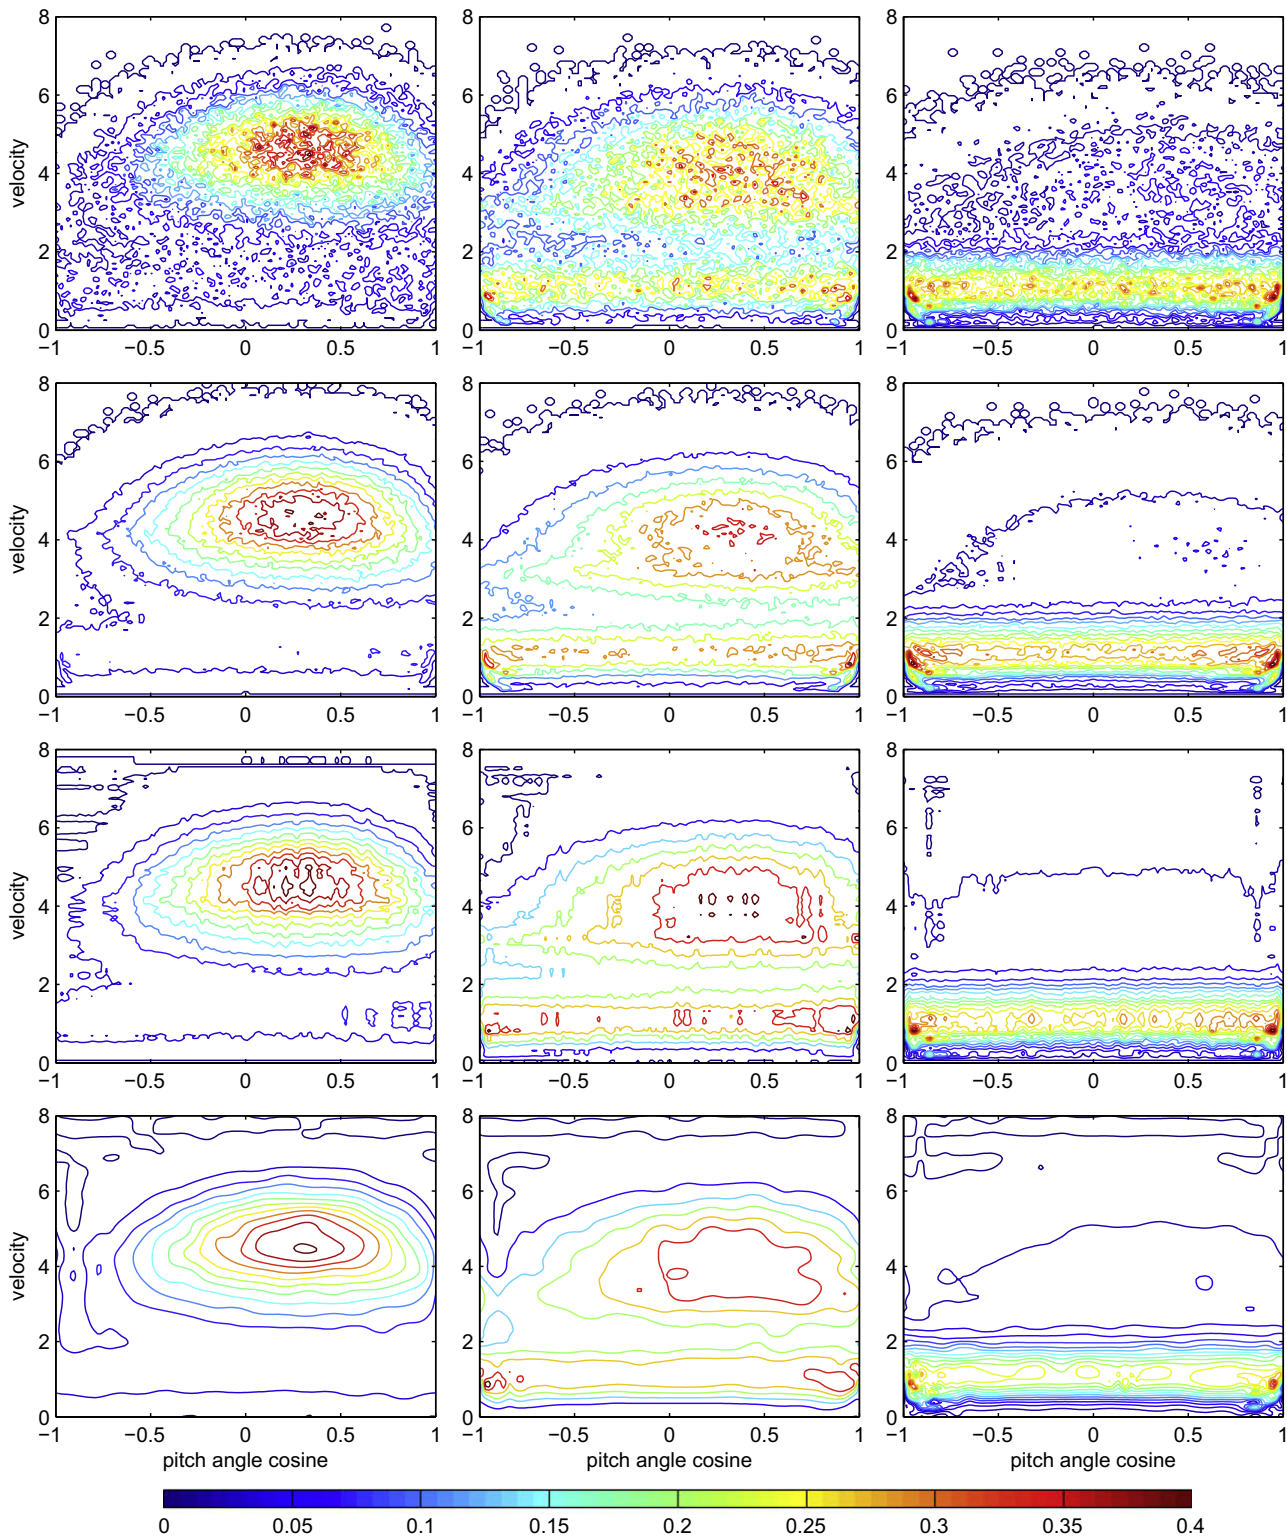


Fig. 4. Contour-plots of estimates of f for the collisional relaxation particle data. First row: histogram method estimated using $N_p = 10^5$ particles. Second row: histogram method estimated using $N_p = 10^6$ particles. Third row: POD method estimated using $N_p = 10^5$ particles. Fourth row: WBDE method estimated using $N_p = 10^5$ particles. The three columns correspond to $t = 28$, $t = 44$ and $t = 72$, respectively. The plots show twenty isolines, equally spaced in the interval $[0, 0.4]$.

the POD method in Ref. [15]. The phase space of the simulation is five-dimensional. However, as in Ref. [15], we limit attention to the denoising of the particles distribution function along two coordinates corresponding to the poloidal angle $\theta \in [0, 2\pi]$ and the cosine of the pitch angle $\mu \in [-1, 1]$. The remaining three coordinates have been averaged out for the purpose of this study. The θ coordinate is periodic, but the pitch coordinate μ is not.

Table 2

Normalized root mean squared error e_0 (25) for the histogram, POD and WBDE estimates of the particle distribution function for $N_p = 10^5$ at three different times of the Maxwellian relaxation problem.

	$t = 28$	$t = 44$	$t = 72$
f^H	0.14	0.17	0.12
f^P	0.068	0.090	0.094
f^W	0.064	0.094	0.088




Fig. 5. Reconstruction error, $\frac{e^{1/2}}{N_g^2}$, as a function of N_p for the collisional relaxation particle data corresponding to the Maxwellian equilibrium state. Bold solid lines correspond to the WBDE method, bold dashed lines correspond to the POD method, and thin dashed lines correspond to the histogram method.

An important issue to consider is that the data were generated using a δf code (DELTA5D). Based on an expansion on $\rho/L \ll 1$ (where ρ is the characteristic Larmor radius and L a typical equilibrium length scale) the distribution function is decomposed into a Maxwellian part f_M and a first-order perturbation δf represented as a collection of particles (markers)

$$\delta f(\mathbf{x}) = \sum_n W_n \delta(\mathbf{x} - \mathbf{X}_n), \tag{30}$$

