Parallelization of particle codes for the simulation of Alfvénic turbulence: Gridless Finite Size Particle versus PIC approach

S. Briguglio[a], G. Vlad[a], B. Di Martino[b] and G. Fogaccia[a]
[a] Associazione Euratom-ENEA sulla Fusione, C.R. Frascati,
C.P. 65 - I-00044 - Frascati, Rome, Italy
[b] Dip. Ingegneria dell'Informazione, Second University of Naples, Italy

One of the most convenient technique for parallelizing *particle-in-cell* (PIC) codes on distributed-memory architectures is represented by the so-called *particle decomposition*. It corresponds to replicating the whole spatial domain on each processor, while distributing the particle population. Different from *domain decomposition*, there is no load-balancing problem, because no particle migrates from one processor to another. Unfortunately, the memory request associated to the grid quantities (equilibrium and fluctuating fields), which are replicated on each processor, gives rise to a bottle-neck on the scalability of physical resolution with processors. Such a constraint is strongly penalizing and motivates a large effort in looking for alternative methods able to reduce the replicated-array memory offset.

In many concrete situations, as a matter of fact, periodic spatial domains are considered, and the problem admits solution in terms of normal modes. This feature allows to solve the equations for the electromagnetic fields in Fourier space. Moreover, in order to get a deeper insight in the physical mechanisms underlying the observed phenomena, the analysis is often focused on the evolution of relatively few harmonics. This is the case of many studies in the field of controlled nuclear fusion research.

Although these facts do not cause the resolution request to be relaxed (one is typically interested in the high mode-number portion of the complete fluctuation spectrum), they offer a different path toward the target of an efficient parallelization and, hence, toward the high-resolution simulations: namely, the gridless *finite-size-particle* (FSP) approach. Instead of transforming the fields, from Fourier to real space, on a spatial grid, then interpolating such fields at each particle position to evolve particle coordinates, the gridless FSP algorithm directly transforms the fields at the particle position. Correspondingly, particle contributions to the pressure perturbation are just Fourier-transformed and summed together, instead of being collected at the nearest grid points. In such a way, the memory resources demanded by the field storage are greatly reduced. Moreover, from the point of view of parallel effciency, such a method is more convenient than the PIC one, because it also reduces the amount of replicated calculations assigned to each computational node. This can be seen from Fig. 1, where the results obtained by the gridless FSP and PIC versions of the Hybrid MHD-Gyrokinetic Code (HMGC) – a code for the investigation of Alfvénic turbulence in tokamak plasmas – are compared. The efficiency η, defined as the speed-up factor divided by the number of processors, is plotted versus the ratio, n_{proc}/N_{ppc}, between the number of processors and the average number of particle per cell (a "virtual" cell, in the gridless case).

