TN

hPIC: A SCALABLE ELECTROSTATIC PARTICLE-IN-CELL FOR PLASMA–MATERIAL INTERACTIONS

Allocation: Exploratory/50 Knh
PI: Davide Curreli[1]
Co-PI: Rinat Khaziev[1]
Collaborators: Cameron Dart[1], Alyssa Hayes[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The hPIC code is a kinetic–kinetic electrostatic particle-in-cell (PIC) application targeted at large-scale simulations of plasma–material interactions (PMI). hPIC models the plasma sheath and presheath in strongly magnetized conditions, as normally encountered at the wall of magnetic fusion reactors. The Poisson solver is based on the PETSc conjugate gradient with BoomerAMG algebraic multigrid preconditioners. Scaling tests on the Blue Waters supercomputer have demonstrated excellent weak scaling and considerable strong scaling.

RESEARCH CHALLENGE

In the edge region of magnetically confined plasmas, plasma–material interactions can compromise the thermos–mechanical integrity of the wall, currently limiting the development of nuclear fusion reactors. When exposed to a fusion plasma environment, plasma-facing materials exhibit evidence of surface morphology modifications, nanostructuring, and neutron damage. The boundary layer between the edge plasma and the material surface is called the plasma sheath, a region where the plasma ions accelerate up to supersonic conditions in order to maintain the quasi-neutrality of the plasma. In the strong magnetic field of a fusion machine (of the order of several Tesla), the plasma sheath is a complex three-dimensional multi-layer structure, where substantial electrostatic forces accelerate the ions toward the wall, causing damage. The hPIC code is a large-scale particle-in-cell under development at the University of Illinois and is able to perform detailed analyses of the plasma sheath in the strongly-magnetized conditions encountered in a fusion reactor. We report progress on the scaling tests (weak and strong scaling) and code optimization performed on Blue Waters. All scaling plots have been reported as a function of the number of Blue Waters' XE6 nodes; one Blue Waters node corresponds to 32 cores.

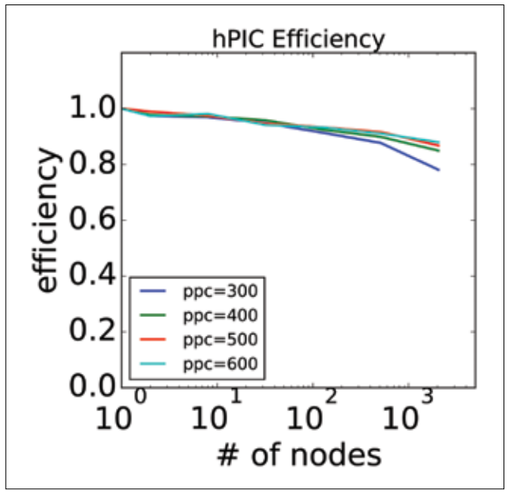

Figure 1: Weak scaling efficiency of the electrostatic version of the full cycle of hPIC for a hydrogen plasma with ion and electron temperature equal to 1eV and electron density equal to 10^{17} m^{-3}; each rank stores 100×100 grid points, with a number of particles per cell (ppc) as indicated in the figure.

METHODS & CODES

As an electrostatic particle-in-cell, hPIC requires the solution of a large-scale Poisson problem at each time step. A well-known problem of elliptic solvers is the intrinsic nonlocality of the differential operator, hindering an efficient and scalable implementation on a large-scale computer. hPIC uses a finite-difference discretization of the Laplace operator solved with conjugate gradient and an algebraic multi-grid (BoomerAMG) as a preconditioner. The PETSc library was used as an interface for the parallel implementation. The PETSc numerical library manages the field functions of the code. The interface between hPIC and PETSc also ingests the charge density inputs required by the Poisson solver and returns electric field and electric potential.

RESULTS & IMPACT

The largest weak scaling test performed was a Poisson problem on 2,048 Blue Waters nodes (65,536 cores). A 500×500 grid ran on each core, resulting in 16.384 billion degrees of freedom. Fig. 1 shows the efficiency (w.r.t. single node) of the full PIC cycle, including both particle and field functions, for different numbers of particles-per-cell (ppc). A weak scaling efficiency greater than 78% and between 85% and 88% for all conditions relevant for practical applications has been attained.

The tests have shown that as a rule of thumb for good performance, the number of particles per cell should be kept larger than 100, with preference to 500–600 particles-per-cell. With a small local grid of 50×50 on each MPI process, hPIC is capable of achieving up to 3,000 particles per cell, enough to achieve a satisfactory level of statistical noise. Currently, hPIC uses double precision for all physical quantities on the XE6 nodes of Blue Waters containing 64 gigabytes of random access memory per node. Thus, each node stores approximately 250 million particles.

Similarly, strong scaling tests were performed on 4 to 8,192 nodes. The problem was a Poisson equation with a spatial grid of $N \times N$ points and a corresponding Laplacian matrix of size $(NN \times NN)^2 = NN^4$. At the largest scale of 8,192 nodes, the time to solution for the two allocatable cases of NN=16,000 and NN=32,000 was 1.6 seconds. The greatest problem of NN=100,000 was solved on 256, 512, and 2,048 Blue Water nodes. This solution required PETSc's index to switch from a 4 byte signed integer to an 8 byte signed long long. The problem of size NN=100,000 has a grid size of $NN^2 = 10^{10}$ nodes, and a corresponding Laplacian matrix of $NN^4 = 10^{20}$ degrees of freedom with matrix indices running from 0 to10^{10}-1, which can be allocated only by using 64-bit unsigned long integers. Such a problem was solved in 13.4, 6.8, and 2.1 seconds respectively on 256, 512, and 2,048 Blue Waters nodes.

WHY BLUE WATERS

Access to Blue Waters has accelerated the development of hPIC, allowing for testing and optimization of a number of massively parallel features of the code. The code is now at the level of meeting the large computational demands of fully kinetic simulations of near-wall plasmas encountered at the edge of fusion reactors, simulations feasible only on a machine the size of Blue Waters. Resolving the plasma–material interface at scales from nano to millimeters and from picoseconds to the time scales of plasma-sheath equilibrium requires computing at the petascale and beyond. Blue Waters offers a unique opportunity to push this research forward.