DENSITY ESTIMATION TECHNIQUES FOR CHARGED PARTICLE BEAMS WITH APPLICATIONS TO MICROBUNCHING INSTABILITY

G. Bassi, University of Liverpool and Cockcroft Institute, Liverpool, UK *
B. Terzić, NICADD, Department of Physics, Northern Illinois University, DeKalb, IL 60115, USA

INTRODUCTION

We present two new techniques to represent particle distributions in beam dynamics simulations. As an application we consider 2D simulations of coherent synchrotron radiation (CSR) with the code developed in [1]. In the current version of the code, a Monte Carlo particle (MCP) method is implemented to represent the particle distribution sampled by N point-charge particles based on a truncated cosine expansion. The particle distribution is then evaluated on a finite grid and stored for computation of retarded potentials. The alternative representations outlined here are both grid-based, so the first step in their implementation is a particle deposition on a finite grid. The resulting gridded representation is fairly noisy and therefore not very competitive in terms of accuracy (about an order of magnitude less accurate than the cosine expansion), but nearly three orders of magnitude more efficient.

The first alternative technique uses Fast Cosine Transforms (FCT) to transform the gridded distribution to Fourier space, where it truncates higher order frequencies by retaining only a fraction of the cosine coefficients. This non-discriminatory removal of the small scale structure results in smoothening of the distribution, and represents an overly simplistic method for noise removal. The truncated FCT technique is equivalent to the Monte Carlo cosine expansion in terms of accuracy, but about *three orders of magnitude more efficient*.

The second alternative technique uses Discrete Wavelet Transforms (DFT) to transform the gridded distribution to wavelet space, where it performs *wavelet thresholding*, thereby largely removing the numerical noise intrinsic in numerical simulations. Transforming back to the physical space yields an approximation to the particle distribution which is *more accurate* than the cosine expansion – as quantified by the L_2 norm of the error and the Signal-to-Noise Ratio (SNR) – and still *about three orders of magnitude faster*.

PARTICLE DEPOSITION AND DENSITY ESTIMATION

The MCP method implemented in [1] is a meshless method that gives a smooth particle distribution at the price of a costly evaluation of the Fourier coefficients. Here we propose a mesh method based on particle charge deposition, in the spirit of traditional PIC codes. The accuracy of the grid representation scales as the square root of the number of particles per cell. Taking it a step further, by properly understanding the nature and profile of the noise associated with the particle deposition scheme, we propose the use of a wavelet basis – which provides a natural way to efficiently separate noise from signal – to denoise and smoothen the grid distribution and compare it to the FCT method.

Filtered Fast Cosine Transform

We introduce the alternative cosine expansion which is considerably faster then the MCP method currently used in [1]. The algorithm can be outlined as follows:

1. Deposit particles on the (N_x, N_z) grid.
2. Apply 2D FCT on the grid, thus yielding (N_x, N_z) cosine coefficients.
3. The high-frequency contribution to the density is then removed by filtering (truncating) coefficients higher than N_{cx} and N_{cz} in x- and z-coordinate, respectively. This removal of the high-frequency components results in a smoother distribution, but it also removes small scale structures that may not be due to noise. The truncation of the Fourier coefficients also restrict spatial resolution of the representation.
4. Apply 2D inverse FCT on the grid, to obtain the smoothed distribution in physical space.

Wavelet-Denoised Density

The wavelet-denoised algorithm for estimating particle density can be outlined as follows:

1. Deposit particles on the (N_x, N_z) grid.
2. Apply Anscombe transformation to convert the signal polluted by Poissonian noise (the gridded particle distribution) to the signal with Gaussian noise.
3. Apply 2D DWT on the grid, thus yielding (N_x, N_z) wavelet coefficients.
4. Perform wavelet thresholding on the wavelet coefficients in order to remove numerical noise and smoothen the particle distribution. Unlike the filtering of the cosine coefficients in Fourier space, this *judicious* noise removal *does not* restrict the spatial resolution of the distribution: small-scale structures which are not deemed to be noise during wavelet thresholding (*i.e.*, the amplitude of the corresponding wavelet

* gabriele.bassi@stfc.ac.uk

coefficients is larger than the threshold) are retained in the denoised distribution.