like in Eq. (1) except that each marker is assigned a time dependent weight W_n whose time evolution depends on the Maxwellian background [38]. The direct use of $\delta f(\mathbf{x})$ is problematic in the WBDE method because δf is not a probability density. To circumvent this problem the WBDE method was applied after normalizing the δf distribution so that $\int |\delta f|^H = 1$, on a 128×128 grid.

Fig. 6 shows contour plots of the histogram f^H corresponding to $N_p = 32 \times 10^3$, $N_p = 64 \times 10^3$, and $N_p = 1024 \times 10^3$ along with the WBDE and POD reconstructed densities. The POD reconstructions were done using $r = 3$ modes, as in Ref. [15]. It is observed that comparatively high levels of smoothness can be achieved with considerably less particles by using either the WBDE or POD reconstruction methods. The WBDE method provides better results for the $\delta f \sim 0$ contours. This is because POD modes are tensor product functions, that have difficulties in approximating the triangular shape of these contour lines. Note that the boundary artifacts due to periodization of the Daubechies wavelets do not seem to be very critical. The large wavelet coefficients associated with the discontinuity between the values of δf at $\mu = \pm 1$ are not thresholded, so that the discontinuity is preserved in the denoised function. Fig. 7 compares the reconstruction errors in the WBDE, POD, and histogram methods as functions of the number of particles. To evaluate the error we used f^H computed with $N_p = 1024 \times 10^3$ as the reference density f^{ref} . As in the collisional transport problem, the error is reduced roughly by a factor 2 for both methods compared to the raw histogram. Note that the scaling with N_p is slightly better for WBDE than for POD.

3.3. Collisionless electrostatic instabilities

In this section we apply the WBDE and POD methods to reconstruct the single particle distribution function from discrete particle data obtained from PIC simulations of a Vlasov–Poisson plasma. We consider a one-dimensional, electrostatic, collisionless electron plasma with an ion neutralizing background in a finite size domain with periodic boundary conditions. The dynamics of the distribution function is governed by the system of equations

