Multi-level Parallel-In-Time Methods for N-body Plasma Physics Applications

(Prospect on upcoming work)

M. Winkel*, R. Speck*†, D. Ruprecht†, P Gibbon*
*Jülich Supercomputing Centre, Jülich, Germany.
†Institute of Computational Science, Università della Svizzera italiana, Lugano, Switzerland.

I. Introduction

Parallel-in-time integration methods are a promising approach to extend strong-scaling limits of simulation code at large-scales. With fixed degrees-of-freedom in space and million-way concurrency at hand, classical space-parallel strategies inevitably lead to inefficient use of compute power as strong-scaling saturates. In this contribution, we will describe the ongoing transition from a large-scale, space-parallel simulation code towards a space-time parallel method ready for extreme-scales. To this end, we will join the space-parallel Barnes-Hut tree code PEPC [1] with the time-parallelization library PFASST [2].

II. Limits of parallelization using multipole-based fast summation schemes

Being dominated by more than 20 Petaflop-class installations, the Top500 list of world's most powerful computers is the evidence of a tremendous upgrade in worldwide computing power.However, this does not solve two very principal problems of all particle-based molecular dynamics simulations. Following Newtonian dynamics, the particle's movement in these simulations is driven by forces from interaction with each of the individual particles. For example, in the electrostatic case, these are mediated through the Coulomb force

$$\vec{F}(\{\vec{r}_i,q_i\},\{\vec{r}_{1\dots N},q_{1\dots N}\}) = q_i \sum_{\substack{j=1\\ j\neq i}}^{N} q_j \frac{\vec{r}_i - \vec{r}_j}{|\vec{r}_i - \vec{r}_j|^3} + \vec{F}^{(\mathrm{ext})} \qquad (1)$$

acting from N particles $\vec{r}_{1\dots N} \in \mathbb{R}^3$ with charges $q_{1\dots N} \in \mathbb{R}$ onto particle i with charge q_i at $\vec{r}_i$. In addition, usually external forces $\vec{F}^{(\mathrm{ext})} \equiv \vec{F}^{(\mathrm{ext})}(\vec{r}_i,\vec{v}_i,t)$ are present that – besides the particle's position $\vec{r}_i$ – can in general also depend on its velocity $\vec{v}_i = \dot{\vec{r}}_i$ (e. g. Lorentz force on particles in an external magnetic field) and time t (e. g. in laser-driven systems).

A direct evaluation of (1) for all particles $i = 1 \dots N$ exhibits a computational complexity of $\mathcal{O}(N^2)$ that for ever increasing simulation size and thus particle numbers cannot be dealt with by means of simply increasing computation power. Instead, efficient fast summation schemes such as the Barnes-Hut tree code [3] and the Fast Multipole Method [4] have been developed. They rely on a multipole-based approach to replace the interactions between individual particles with interactions between particles and particle groups via their multipole expansion. Thus, the contribution of distant particles is approximated by systematically grouping them into clusters and representing their charge distribution through a series expansion. The coarseness of the spatial approximation can be adjusted through a so-called multipole-acceptance-parameter that determines which sizes of particle groups at which distance are to be accepted as interaction partners. Thus, an overall complexity of $\mathcal{O}(N \log N)$ or even $\mathcal{O}(N)$ can be achieved [5].

During the last years, a highly scalable hybrid parallelized implementation of the original Barnes-Hut tree algorithm using MPI and Pthreads has been developed at Jülich Supercomputing Centre under the label *Pretty Efficient Parallel Coulomb solver* (PEPC) [1].

Today, it is not only limited to pure Coulomb interaction but includes numerous interaction-specific backends and respective frontends. Applications are reaching from simulation of interstellar dust dynamics, plasma boundary layer instabilities, laser-driven plasmas and plasma-wall interactions in fusion vessels to hydrodynamics applications with the SPH or the vortex-particle method.[1] PEPC uses a work separation scheme between one communicator thread performing rapid low-latency point-to-point communication for tree node exchange and an adjustable number of worker threads that are responsible for tree structure interpretation and force evaluation on every shared-memory node. For large systems with $N = 32 \dots 2{,}048$ million particles, efficient scalability across $294{,}912$ processors of the Blue Gene/P system JUGENE at JSC has been shown recently [1]. However, the inevitably synchronized collective tree construction phase as well as point-to-point communication during tree traversal are becoming dominant as the number of particles N per processor core C falls below $\frac{N}{C} \lesssim 2{,}000$. Thus, for smaller problem sizes, the algorithm's efficient operating point is limited to small supercomputer partitions, leading to a strong lower limit for the time-to-solution that is achievable. This strong-scaling limit is the second issue experienced by all parallel particle-based simulation methods. Their parallel scaling is intrinsically constrained and saturates as soon as the number of particles per processor is becoming too low.