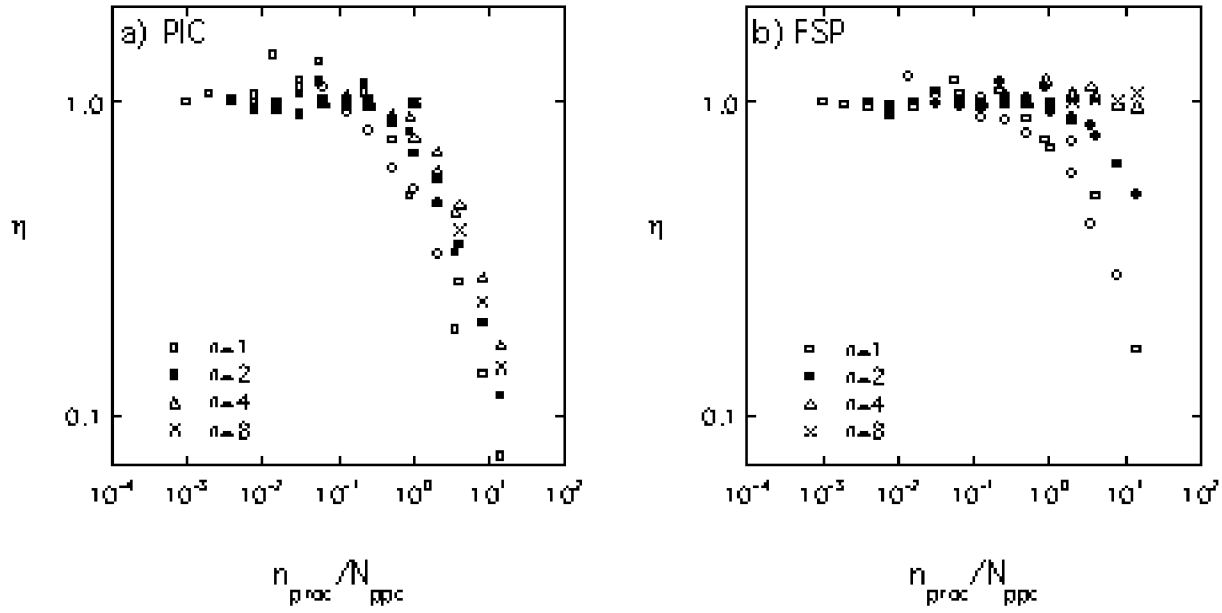

Figure 1: Efficiency η, defined as the speed-up factor divided by the number of processors, versus n_{proc}/N_{ppc}. Simulations retaining fluctuations with a single toroidal mode number are considered. The results obtained by PIC HMGC (a) are compared with those obtained by FSP HMGC (b), for several different choices of n.

Simulations retaining modes with a single toroidal number n are considered here, for several cases corresponding to different choices of n. It is possible to note that, for each value of n, the efficiency essentially maintains its ideal value ($\eta \approx 1$) up to a certain threshold value in n_{proc}/N_{ppc}. Such a value comes out to be higher in the FSP case than in the PIC one, as expected, because of the lesser amount of replicated calculations that characterizes the FSP code. The most relevant features, here, is represented by the fact that, while for the PIC code the n_{proc}/N_{ppc} threshold value exhibits a negligible dependence on n, for the FSP one such a dependence is positive and much stronger. This fact, which can be understood by considering that the threshold value is approximately obtained by imposing that the replicated calculations exceed the distributed ones, implies that the advantages of the FSP method in terms of parallelization efficiency become more apparent with increasing n. We can then conclude that the particle-decomposition parallelization of FSP codes is potentially efficient even for massively parallel architectures, not only for simulations characterized by high resolution in the velocity space (high N_{ppc} values), but also – different from the same parallelization of PIC codes – for simulations with high resolution in the real space (high n values).

The FSP method can be advantageous also concerning the absolute CPU-time performances, as far as the number of processors exceeds a certain threshold. This is due to the fact that, although the FSP method requires much heavier calculations for each particle (the electromagnetic fields are transformed back from the Fourier space directly at the particle position), such calculations are distributed among processors. If the FT algorithm is adopted, instead of the

FFT one, in order to save memory space, in the PIC simulations, it can be shown that the threshold in the number of processors is essentially given by the average number of particle per cell, N_{ppc}, and is easily overcome, even in moderately parallel simulations. This is confirmed by the experimental results shown in Fig. 2, where the Central Processor Unit (CPU) *user time* re-

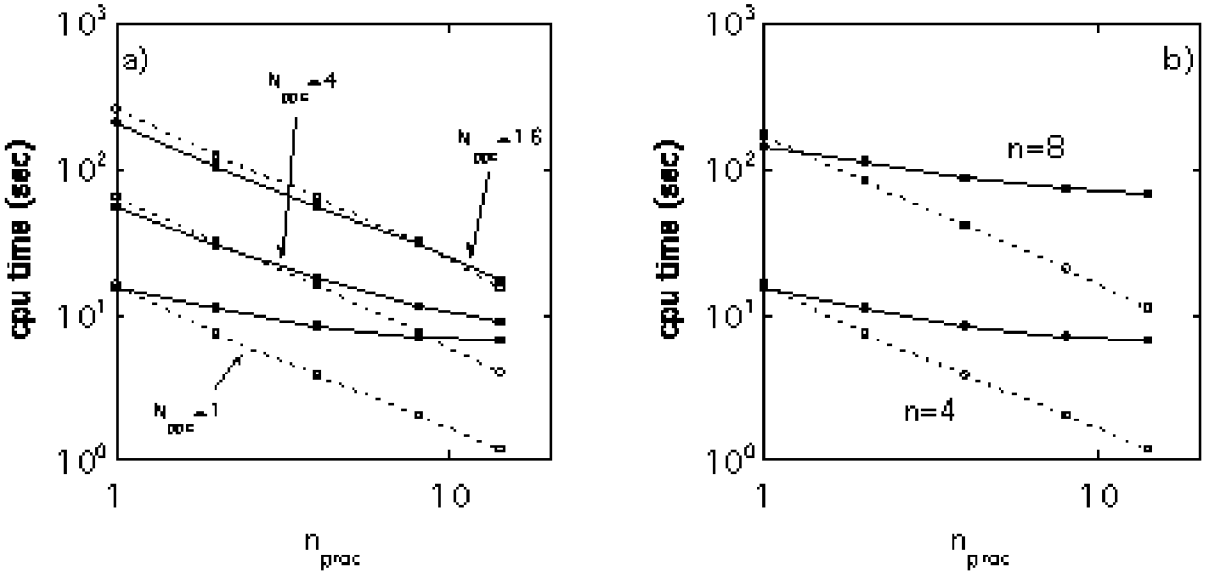

Figure 2: CPU *user time* per time step versus n_{proc}. Results obtained by FSP (empty symbols) and PIC (full symbols) simulations with fixed mode number ($n = 4$) and different values of N_{ppc} (a), or fixed velocity-space resolution ($N_{ppc} = 1$) and different values of n (b) are compared.

