QUANTUM ESPRESSO ON GPU ACCELERATED SYSTEMS

Massimiliano Fatica

NVIDIA, USA

E-mail: mfatica@nvidia.com

Quantum Espresso (QE) is a popular integrated suite of open source software for first principle simulations of materials, simulations which frequently require a large amount of computational resources.
GPU accelerated systems have emerged as an attractive alternative to conventional CPU systems for high performance computing.
They have proved to be very effective, in regards to both power efficiency (a critical aspect for exascale-class machines) and computational throughput, at accelerating compute intensive and memory bandwidth intensive applications.
This talk will present the porting effort of the Plane-Wave Self-Consistent Field (PWscf) solver, one of the main workhorse solvers in the QE suite, to GPU accelerated systems using CUDA Fortran.
Performance of the GPU-accelerated code on real-world benchmark cases versus a baseline CPU system using Intel Xeon processors systems will be reported, including the effect of GPU node topology (CPU to GPU ratio, availability of fast P2P connections via NVLink, x86 and POWER host CPUs).
A performance comparison with QE on Intel KNL systems and of our code versus SIRIUS, a port of this program by other authors using a mixed source approach, will also be provided.