III. Connecting PEPC to PFASST

Instead of decreasing $\frac{N}{C}$ any further, the additional direction of parallelization offered by parallel-in-time methods provides an appealing possibility to further reduce time-to-solution.

In a previous study, we already applied the *Parallel Full Approximation Scheme in Space and Time* (PFASST) [6], [2] to PEPC, making use of two different levels of spatial discretization [7], [8].

PFASST is an iterative parallel-in-time solver which bases on a coupling of Parareal iterations [9] with the sweeps of a multi-level spectral deferred correction method (MLSDC, see [10], [11]). Similar to Parareal, it profits from decreasing the runtime of the coarse scheme by not only coarsening in time but also using a coarser spatial discretization. PFASST further improves this strategy by using a full approximation scheme (FAS) technique to recycle coarse grid information efficiently and to increase the accuracy of the coarse grid SDC sweeps. While coarsening in time can be easily obtained in PFASST by using less intermediate collocation points, efficient and effective strategies for spatial coarsening are an area of active research.

With its multipole-acceptance-parameter, the Barnes-Hut approach as used in PEPC already offers a convenient possibility to choose between fast, inaccurate force computation for the serial initial value propagation and a precise but slow evaluation for the parallel fine propagator. This strategy is similar to the operator-based coarsening strategy described in [11]. Another approach sketched in [11] but not investigated there is the implementation of *different physics* on both levels, which is envisaged in the current project. For the coarse level

[1] http://www.fz-juelich.de/ias/jsc/pepc

Workshop "Innovative space-time-parallel methods: Analysis and Applications"
Manchester, UK, June 18-19, 2013.

an adiabatic approximation, i. e. decoupling of electronic and ionic motion in plasma simulations, seems natural at least for moderate and weak coupling between the species. Thus, the ions will for example move in front of an electronic background with reduced dynamics. On the fine level, the full electron dynamics will be taken into account. For PEPC, this will go beyond the previously used simple approach and has the potential to further promote parallel molecular dynamics simulations at their strong scaling limit. Thus it will be possible to study previously strictly separated time scales in a consistent model within one simulation.

IV. Coarsening via reduced physics for plasma-physical applications

The field of plasma physics comprises a broad range of physical configurations as the presence of free charges is common to a multitude of physical regimes. From the cold and ultra-thin interstellar matter to high densities in bulk metals and high temperatures in fusion reactions, densities of 10^{-5} to 10^{25} particles per cm^3 and temperatures from 10^{-2} to $10^{10}\,\mathrm{K}$ are covered.

The relevance of mutual interactions in the system is characterized by the coupling parameter $\Gamma = \frac{E^{(\mathrm{Coul})}}{E^{(\mathrm{therm})}}$, which is the ratio of (Coulomb) interaction energy $E^{(\mathrm{Coul})} = \frac{q^2}{4\pi\varepsilon_0 \cdot d}$ to thermal energy $E^{(\mathrm{therm})} = k_\mathrm{B} T$. Here, $d = \left(\frac{4}{3}\pi n\right)^{-\frac{1}{3}}$ is the average interparticle distance, n the particle's number density, T their temperature, ε_0 the vacuum permittivity, and k_B Boltzman's constant.

For systems of low density or high temperature, i. e. $\Gamma \ll 1$, the system can be treated as an ideal gas. Approaching $\Gamma \lesssim 1$, a rigorous perturbation theoretical treatment that includes particle interactions in lower orders has been proven successful, see e. g. [12] for a review. However, for stronger coupling $\Gamma \gtrsim 1$, matter is approaching fluid and finally solid state and these methods fail as – with the exception of regular lattices – the emerging correlated structure cannot be expressed analytically.

In this regime, numerical methods such as molecular dynamics simulations are today's methods of choice to proceed towards studying cold and dense plasmas. However, the constituent particles – usually negatively charged electrons and positively charged ions – do not only differ in charge but also carry considerably different masses. This introduces very different dynamic time scales as the mass ratio between electrons and protons $\frac{m_\mathrm{p}}{m_\mathrm{e}} \approx 1836$ is a lower bound for $\frac{m_\mathrm{ion}}{m_\mathrm{e}}$. For accurately resolving dynamics on all relevant scales, short physical time steps and thus long total simulation times are equally necessary.

In order to further extend the approach of combining the large-scale spatial parallelization of PEPC with the parallel-in-time method PFASST, we want to make use of this vast dynamic range of time scales for the particle species by decoupling electronic and ionic motion. To this end, the coarse level will fully resolve slow processes while fast dynamics is treated in an averaged or simplified model. The fine level will include full dynamics for slow and rapidly moving particle species.