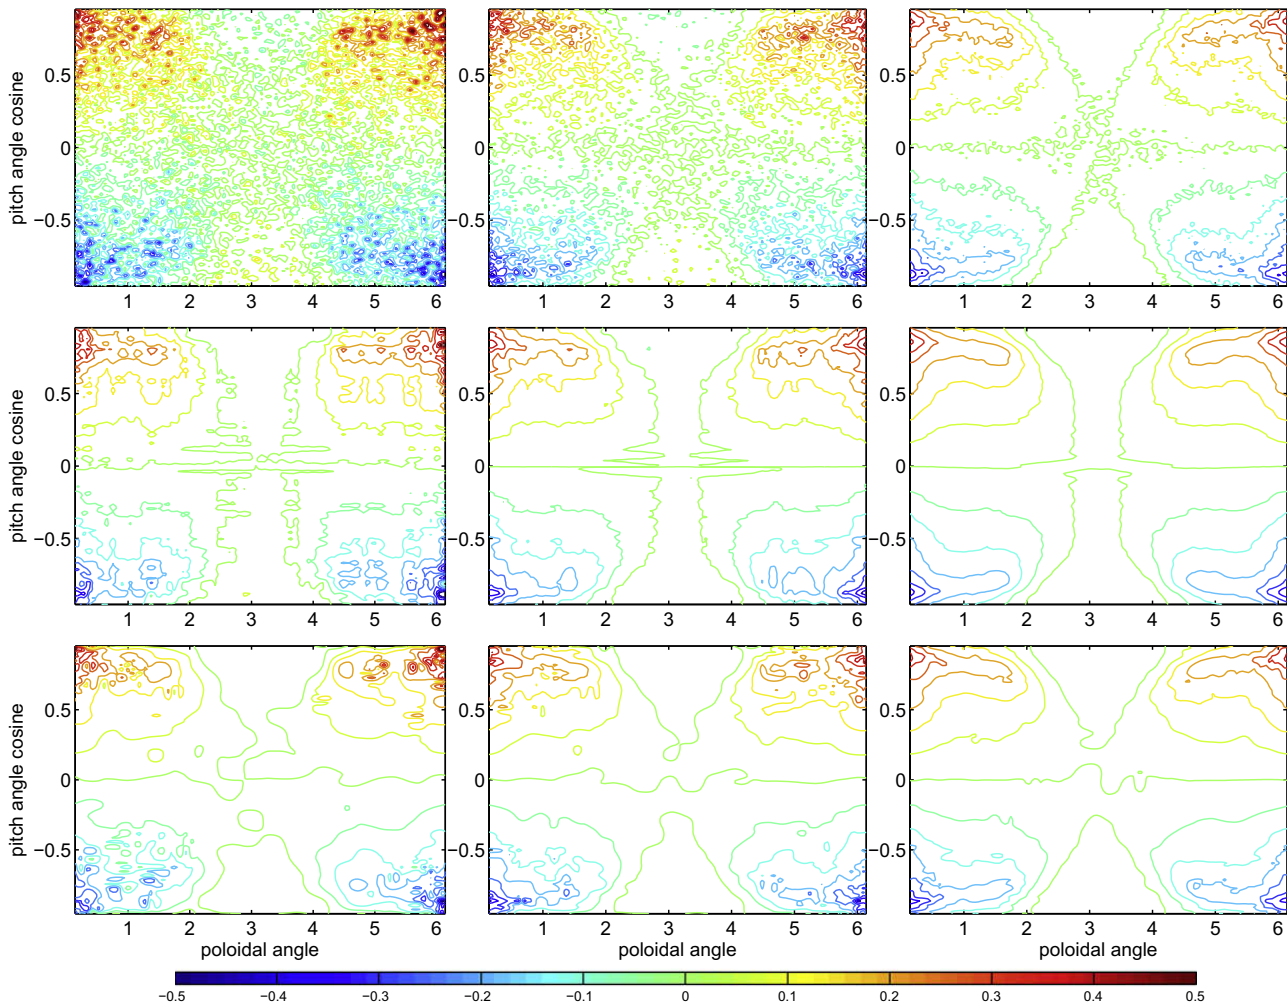


Fig. 6. Contour plots of estimates of δf for the collisional guiding center transport particle data: histogram method (first row), POD method (second row), and WBDE method (third row). The left, center and right columns correspond to $N_p = 32 \cdot 10^3$ (left), $N_p = 128 \cdot 10^3$ (middle) and $N_p = 1024 \cdot 10^3$ (right), respectively. The plots show 17 isolines equally spaced within the interval $[-0.5, 0.5]$.

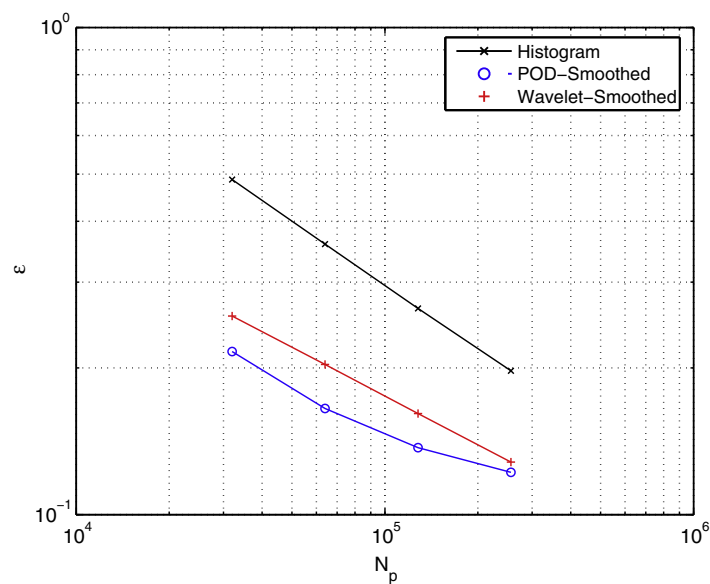


Fig. 7. Error estimate, $\frac{\epsilon^{1/2}}{N_p}$, for collisional guiding center transport particle data according to the histogram, the POD, and the wavelet methods.

$$\partial_t f + v \partial_x f + \partial_x \phi \partial_v f = 0 \tag{31}$$

$$\partial_x^2 \phi = \zeta \int f(x, v, t) dv - 1 \tag{32}$$

where the variables have been non-dimensionalized using the Debye length as length scale and the plasma frequency as time scale, and ζ is the normalized length of the system. Following the standard PIC methodology [1], we solve the Poisson equation on a grid and integrate the particle equations in time using a leap-frog method. The reconstruction of the charge density uses a triangular shape function. We consider two initial conditions: the first one leads to a bump-on-tail instability, and the second one to a two streams instability.