5. Apply 2D inverse DWT on the grid, to obtain the smoothed distribution in physical space.

6. Apply inverse Anscombe transformation.

FAST COSINE TRANSFORM VS DISCRETE WAVELET TRANSFORM

We now compare the accuracy and efficiency of the wavelet-denoised grid distribution and the filtered fast cosine approximation considering the initial distribution used in [1] to study microbunching instability, see Fig. 1 (top left). The amplitude of the initial modulation is $A = 0.05$ and the wavelength $\lambda = 100\mu$m. All the results shown here are with $N_z = 1024$, $N_x = 128$ and $N = 10^8$ particles. The number of cosine basis functions (highest order) in the Fourier expansion is $N_{cx} = 40$ and $N_{cz} = 100$, which determines the smallest structure representable by a finite cosine approximation.

The argument of the highest order basis functions is

$$N_{cx}\pi\bar{x} = \frac{N_{cx}\pi x}{L_x} + \phi_x, \quad N_{cz}\pi\bar{z} = \frac{N_{cz}\pi z}{L_z} + \phi_z, \quad (1)$$

where ϕ_x and ϕ_z are constants (and therefore just simple phase shifts). Therefore, the smallest wavelengths representable by this expansion are

$$\lambda_x^{\min,\,\cos} = \frac{2L_x}{N_{cx}}, \quad \lambda_z^{\min,\,\cos} = \frac{2L_z}{N_{cz}}, \quad (2)$$

The cosine approximation will be able to accurately approximate only particle distributions with small-scale structure of wavelengths λ (of the order of the initial modulation wavelength) larger than the smallest wavelength representable by the cosine expansion $\lambda_z^{\min}$. This imposes a limit on usability of the cosine expansion

$$\bar{\lambda} > \lambda_z^{\min,\,\cos} = \frac{2L_z}{N_{cz}}, \quad N_{cz} > \frac{2L_z}{\bar{\lambda}} = \frac{7200}{\lambda} \quad (3)$$

Figure 1 (top center) shows the accuracy in approximating the distribution. The relation (3) predicts that the precipitous drop in accuracy (increase in error) for the cosine approximation with $M = 100$ should occur for $\lambda < 7200/N_{cz} = 72$ which is exactly what is observed.

The smallest structure representable by a finite grid is determined by the grid's resolution. The smallest wavelength representable on the grid is simply two spacings of the grid (three gridpoints), describing the "saw-tooth" structure $\lambda_z^{\min,\,\mathrm{grid}} = 2h_z = 2L_z/N_z$ The smallest structure representable on the grid is a simple "saw tooth", which has a lengthscale of 4 grid spacings. This means that the requirement for the accurate representation of the distribution on the finite grid is then given by

$$\bar{\lambda} > \lambda_z^{\min,\,\mathrm{grid}} = \frac{4L_z}{N_z}, \quad N_z > \frac{4L_z}{\bar{\lambda}} = \frac{3600}{\lambda}. \quad (4)$$

For the simulations in Fig. 1 (top center), where $N_z = 1024$, the relation (4) predicts that for $\lambda < 3600/N_z \approx 14$, the grid approximation will be inaccurate, which is exactly what is implied by the increase in the error of the grid approximation.

Therefore, the spatial resolution of the cosine expansion is determined by the number of expansion coefficients N_{cx} and N_{cz}, while for the grid approximation it is given by the number of gridpoints N_x and N_z.

Figure 1 (top center) shows the error E for the cosine, regular grid and wavelet-denoised grid approximations as a function of the number of particles N used in the simulation. It is evident that the square of the L_2-norm of the error (E) for all three approximations scales as N^{-1}, consistent with the well-known finding that the signal quality (as measured by the normalized inverse of the L_2-norm of the error, also known as the signal-to-noise ratio) scales as $N^{1/2}$ with the number of particles. It is also clear that the wavelet-denoising significantly improves the accuracy of the approximation, surpassing that of the cosine expansion. The increase in accuracy when using the denoised grid approximation is even more pronounced for other values of the modulation wavelength λ (see Fig. 1 (top center)).

