A White Paper for

Workshop on Turbulent Flow Simulation at the Exascale: Opportunities and Challenges

August 4-5, 2015, Washington DC

Code Name: XGC1
Application Area: Tokamak Magnetic Fusion
Supporting DOE Application Office: Office of Fusion Energy Science (OFES)

XGC1 is a multiscale turbulence code for tokamak fusion plasma, being supported by Office of Fusion Energy Science, US Department of Energy. Plasma is an ionized gas at high temperature, with the force dominated by classical electromagnetic interactions. The kinetic dynamics of the charged particles in a strong and inhomogeneous toroidal magnetic field is an important driver of the plasma turbulence that leads to deterioration of the confinement. Thus, differently from the traditional fluid turbulence codes, XGC1 does not solve the 3D Navier-Stokes equation but the more fundamental gyro-kinetic Boltzmann equation in 5D space (3D in configuration space and 2D in velocity space) in strong magnetic field. The third dimensional velocity variable, which is the gyro-angle variable, has been analytically removed by restricting the electromagnetic fluctuations to be below the fast gyro-frequency, where the confinement-relevant turbulence behaviors exist in steady fusion reactor conditions.

Unlike most of the existing fusion gyrokinetic codes --which assume scale separation between the background plasma profile and turbulent perturbation, and calculate the perturbed distribution function and kinetic turbulence driven by a fixed fluid background gradient (perturbed delta-f method)-- XGC1 calculates the full kinetic plasma all together (total-f method). Thus, XGC1 obtains self-organized total solution between the background kinetic plasma and turbulence. The conventional delta-f method was designed to be useful on small-scale computers in 1980s and is still predominant in the fusion turbulence simulations on capacity computing platforms. Since XGC1 solves for the total kinetic distribution function (total-f), instead of the small perturbed part of the distribution function in the delta-f method, it requires much more computing power than the conventional delta-f method does. The more powerful the computer becomes, the more physics XGC1 can contain.

Another important feature, which distinguishes XGC1 from the rest of the fusion gyrokinetic codes and which requires the most extreme scale leadership class computers, is that XGC1 includes the edge region of tokamak across the magnetic separatrix to the material wall. It has been well-known experimentally that the fusion quality in the burning core is determined by the plasma confinement at the edge that is governed by wide range of scale inseparable non-thermal multi-scale physics, including the neutral particle recycling at the material wall in complicated edge geometry.

Out of these necessities, XGC1 has been constructed from the beginning as a modern extreme scale code in the past SciDAC-2 program titled "Prototype Fusion Simulation Project: Center for Plasma Edge Simulation (CPES)," in close collaboration with applied mathematicians, data management scientists, and code performance engineering scientists. Further development and performance engineering of XGC1 is continuing in the SciDAC-3 progrom: Partnership Center for Edge Plasma Simulation (EPSI). XGC1 scales excellently in both weak and strong sense to the maximal capability of the heterogeneous Titan (~27 PF peak, load-sharing between GPUs and CPUs), homogeneous Mira (~10 PF peak), heterogeneous Blue Waters, and homogeneous Edison, with good portability.