quired by each simulation time step is plotted versus the number of processors. Results obtained by FSP (empty symbols) and PIC (full symbols) simulations with fixed mode number ($n = 4$) and different values of N_{ppc} (a), or fixed velocity-space resolution ($N_{ppc} = 1$) and different values of n (b) are compared. It can be seen that the threshold value is approximately equal to N_{ppc} and essentially independent on n, in agreement with the theoretical prediction. If the more efficient FFT algorithm were instead adopted, the threshold would also depend on the number of harmonics retained in the simulation and, more weakly, on the required spatial resolution. In concrete situations, the number of processors can easily come out to be below this threshold if the full set of harmonics corresponding to a given spatial resolution is retained. If, however, for the sake of simplifying the investigation, only a limited number of harmonics is retained, without reducing the resolution requirement, the threshold value sensibly decreases, and the FSP method can prevail also over the FFT-PIC one. In practice, the resort to such simplified *few-harmonic* simulations is often forced by the limitations on the memory available on each computational node.

Linear and nonlinear properties of moderate-toroidal-number (n) shear-Alfvén modes in tokamaks have been investigated by using FSP HMGC. The existence of unstable Toroidal Alfvén Eigenmodes (TAE's) and their kinetic counterpart has been shown for low values of the energetic-ion pressure gradient. Above a certain threshold value, the Energetic Particle continuum Mode

(EPM) is destabilized, with growth rate fast increasing with increasing energetic-particle pressure gradient.

Nonlinear TAE saturation appears to be due to the trapping of resonant energetic ions in the potential well of the wave. Saturation of the EPM is associated to a macroscopic outward displacement of the energetic-particle distribution, with potentially dramatic consequences on α-particle confinement. Our simulations indicate that, for circular magnetic surface equilibria, the EPM excitation threshold in β_H (with β_H being the ratio between the energetic-ion pressure and the magnetic pressure) has a minimum, versus the toroidal number n, at values $n \approx 10$, as shown in Fig. 3. The case of a plasma with aspect ratio $R_0/a = 10$ (with R_0 and a being, respectively, the major and the minor radius of the torus) and energetic-particle Larmor radius ρ_H such that $\rho_H/a = 0.01$ is considered here. A density profile of the form $\exp\left[-\left(r^2/L_n^2\right)^{\alpha_n}\right]$ has been assumed for the energetic ions, with $a^2/L_n^2 = 2$ and $\alpha_n = 2$. Results obtained by FSP (empty symbols) and PIC (full symbols) simulations are compared. It should be noted here that parallel FSP simulations allow to push the investigation up to higher values of n. The existence of a minimum in the threshold could not be appreciated, with the given size of the single node Random Access Memory, in the framework the PIC approach.

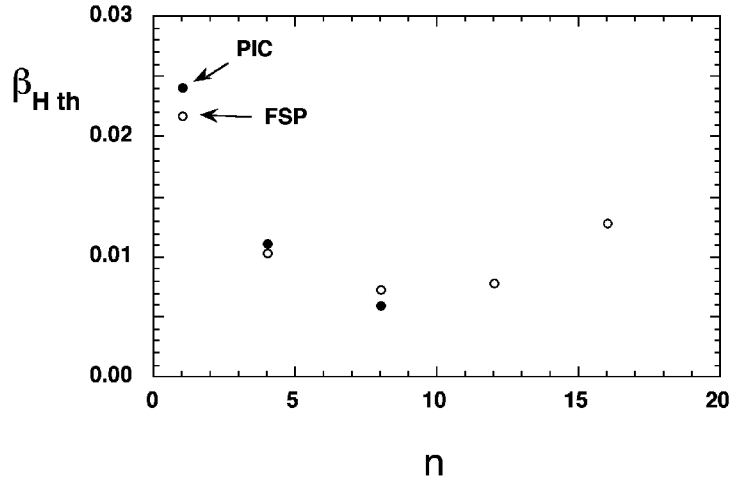

Figure 3: Threshold β_H values for the linear EPM destabilization at different values of the toroidal number n of the mode. Results obtained by FSP (empty symbols) and PIC (full symbols) simulations are compared.

Because of the correlation between the linear excitation of EPM and the strong nonlinear redistribution in the energetic particle source, it becomes a crucial issue to extend the analysis to realistic equilibria and investigate whether such a minimum can become comparable with the values which are relevant for ignited plasmas.