3.3.1. Bump-on-tail instability

To trigger a bump-on-tail instability we initialized ensembles of particles by sampling the distribution function

$$f_0(x, v) = \frac{2}{3\pi\zeta} \frac{1 - 2qv + 2v^2}{(1 + v^2)^2} \tag{33}$$

using a pseudo-random number generator. This equilibrium is stable for $q \leq 1$ and unstable for $q > 1$. The dispersion relation and linear stability analysis for this equilibrium studied in Ref. [36] was used to benchmark the PIC code as shown in Fig. 8. In all the computations presented here $q = 1.25$ and $N_p = 10^4, 10^5$ and 10^6 . The spatial domain size was set to $\zeta = 16.52$ to fit the wavelength of the most unstable mode.

Since the value of q is relatively close to the marginal value, the instability grows weakly and is concentrated in a narrow band in phase space centered around the point where the bump is located, $v \approx 1$ in this case. In order to unveil the non-trivial dynamics we focus the analysis in the band $v \in (-3, 3)$, and plot the departure of the particle distribution function from the initial background equilibrium. The POD method is applied directly to $\delta f^H = f^H(x, v, t) - f_0(x, v)$, but the WBDE method is applied to the full $f^H(x, v, t)$, and $f_0(x, v)$ is subtracted only for visualization. Note that because we are considering only a subset of phase space, the effective numbers of particles, $N_p = 7318, N_p = 73, 143$ and $N_p = 731, 472$, are smaller than the nominal numbers of particles, $N_p = 10^4, N_p = 10^5$ and $N_p = 10^6$, respectively.

Fig. 9 shows contour plots of δf for different number of particles. Since the instability is seeded only by random fluctuations in the initial condition, increasing N_p delays the onset of the linear stability and this leads to a phase shift of the nonlinear saturated regime. To aid the comparison of the saturated regime for different numbers of particles we have eliminated this phase shift by centering the peak of the particle distributions in the middle of the computational domain. A 256×256 grid was used in the WBDE method, and a 50×50 grid was used for the histogram and the POD methods. The thresholds for the POD method where $r = 1, r = 2,$ and $r = 3$ for $N_p = 10^4, N_p = 10^5$ and $N_p = 10^6$, respectively. Except for the case where $N_p = 10^4$, both the POD and WBDE estimates are very smooth, in agreement with the expected behavior of f for this instability. It is observed that the level of smoothness of the histogram estimated using 10^6 particles is comparable to the level of smoothness achieved after denoising using only 10^5 particles. One should mention that for scales between L and J occurring in the WBDE algorithm we find that none of the wavelet coefficients are above the threshold at each scale. In fact, a simple KDE estimate with a large enough smoothing scale would probably do the job pretty well for this kind of instabilities which do not induce abrupt variations in f . Table 3 shows the POD and WBDE reconstruction errors for $N_p = 10^4$ and $N_p = 10^5$. The error is computed using formula (25), taking for f_{ref} the histogram obtained from the simulation with $N_p = 10^6$.




Fig. 8. Electrostatic energy as a function of time in the Vlasov–Poisson PIC simulations of the bump-on-tail instability for different numbers of particles. The straight lines denote the growth rate predicted by linear stability theory [36].