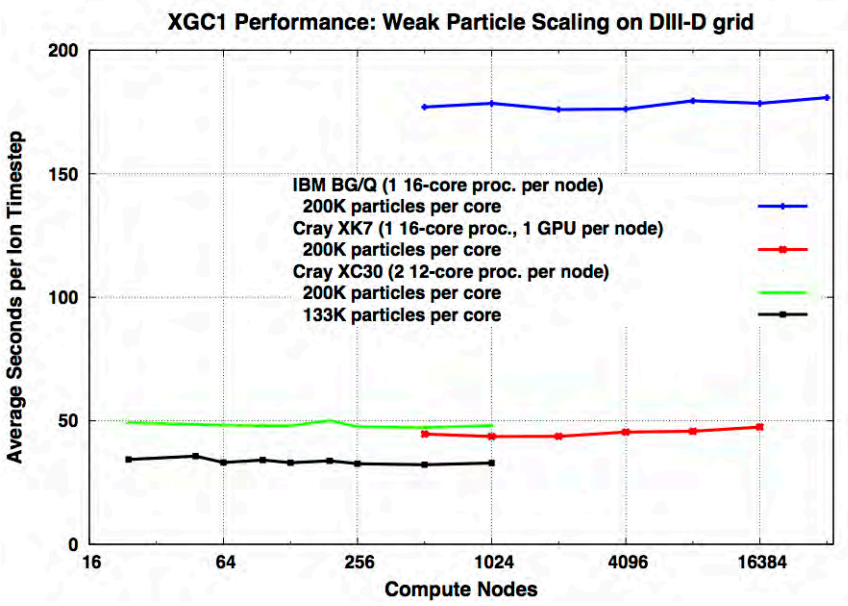

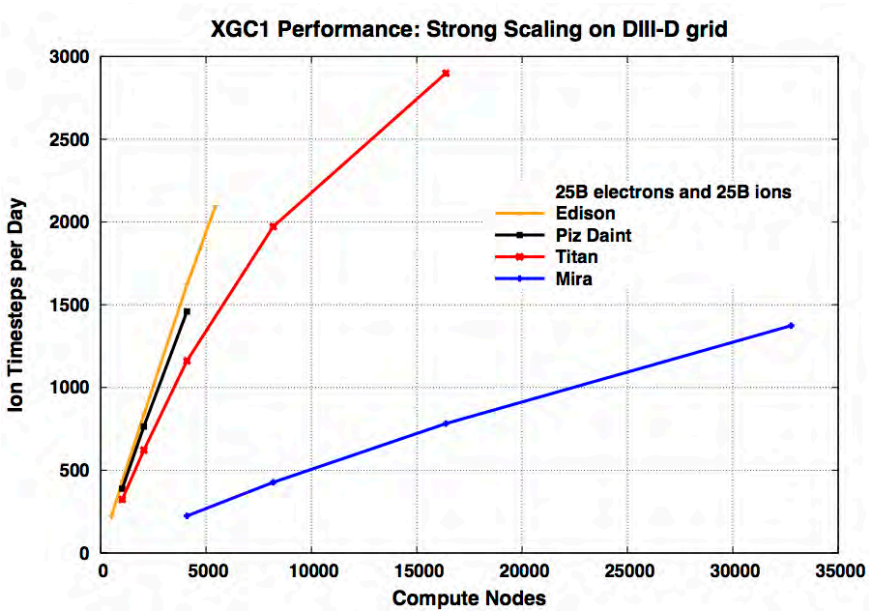

Figure 1. Weak and strong scaling of XGC1 on the US and EU open-science leadership class computers. Neither weak nor strong scaling show degradation all the way the maximal number of nodes, except for the strong scaling case on heterogeneous Titan where the reduction of the GPU workload causes the GPU-CPU communication overhead to surface up at high node count.

XGC1 uses the particle-in-cell technology on unstructured triangular grid. The grid vertices follow the equilibrium magnetic field lines, approximately. It also combines the particle-in-cell technology with a creative 5D continuum grid scheme to enable nonlinear collisions, sources, sinks, reduction of particle noise, easier physics analysis, and reduction of output data. Finite difference and finite volume methods are used on the continuum operations. For portability, XGC1 does not use commercial libraries. XGC1 is an open-source code and uses only open-source libraries.

On the world #2 HPC Titan, XGC1 is able to perform heroic simulations of ion gyroradius-scale turbulence with ~1mm grid resolution in the whole ITER volume. One of these heroic simulations occupies about 5 days of the entire Titan CPU cores and GPUs to achieve an approximate turbulence saturation, with the background kinetics and the neutral particle recycling included in diverted edge geometry. On the upcoming pre-exascale computers, one of these simulations will become a routine 10 hour job, and on the exscale computers, it will be a routine ~1 hour job if the scalability can be extended as desired. These simulations, however, cannot be called a true whole-device simulation that contains all the important multiscale ITER physics since the ion-scale resolution does not cover the electron scale turbulence. The electron scale turbulence can resides in ~1/10 of the plasma volume requiring 2D grid resolution of ~0.1mm, hence demanding ~10X compute power than the ion-scale turbulence alone does. This means that an exa-scale computer can enable a whole-device all-scale physics simulation of fusion plasma from magnetic axis to material wall in one-day wall-clock time. The 6D physics that cannot be described by the 5D gyrokinetic simulations can be imbedded into specific part of 5D space as needed, or coupled into it using special-purpose codes (such as the rf wave propagation codes) on an exacale computer.

Such a high fidelity understanding and prediction of the whole-device burning plasma behavior can yield more economical and accelerated achievement of ITER's goal, and help realize commercial reactors sooner.

There are challenges to be resolved in order to continue the present efficient scalability to exascale computers: i) Efficient implicit algorithm to solve the fast electron time advance at minimized overhead, ii) Efficient usage of the deep memory structure, iii) Clever usage of NVRAMs for in-situ data analysis and check-point restart in order to alleviate the big data issue, iv) Make XGC1 to be resilient to fault, v) Efficient 3D grid-particle domain decomposition, and others.