Post-doctoral position at ECR Lab. (Intel, CEA, UVSQ)

Campus Teratec, Bruyere-le-Chatel, France

Title: Post-doctoral position on accelerating Quantum chemistry HPC simulation

Contacts :

Eric.Petit@intel.com

Marie-Christine.Sawley@intel.com

Keywords:

ABINIT, MPI, OpenMP, GASPI, TBB, parallelism, runtime, algorithm, programming model.

Context:

The proposed post-doctoral position will take place in the Exascale Computing Research Lab in Terratec, Bruyere-lel-Chatel, France. This structure is a collaboration between three players: CEA DAM, Intel, and the University of Versailles (UVSQ).

CEA, the French Atomic Energy Agency (Commissariat à l'Energie Atomique et aux Energies Alternatives), performs research, development and innovation in four main fields: energy, fundamental research, information and health technologies, and defense and national security.

Intel is a technology manufacturing firm providing silicon and software enabling technologies to various industry and market segments worldwide. Intel is developing a new architecture, Xeon Phi and is willing to engage into collaboration with leadership class application developers to demonstrate the potential of the Xeon Phi for exascale computing.

UVSQ, a leading French University, is involved in numerous fields and has expertise on research on software tools to optimize processor and system performance, and works on new optimization models and tools for many core environments for exascale computing.

Scientific background:

ABINIT is a software for atomistic modelling of the properties of solids and nanostructures from first-principles theory. It is an international collaborative project and the CEA is one of its main developer.

ABINIT is used by CEA for many applications in material science related to electronic properties: determination of Equations of States (EoS), computation of thermo-dynamical properties, study of defects (structural point defects, surfaces), computation of response of material to solicitations (elastic or vibrational properties, response to electric filed…)
Most of these applications involve simulations using the so-called Quantum Molecular Dynamics technique: the atomic nuclei trajectories are simulated using forces based on the electronic structure. The characteristic times of such simulations (physical simulated times) are of the range of a few picoseconds.

Such time scales are not necessarily sufficient to access to rare events like, for instance, melting activation, motion of an interstitial atom, energetic barrier crossing, etc.
The main goal of our project is to access to longer simulation trajectories; to reach this goal, the objective is simple: decrease the time needed to determine the forces on nuclei at each time step (forces are derivatives of the energy). For this, we need to drastically reduce the *time to solution*.

The energy of a system made of electrons depends on the "electronic wave function". The latter is a vector obtained from the resolution of a diagonalization problem. Unfortunately, the operator (matrix) to be diagonalized is not known (too large to be stored), only its application to a vector.

The position will require strong knowledge in computer science, especially in the field of parallel programing, parallel algorithms, parallel runtimes, and performance tools. Publications are expected both in computer science and quantum chemistry

Research and development objectives

The efforts will be concentrated on the following tasks:

1. Finalize an abstraction layer in ABINIT, allowing modifications of the low-level routines of the code without changing the coding of the physics. This abstraction layer should have, in principle, no impact on performances, but will facilitate the adaptation of algorithm on new architecture.
2. Speed-up the application of the operator to be diagonalized, especially on Xeon-phi KNL architecture. This involves exploring new diagonalization algorithms and efficient shared memory parallism (e.g. OpenMP) using highly scalable paradigms.
3. Scalable FFTs. Several collective communications (*alltoall*) are performed during the computation to transfer the vectors from one representation to another. At large scale, these communications can be improved using asynchronous one sided scheme, task programming and shared memory.