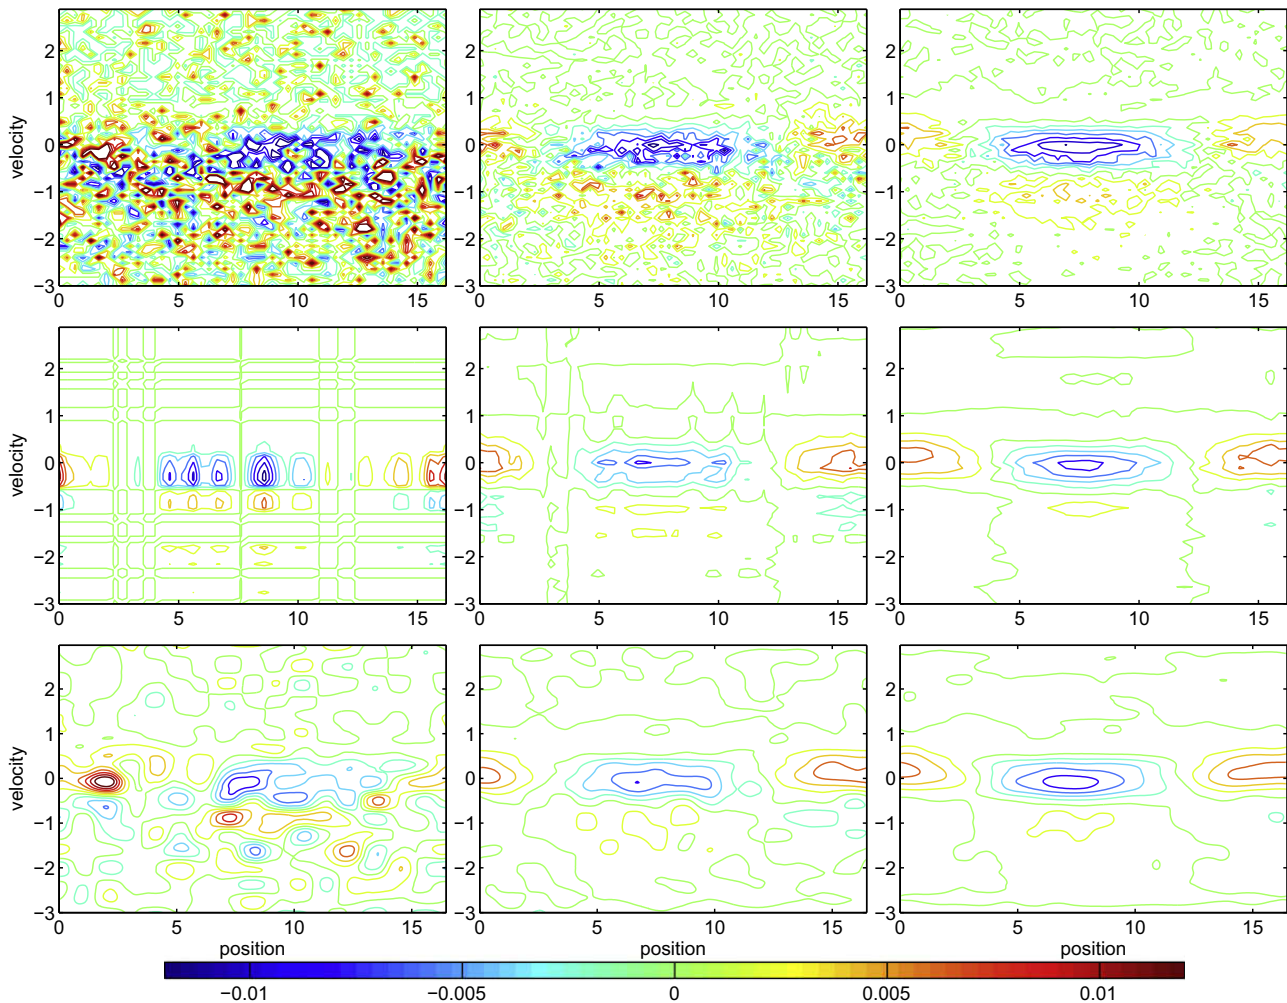


Fig. 9. Contour plots of estimates of δf for the bump-on-tail instability PIC data at $t = 149$: histogram method (first row), POD method (second row), and WBDE method (third row). The left, center and right columns correspond to $N_p = 10^4$, $N_p = 10^5$ and $N_p = 10^6$ particles, respectively. The plots show 13 contour lines equally spaced within the interval $[-0.012, 0.012]$.

Table 3

Comparison of normalized root mean square errors e_0 (25) for the raw histogram and for the WBDE and POD methods, for the bump-on-tail instability at $t = 149$, depending on the number of particles. The simulation with $N_p = 10^6$ is used as a reference to compute the error.

	$N_p = 10^4$	$N_p = 10^5$
f^H	0.443	0.140
f^P	0.163	0.090
f^W	0.173	0.086

Fig. 10 shows the relative error on the second-order moment:

$$\frac{|\mathcal{M}_{v,2}^W - \mathcal{M}_{v,2}^\delta|}{\mathcal{M}_{v,2}^\delta}$$

where $\mathcal{M}_{v,2}^W$ is defined by (19). A similar quantity is also represented for f^H and f^P . Time and number of particles are kept fixed at $t = 149$ and $N_p = 10^6$, and only grid resolution is varied. As expected, f^H and f^W conserve the second-order moment with accuracy $O(N_g^{-1})$. The errors corresponding to f^P is of the same order of magnitude but seems to reach a plateau for $N_g \simeq 1024$. This may be due to the fact that for $N_g \geq 1024$, there is less than one particle per cell of the histogram used to compute f_p .

3.3.2. Two-streams instability

As a second example we consider the standard two-streams instability with an initial condition consisting of two counter-propagating cold electron beams initially located at $v = -1$ and $v = 1$. This case is conceptually different from the pre-

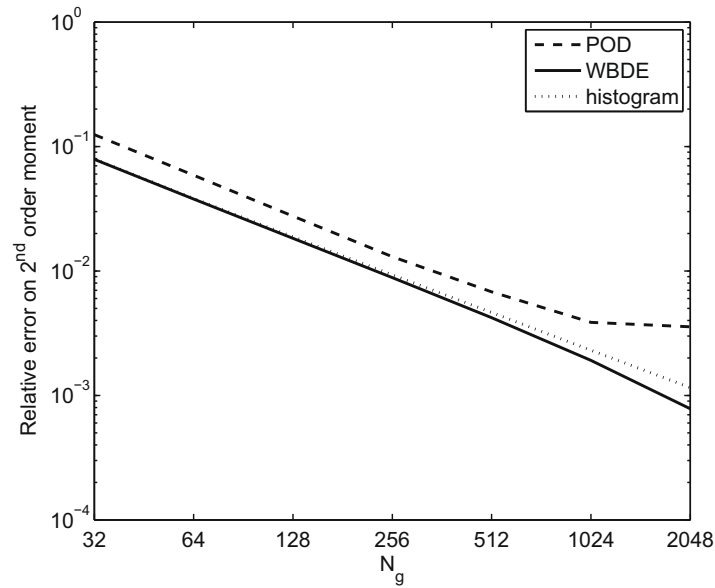


Fig. 10. Relative error on the second-order moment as a function of the grid resolution, N_g , in the POD, WBDE, and histogram methods for the bump-on-tail instability particle data at $t = 149$, with $N_p = 10^6$ particles.