Figure 1 (bottom left) shows a small part of the cross-section of the particle distribution approximated with the various schemes. It is visually evident that the wavelet thresholding smoothens the small-scale noise present in the gridded distribution, while maintaining high-fidelity signal.

Figure 1 (bottom center) shows how the execution times of the different methods and their constituent parts scale with the number of particles N. The Monte Carlo-based computation of cosine coefficients requires integration over N, therefore scaling as $\propto \mathcal{O}(N_{cx}N_{cz}N)$. Using cosine coefficients to approximate the particle distribution on the grid requires summing over all the cosine coefficients, thus scaling as $\propto \mathcal{O}(N_{cx}N_{cz}N_{\mathrm{grid}})$. Since, in realistic simulations, $N \gg N_{\mathrm{grid}}$, the computation of cosine coefficients will be by far the most time-consuming of the two.

Both grid-based methods use particle deposition, which scales as $\propto \mathcal{O}(N)$, and a fast transform FCT and DWT which scale as $\propto \mathcal{O}(N_{\mathrm{grid}}\log(N_{\mathrm{grid}}))$ and $\propto \mathcal{O}(MN_{\mathrm{grid}})$, respectively, where M is the width of the wavelet family. This means that for large N used in realistic simulations, the most time-consuming part of the approximation is a particle deposition (Fig. 1 (bottom center)). This is also why the execution times of the grid methods become quite similar for large N (Fig. 1 (bottom right)).

In Fig. 1 (bottom right), the three methods are directly compared, by plotting the execution time for the two alternative methods scaled with the execution time of the Monte Carlo cosine approximation as a function of N. It is evident that the ratio of execution times asymptotically approaches $1/(N_{cx}N_{cz})$, since Monte Carlo cosine $\propto \mathcal{O}(N_{cx}N_{cz}N)$ and the grid-based methods (for large N used in realistic simulations) as $\propto \mathcal{O}(N)$.

An honest comparison of the efficiency of cosine expansion and the wavelet-denoised grid expansion has to take

