High-Performance Computing Challenges in Space Physics Simulations

Nikolai V. Pogorelov
Department of Space Science and Center for Space Plasma and Aeronomic Research,
University of Alabama in Huntsville

The subject of space physics covers a wide range of topics from the solar interior, solar wind (SW) acceleration near the Sun, Sun–Earth connection, planetary magnetospheres, to the SW interaction with the magnetized local interstellar medium (LISM). Physical problems involved include turbulence, kinetic processes in the collisionless SW plasma, ion acceleration, magnetic reconnection, various space weather topics, non-thermal ion behavior, charge-exchange effects, ionosphere-magnetosphere coupling, the physics of coronal mass ejections, galactic cosmic ray transport throughout the heliosphere, the origin and behavior of anomalous cosmic rays, and many others. The range of problems stipulates the necessity to address different space and time scales and use different sets of equations, e.g., MHD, Boltzmann, Fokker–Planck, Parker, focused-transport, and other equations, which very often should be solved together, in a coupled fashion. Numerical methods used to solve those equations are also different and can be based on both deterministic and stochastic approaches. Coupling of different scales frequently requires overlapping grids, which is impossible without adaptive mesh refinement (AMR) on the largest scale. An increase in the AMR level creates challenges for data exchange between computing nodes. Coupling of different sets of equations requires more sophisticated algorithms that take into account computer architectures, e.g., the need in hybrid parallelization strategies.

Although kinetic and hybrid simulations provide invaluable insight into the microscopic nature of such physical phenomena as, e.g., magnetic reconnection, scaling of the obtained results to macroscopic scales may be very difficult, which is especially well seen in the attempts to explain Voyager 1 observations at the SW–LISM boundary. It is interesting that coupling between microscopic and macroscopic scales occurs through the turbulence intrinsic to collisionless space plasma. The mere possibility to apply the equations of continuum mechanics (gas dynamics or MHD equations) is due to a considerable decrease in the ion mean free path due to their scattering on the magnetic field fluctuations. The presence of turbulence undermines the frozen-in conditions for magnetic field lines and therefore affects magnetic reconnection at a macroscopic level. To couple the scales in global simulations one needs to identify the spectrum of turbulent fluctuations in the regions of possible reconnection and instability. This can be done by performing hybrid (ions are treated as particles while electrons behave as the neutralizing fluid) simulations on a grid overlapping with the global grid (see Fig. 1). This is feasible with AMR. On the other hand, hybrid simulations may overlap in scale with full particle simulations that take into account electron kinetic effects. A hierarchy of overlapping meshes makes it possible to couple macro- and micro-scales in many problems. Available computing power does not allow us to cover all scales self-consistently, but the setup for each lower-scale problem may be defined by a one-level-higher-scale problem solution. Additionally, the knowledge of the properties of MHD turbulence determined locally in problem-specific regions makes it possible to describe ion acceleration and GCR transport using realistic transport coefficients, and relevant plasma and magnetic field background. Below is a brief list of our research needs.

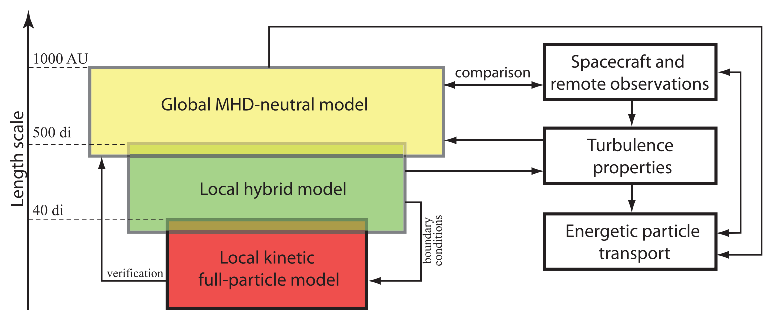

Figure 1: The hierarchy of scales and their dependence as described in my Heliospheric Grand Challenges project supported by NASA for 2014–2016 (d_i is the ion skin depth).

1. Global AMR simulations based on MHD equations coupled with the Boltzmann and Fokker–Planck equations are necessary to model the SW–LISM interaction and a variety of related fundamental physical phenomena (ion acceleration, MHD and kinetic instabilities, magnetic reconnection, etc., intrinsic to most problems involving collisionless space plasma and essential for space weather prediction as one of the NSF priority research subjects) measured in situ by Voyagers and remotely by IBEX. Such simulations belong to the capability computing area. They require efficient system-level support for parallel code-to-code communication, so that two different parallel codes could run simultaneously and be able to communicate with each other. So far we have been dealing with this problem taking advantage of the fact that the kinetic part requires considerably greater time than its MHD counterpart. In this case the parallelization deficiencies are not visible and it suffices only to arrange an optimum distribution of the MHD data among the cores solving the Boltzmann equation. This was done using a mixed MPI/OpenMP parallelization strategy and special data buffering.

2. Substantial computing resources are necessary to perform global, 3D, hybrid simulations of the SW plasma, especially in the Earth magnetosphere. While some results on such modeling have been reported recently, most of hybrid simulations are either local or 2D. These problems are of major importance for understanding the SW interaction with planets, including Earth. There are attempts to embed hybrid simulations into AMR MHD codes, which saves computer resources (see, however, the previous paragraph).

3. Full-particle (particle-in-cell, PIC) simulations are necessarily limited to very small regions of the order of a few tens of ion gyroradii. These simulations are among the most computationally demanding calculations that are run on the NSF and DOE supercomputers. PIC codes are easily parallelized, but are data and computationally expensive. They are important for phenomena dependent on the electron dynamics. An implicit approach has been proposed recently for PIC simulations, which however essentially smears electron scales.

4. Solar Probe Plus (SPP) is a new NASA mission with the launch scheduled in 2018. Among other high-resolution data, this mission will perform time-dependent measurements of the 3D, ion, electron and α-particle distribution functions. On the one hand, this will result in an enormous amount of data, considering the 7-dimensional character of the distribution. On the other hand, this will require the development of new, efficient numerical methods to interpret the observations. This means the necessity of developing 6-dimensional parallelization strategies and AMR both in the physical and velocity . The SPP mission will give rise to a data-intensive simulations in space physics.

5. To make computational efforts more efficient, the NSF should prioritize funding research efforts that focus on new intelligent numerical algorithms, i.e., asynchronous time integration, AMR, load-balancing, etc.

6. Efficient, parallel visualization software for analyzing multi-dimensional numerical results should be supported and made accessible on every supercomputer.

7. An area that is poorly covered by existing systems is data storage and sharing. Most large-scale computational centers do have adequate storage facilities, but none of them provides reasonable resources for collaborative data analysis and data sharing, while many large-scale simulations generate data of interest to many research groups.

High-performance resources for the space physics community are available on NASA, NSF, and DOE supercomputers. However, computing power available for scientific simulation on NASA supercomputers is rather limited. On the other hand, the DOE INCITE program imposes requirement on the code scalability so strong that getting allocations on Titan or Mira is very difficult for most of users. In my opinion, NSF allocations are essential for the majority of simulation-related NSF projects. The demand in computing resources will continue exceeding available capabilities. The progress on the hardware side greatly outpaces the progress in the algorithms and software development, which is underfunded across the agencies. It is desirable for the NSF to support the development of a few high-quality community codes and for the space physics community to take advantage of the already developed software.