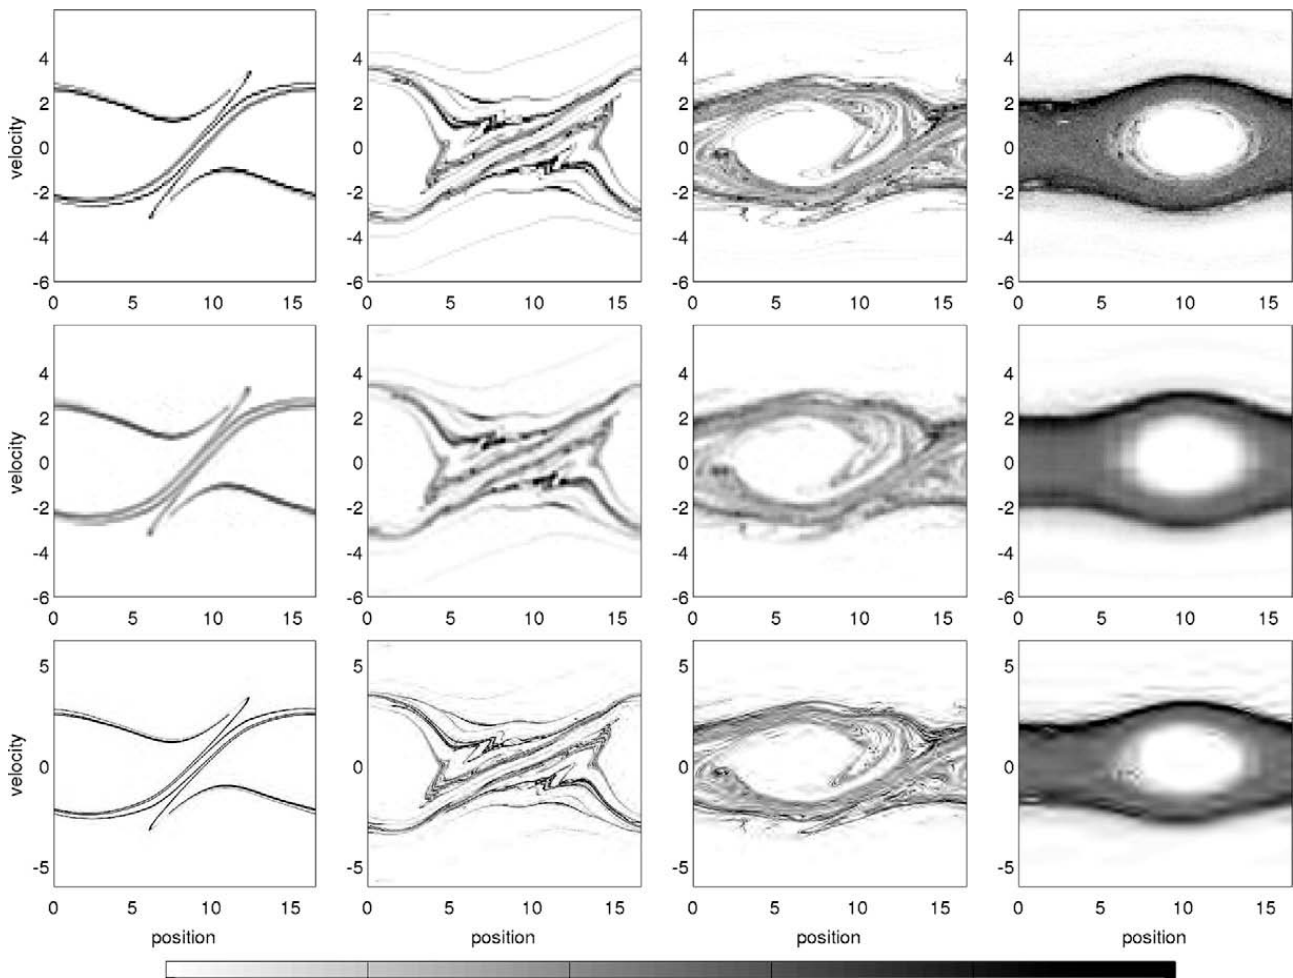


Fig. 11. Contour plots of estimates of f for the two streams instability PIC data at times $t = 40$, $t = 60$, $t = 100$ and $t = 400$ (left to right). Histogram method (first row), WBDE method (second row), and POD method (third row). The gray level tone varies uniformly in the interval $[0, A]$, where $A = 0.15$, $A = 0.08$, $A = 0.05$ and $A = 0.025$ in the first, second, third and fourth columns, respectively.

vious one because the initial condition depends trivially on the velocity. Therefore, there is no statistical error in the sampling of the distribution and the noise builds up only due to the self-consistent interactions between particles. In other words, there is initially a strong correlation between particles' coordinates, which will eventually almost vanish. This situation offers a way to test robustness of the WBDE method with respect to the underlying decorrelation hypothesis.

The analysis is focused on four stages of the instability, corresponding to $t = 40, 60, 100,$ and 400 . Fig. 11 shows a comparison of the raw histogram, the POD and the WBDE reconstructed particle distribution functions at these four instants. Grid sizes were $N_g = 1024$ for the WBDE estimate, and $N_g = 128$ for the two others. For $t = 40$, no noise seems to have affected the particle distribution yet, therefore a perfect denoising procedure should conserve the full information about the particle positions. Although WBDE introduces some artifacts in regions of phase space that should contain no particles at all, it remarkably preserves the global structure of the two streams. This is possible thanks to the numerous wavelet coefficients close to the sharp features in f that are above the thresholds, in contrast to the bump-on-tail case. On the next snapshot at $t = 60$, the filaments have overlapped and the system is beginning to loose its memory due to numerical round-off errors. The fastest filaments still visible on the histogram are not preserved by WBDE, but the most active regions are well reproduced. At $t = 100$, the closeness between the histogram and the WBDE estimate is striking. To put it somewhat subjectively, one may say that WBDE did not consider most of the rough features present at this stage as 'noise', since they are not removed. Only with the last snapshot at $t = 400$ does the WBDE estimate begin to be smoother than the histogram, suggesting that the nonlinear interaction between particles has introduced randomization in the system.

The POD method is able to track very well the fine and coarse scale structures of the particle density using a small number of modes. In particular, for $t = 40, 60, 100,$ and 400 only $r = 28, r = 27, r = 18,$ and $r = 5$ modes were kept. The decrease of the number of modes with time results from the lost of fine scale features in the distribution function.

4. Summary and conclusion

Wavelet-based density estimation was investigated as a post-processing tool to reduce the noise in the reconstruction of particle distribution functions starting from discrete particle data. This is a problem of direct relevance to particle-based transport calculations in plasma physics and related fields. In particular, particle methods present many advantages over continuum methods, but have the potential drawback of introducing noise due to statistical sampling.