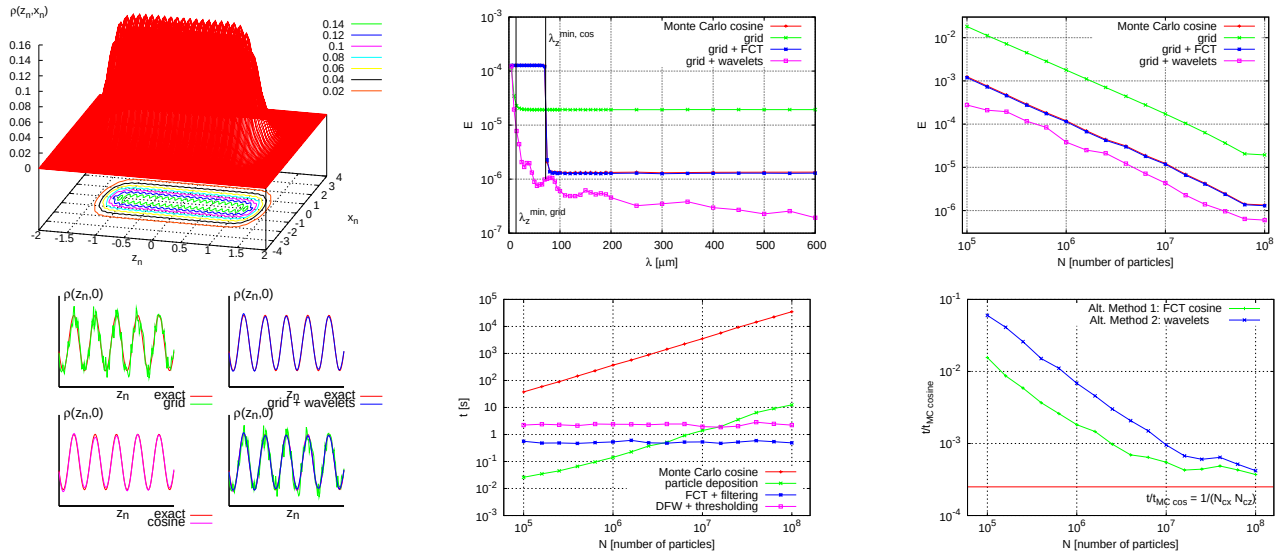

Figure 1: Top left: normalized spatial density with flattop longitudinal profile [1]. Top center: accuracy of FFT-based cosine expansion (red line), simple grid deposition (green line), grid-based truncated FCT (blue) and a wavelet-denoised grid deposition (purple). Vertical lines denote minimal wavelengths representable with grid and cosine approximations (from left to right). Wavelet family used is Daubechies of order 10. Top right: error E as a function of N for the approximations: Monte Carlo cosine (red); grid (green), FCT on a grid (blue) and grid with wavelet denoising, using Daubechies wavelets of order 10 (purple). Bottom left: cross-section ($x = 0$) of the approximations of the charged particle distribution on the grid, compared the the exact distribution: simple grid deposition approximation (top left); wavelet-denoised grid deposition approximation (top right); cosine approximation (bottom left); and all three superimposed. Wavelet family used is Daubechies of order 10. Bottom center: execution time of different approximations and their parts, as a function of N. Wavelet family used is Daubechies of order 2. Bottom right: direct comparison of efficiency of the three techniques: execution times of the two alternative techniques scaled with the execution time of the Monte Carlo cosine approximation as a function of N. Horizontal red line denotes $t/t_{MCcosine} = 1/(N_{cx}N_{cz})$, to which the ratios tend. Wavelet family used is Daubechies of order 2.

into account the optimum number of basis functions in the cosine expansion. The execution times for the computation of cosine coefficients and grid deposition methods is given by $t_{\mathrm{MC\,cos}} = k_1 N_{cx} N_{cz} N$, $t_{\mathrm{grid}} = k_2 N$ where k_1 and k_2 are some constants which depend on the computer executing the simulation. Therefore, the ratio of the execution times is simply $r \equiv t_{\mathrm{MC\,cos}}/t_{\mathrm{grid}} = k N_{cx} N_{cz}$, where the constant k, and therefore the ratio r should be computer-independent (as it can be seen from Fig. 1 (bottom right), it is quite close to unity). For the near-optimal number of basis function in the cosine expansion along the longitudinal direction $\bar{N}_{cz} = 7200/\lambda$, the ratio is $\bar{r} = k N_{cx} \bar{N}_{cz} = k N_c N_{cz} \bar{N}_{cz}/N_{cz}$. From our earlier analysis and Fig. 1 (bottom right), we get that for large N and $N_{cz} = 100$, $r \approx N_{cx} N_{cz}$ (ignoring the contributions from the cosine evaluation on the grid and wavelet denoising, which are negligible for large N) so $\bar{r} \approx 7200/\lambda N_{cx}$, which means that the cost-effectiveness of the wavelet-based grid over the cosine expansion will be even more pronounced for simulations of small-scale perturbations, as expected. For example, this means that for $N_{cx} = 40$, the increase in computational efficiency for $\lambda < 100$ should be at least a factor 2880, and the double at $\lambda < 50$.

DISCUSSION AND CONCLUSION

We discussed approximation schemes for charged particle distributions with applications to beam dynamics. We found that the wavelet method provides a superior alternative to the Monte Carlo cosine expansion currently in use in [1] and to the fast cosine expansion. The new method deposits particles on a grid – an approximation to the charged particle distribution which is intrinsically plagued by numerical noise – and then uses wavelet thresholding to remove judiciously the noise. The resulting wavelet-based grid approximation is appreciably more accurate – quantified by the square of the L_2-norm of the error – and significantly more efficient (as measured by execution times) by a factor of 5-150, than the Monte Carlo method.

In the future, it will be very beneficial to explore possibilities of further optimizing the computational efficiency of the new wavelet-based approach by exploiting the sparsity of the charge representation. Storing the entire history of the charge distribution as a small set of sparse wavelet coefficients instead of the full grid (2D in space and 1D in time) could provide a substantial savings in memory and processor communication, thus significantly reducing simulation times. This study is currently underway.

REFERENCES

[1] Bassi G, Ellison J.A, K. Heinemann and R. Warnock. These proceedings, paper TU1PBIO3

[2] Terzić B., Pogorelov I.V. and Bohn C.L. 2007, Phys. Rev. ST Accel. Beams 10, 034201.