Among others, the following three applications seem to provide an ideal setup for upcoming studies:

A. Laser-driven ion acceleration. In laser-driven ion acceleration, usually charged particles are expelled from a micrometer-sized metallic target by a strong external field. Using an appropriate configuration of laser and target parameters, good directionality of the escaping particles can be achieved. The interaction between travelling ions and electrons is expected to keep the particle bunches compact after leaving the target. Beam propagation outside the target requires long-running open-boundary simulations that have to resolve the dynamics of the electronic and the ionic subsystem, which is currently impossible for the real physical mass ratio given above. Here, electrons can for example be treated as a single charge cloud on the coarse level.

B. Instabilities at a plasma-vacuum boundary for magnetized systems. Due to thermally driven charge separation at a plasma-vacuum interface and the $\vec{E} \times \vec{B}$ drift resulting from the restoring motion, instabilities can emerge at a plasma-vacuum interface in an external magnetic field. These are currently studied using a two-dimensional simulation with PEPC with the constraint of resolving the gyromotion of electrons only as the period of the cyclotron motion scales with the particles mass. To sufficiently include ionic gyration that will presumably have considerable influence onto the instability pattern structure, much longer simulation times than currently reachable would be necessary. Here, on the coarse propagator level a guiding-center movement for the rapid electron gyration can be used to allow for larger timesteps and progress of the ionic motion. On the fine level, fully resolved dynamics will include electronic gyration.

C. Optical and transport properties in strongly correlated media. Optical properties such as the dielectric function that finally describes propagation of electromagnetic waves inside matter is primarily governed by the rapidly moving quasi-free electrons. However, their trajectories and collision rates strongly depend on the ionic structure and the correlations between electrons and ions. The preparation of strongly correlated, i. e. cold high-density systems in simulations is difficult, though. The formation of long-range structures at phase transitions from gaseous via liquid to bulk state is a slow process as the ions involved are rather heavy. Again, using a parallel-in-time approach for finally resolving both – fast and slow dynamics – will improve our understanding of these phase transitions, the resulting structures and finally optical properties of strongly correlated media.

The projected combination of PFASST and PEPC with spatial coarsening via a *reduced physics approach* will not only boost these plasma physics applications. It will also be an important cornerstone towards efficiently utilizing the ever-increasing computational power of today's supercomputers and their even more powerful successors and will thus pave the ground for even shorter time-to-solution.

References

[1] M. Winkel, R. Speck, H. Hübner, L. Arnold, R. Krause, and P. Gibbon, "A massively parallel, multi-disciplinary Barnes-Hut tree code for extreme-scale N-body simulations," *Comput. Physi. Commun.*, vol. 183, no. 4, pp. 880–889, Apr. 2012.

[2] M. Emmett and M. L. Minion, "Toward an efficient parallel in time method for partial differential equations," *Comm. App. Math. and Comp. Sci.*, vol. 7, pp. 105–132, 2012.

[3] J. E. Barnes and P. Hut, "A hierarchical O(N log N) force-calculation algorithm," *Nature*, vol. 324, no. 6096, pp. 446–449, Dec. 1986.

[4] L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *J. Comput. Phys.*, vol. 73, no. 2, pp. 325–348, Dec. 1987.

[5] S. Pfalzner and P. Gibbon, *Many-Body Tree Methods in Physics*. Cambridge: Cambridge University Press, 1996.

[6] M. L. Minion, "A hybrid parareal spectral deferred corrections method," *Comm. App. Math. and Comp. Sci.*, vol. 5, no. 2, pp. 265–301, 2010.

[7] R. Speck, D. Ruprecht, R. Krause, M. Emmett, M. L. Minion, M. Winkel, and P. Gibbon, "A massively space-time parallel N-body solver," in *Proc. of the Int. Conf. on High Performance Computing, Networking, Storage and Analysis*, ser. SC '12, Salt Lake City, Utah. Los Alamitos, CA, USA: IEEE Computer Society Press, 2012, pp. 92:1–92:11.

[8] ——, "Integrating an N-body problem with SDC and PFASST," in *21st Int. Conf. on Domain Decomp. Meth. [in press]*, Rennes, France, 2012.

[9] J.-L. Lions, Y. Maday, and G. Turinici, "A "parareal" in time discretization of PDE's," *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, vol. 332, pp. 661–668, 2001.

[10] A. Dutt, L. Greengard, and V. Rokhlin, "Spectral Deferred Correction Methods for Ordinary Differential Equations," *BIT Numerical Mathematics*, vol. 40, no. 2, pp. 241–266, 2000.

[11] M. Bolten, M. Emmett, R. Krause, M. Minion, D. Ruprecht, and R. Speck, "A multi-level spectral deferred correction method," (In preparation).

[12] H. Reinholz, "Dielectric and optical properties of dense plasmas," *Annales de Physique*, vol. 30, no. 4-5, pp. 1–187, 2005.