In the context of particle-in-cell methods this problem is typically approached using finite size particles. However, this approach, which is closely related to the kernel density estimation method in statistics, requires the choice of a smoothing scale, h (e.g., the standard deviation for Gaussian shape functions) whose optimal value is not known a priori. A small h is desirable to fit as many Debye wavelengths as possible, whereas a large h would lead to smoother distributions. This situation results from the compromise between bias and variance in statistical estimation. To address this problem we proposed a wavelet-based density estimation (WBDE) method that does not require an a priori selection of a global smoothing scale and that is able to adapt locally to the smoothness of the density based on the given discrete data. WBDE was introduced in statistics [16]. In this paper we extended the method to higher dimension and applied it for the first time to particle-based calculations. The resulting method exploits the multiresolution properties of wavelets, has very weak dependence on adjustable parameters, and relies mostly on the raw data to separate the relevant information from the noise.

As a first example, we analyzed a plasma collisional relaxation problem modeled by stochastic differential equations. Thanks to the sparsity of the wavelet expansion of the distribution function, we have been able to extract the information out of the statistical fluctuations by nonlinear thresholding of the wavelet coefficients. At late times, when the particle distribution approaches a Maxwellian state, we have been able to quantify the difference between the denoised particle distribution function and its analytical counterpart, thus demonstrating the improvement with respect to the raw histogram. The POD-smoothed and wavelet-smoothed particle distribution functions were shown to be roughly equivalent in this respect. These results were then extended to a more complex situation simulated with a δf code. Finally, we have turned to the Vlasov–Poisson problem, which includes interactions between particles via the self-consistent electric field. The POD and WBDE methods were shown to yield quantitatively similar results in terms of mean squared error for a particle distribution function resulting from nonlinear saturation after occurrence of a bump-on-tail instability. We have then studied the denoising algorithm during nonlinear evolution after the two-streams instability starting from two counter-streaming cold electron beams. This initial condition violates the decorrelation hypothesis underlying the WBDE algorithm, and thus offers a good way to test its robustness regarding this aspect. The WBDE method was shown to yield qualitatively good results without changing the threshold values.

Defining a 'mode' as one term in the decomposition of the particle distribution function into a sum, it appears that the number of modes used in the POD reconstruction algorithm is considerably less than the number of modes needed using wavelets. But to make a systematic comparison when it comes to storage requirements or data compression, one has to keep in mind that POD modes are empirical and problem dependent whereas wavelet modes are known a priori. Because of this, the specification of each POD mode requires $2N_g + 1$ components whereas the specification of each wavelet component only requires one number: the amplitude of the mode. A potential limitation of the POD method is the lack of systematic a priori thresholding criteria to determine the optimal number of modes. However, the information contained in the decay of the spectrum of the singular values can be used to determine optimal ranks. For example, in Ref. [15] a criterion based on the relative rate of decay of the spectrum provided consistent good results for denoising Monte-Carlo particle data.

To test and compare the proposed denoising algorithms we have focused on the distribution function because it is in some sense the most basic and fundamental quantity in PIC and Monte-Carlo simulations. However, it would be interesting to explore the use of other fields like for example fluxes, charge densities, or derivatives of f rather than f itself. Lack of space did not allow us to elaborate more on this issue that we look forward to address in a future publication. Regarding this, it is important to keep in mind that what to choose to compare the algorithms, beyond the obvious choice of the raw f , depends on the physics of the specific problem of interest. For example, calculations that require evaluation of quasilinear fluxes in wave-particle interaction problems might benefit from denoising gradients of f . On the other hand, active denoising in Vlasov-Poisson codes might only require denoising of the charge distribution (i.e. the integral of f in velocity) whereas denoising the current (first velocity moment of f) would be important in Vlasov-Maxwell codes. Other specific needs might arise when coupling particle codes to continuum codes, e.g. when performing extended MHD calculations. It may also be of interest to focus on the reconstruction error of the force fields, which determine the evolution of the simulated plasma. These forces depend on f through integrals, and statistical analysis of the estimation of f using weak norms [39] might be beneficial.

The computational cost of our method scales linearly with the number of particles and with the grid resolution. Therefore, WBDE is an excellent candidate to be performed at each time step during the course of a simulation. Once the wavelet expansion of the denoised particle distribution function is known, it is possible to continue using the wavelet representation to solve the Poisson equation [40] and to compute the forces. The moment conservation properties that we have demonstrated in this paper should mitigate the unavoidable dissipative effects implied by the smoothing stage. In Ref. [5], a dissipative term was introduced in a global PIC code to avoid unlimited growth of particle weights in δf codes, and this was shown to improve long time convergence of the simulations. It would be of interest to assess if the nonlinear dissipation operator corresponding to WBDE has the same effect.

There are several potential extensions and applications of the techniques and results presented. Some of these include high dimensional problems, active denoising, and applications to more complex plasma models. The implementation of the WBDE algorithm to high (greater than two) dimensions is in principle straightforward. The POD method on the other hand can be more challenging since the SVD is applied to matrices. One simple way to circumvent this problem is to “fold” high dimensional data into matrices. However, this straightforward approach can be numerically inefficient. A promising alternative would be to use tensor decomposition techniques like the method of generalized low rank approximation of matrices used in [15]. The problem of active denoising, and the application of the reconstruction algorithms to more complex plasma models like Vlasov-Maxwell and gyrokinetics is a key follow up of the results presented here. Also, it would be of interest to explore applications to the problem of coupling particle and continuum codes in extended MHD and radio-frequency heating studies in fusion plasmas.

Acknowledgments

We thank Don Spong for providing the DELTA5D Monte-Carlo guiding center simulation data in Fig. 6, originally published in Ref. [15]. We also thank Xavier Garbet for his comments on the paper and for pointing out several key references, and Bill Dorland for motivating discussions. M.F. and K.S. acknowledge financial support by ANR under contract M2TFP (Méthodes Multiéchelles pour la Turbulence dans les Fluides et les Plasmas). M.F. is grateful to the Wissenschaftskolleg zu Berlin for hospitality while revising this paper. D.C.N. and G.Ch. acknowledge support from the Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the US Department of Energy under contract DE-AC05-00OR22725. D.C.N. also gratefully acknowledges the support and hospitality of the École Centrale de Marseille for the three, one month visiting positions during the elaboration of this work. This work, supported by the European Communities under the contract of Association between EURATOM, CEA and the French Research Federation for fusion studies, was carried out within the framework of the European Fusion Development Agreement. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

References

- [1] C.K. Birdsall, A.B. Langdon, *Plasma Physics via Computer Simulation*, McGraw-Hill, New York, 1985.
- [2] R.W. Hockney, J.W. Eastwood, *Computer Simulation using Particles*, IOP, Bristol, Philadelphia, 1988.
- [3] W.M. Nevins, G.W. Hammett, A.M. Dimits, W. Dorland, D.E. Shumaker, *Phys. Plasmas* 12 (2005) 122305.
- [4] J.A. Krommes, *Phys. Plasmas* 14 (2007) 090501.
- [5] B.F. McMillan, S. Jolliet, T.M. Tran, L. Villard, A. Bottino, P. Angelino, *Phys. Plasmas* 15 (2008) 052308.
- [6] J.A. Krommes, *Phys. Fluids B* 5 (1993) 1066–1100.
- [7] A.Y. Aydemir, *Phys. Plasmas* 1 (1994) 822–831.
- [8] R.W. Hockney, *Phys. Fluids* 9 (1966) 1826–1835.
- [9] A.B. Langdon, C.K. Birdsall, *Phys. Fluids* 13 (1970) 2115.
- [10] S. Brunner, E. Valeo, J.A. Krommes, *Phys. Plasmas* 6 (1999) 4504.
- [11] S. Jolliet, A. Bottino, P. Angelino, R. Hatzky, T.M. Tran, B.F. McMillan, O. Sauter, K. Appert, Y. Idomuram, L. Villard, *Comp. Phys. Commun.* 177 (2007) 409.
- [12] Y. Chen, S.E. Parker, *Phys. Plasmas* 14 (2007) 082301.
- [13] E. Cormier-Michel, B.A. Shadwick, C.G.R. Geddes, E. Esarey, C.B. Schroeder, W.P. Leemans, *Phys. Rev. E* 78 (2008) 016404.
- [14] J.L.V. Lewandowski, *Phys. Plasmas* 12 (2005) 052322.
- [15] D. del-Castillo-Negrete, D.A. Spong, S.P. Hirshman, *Phys. Plasmas* 15 (2008) 092308.
- [16] D.L. Donoho, I.M. Johnstone, G. Keryacharian, D. Picard, *Ann. Stat.* 24 (1996) 508–539.
- [17] B.W. Silverman, *Density Estimation for Statistics and Data Analysis*, Chapman and Hall, 1986.
- [18] E. Parzen, *Ann. Math. Stat.* 33 (1962) 1065–1076.

- [19] S. Tsong Chiu, *Ann. Stat.* 19 (1991) 1883–1905.
- [20] M. Farge, *Ann. Rev. Fluid Mech.* 24 (1992) 395–457.
- [21] S. Jaffard, *Publ. Mat.* 35 (1991) 155–168.
- [22] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, École Polytechnique, 1999.
- [23] A. Cohen, I. Daubechies, P. Vial, *Appl. Comput. Harmon. Anal.* 1 (1993) 54–81.
- [24] D. Donoho, I. Jonhstone, *Biometrika* 81 (1994) 425–455.
- [25] M. Farge, K. Schneider, P. Devynck, *Phys. Plasmas* 13 (2006) 042304.
- [26] R. von Sachs, K. Schneider, *Appl. Comput. Harmon. Anal.* 3 (1996) 268–282.
- [27] M. Vannucci, B. Vidakovic, *J. Ital. Stat. Soc.* 6 (1998) 15–19.
- [28] B. Vidakovic, *Statistical Modeling by Wavelets*, Wiley, 1999.
- [29] S. Gassama, E. Sonnendrücker, K. Schneider, M. Farge, M.O. Domingues, *ESAIM: Proceedings* 16 (2007) 196–210.
- [30] A. Azzalini, M. Farge, K. Schneider, *Appl. Comput. Harmon. Anal.* 18 (2004) 177–185.
- [31] I. Daubechies, *Ten Lectures on Wavelets*, SIAM, 1992.
- [32] I. Daubechies, *SIAM J. Math. Anal.* 24 (1993) 499–519.
- [33] J. Denavit, *J. Comput. Phys.* 9 (1972) 75–98.
- [34] A. Ihler, KDE toolbox for Matlab, <<http://www.ics.uci.edu/~ihler/code/kde.html>>.
- [35] N. Kingsbury, *Appl. Comput. Harmon. Anal.* 10 (2001) 234–253.
- [36] D. del-Castillo-Negrete, *Phys. Plasmas* 5 (1998) 3886.
- [37] G.H. Golub, Van D.F. Loan, *Matrix Computations*, third ed., The John Hopkins University Press, London, 1996.
- [38] S.E. Parker, W.W. Lee, *Phys. Fluids B* 5 (1993) 77.
- [39] H.D. Victory, E.J. Allen, *SIAM J. Numer. Anal.* 28 (1991) 1207–1241.
- [40] S. Jaffard, *SIAM J. Numer. Anal.* 29 (1992) 965–986.